THE RING AMPLIFIER: SCALABLE AMPLIFICATION WITH RING OSCILLATORS

Benjamin Hershberg[1] and Un-Ku Moon[2]

[1]Imec, Leuven, Belgium
[2]Oregon State University, Corvallis, USA

Abstract

Ring amplification is a technique for performing efficient amplification in nanoscale CMOS technologies. By using a cascade of dynamically stabilized inverter stages to perform accurate amplification, ring amplifiers are able to leverage the key benefits of technology scaling, resulting in excellent efficiency and performance. A generalized view of basic small-signal theory is first presented, followed by a deeper discussion of the time-domain operation of a ringamp in the context of a specific ringamp structure. We conclude with a survey of existing ringamp implementations and techniques reported in literature.

1. Introduction

In ecology, bio-diversity is often a key indicator of the fitness and resilience of an ecosystem. In a similar way, the diversity of viable solutions and approaches available in the world of analog circuits is an indicator of the health of the analog design ecosystem. The range of technical requirements for analog signal processing blocks in practical design applications is as broad and vast as the ways in which their commercial implementations are used to enhance the many facets of society, work, and leisure. There is no one-size-fits-all solution here. Rather, a variety of solutions allow us to select the best tool for the job and leverage technology scaling to the fullest.

For design in nanoscale CMOS, amplifiers are one area where diversity has arguably been lost. Only a small sub-set of the viable amplifier topologies that once ruled in micron and submicron CMOS design remain competitive in nanoscale CMOS. The impact of this is readily observed in the study of ADCs, where amplifier-less SAR ADC topologies have been able to leverage technology scaling to consistently achieve conversion efficiencies far surpassing ADC topologies that rely heavily on amplification [1, 2].

A common misconception is that this poor scaling performance is a fundamental flaw of amplification itself. In actuality, it is mainly due to an adherence to the old paradigms of how amplifiers should be built. Conventional opamp topologies were conceived of at a time when 2.5V supplies were considered low-voltage, and the intrinsic properties of transistors were quite different from that of a 14nm FinFET [3]. Transistors were physically much larger with larger internal parasitic capacitances often approaching even that of the load capacitance being charged. Supply voltages provided much more headroom for stacking devices, and intrinsic device gains were higher. All these factors created a design world where cascoding was preferable to cascading, current biasing was necessary, and small-signal analysis was the prevailing and sufficient design paradigm.

Since those times, a lot has changed. Semiconductor technology has continued to evolve in a direction aimed at improving density, efficiency, and speed for digital logic gates. The ability of a conventional opamp topology to flourish in this new environment is fundamentally limited due to inherent incompatibilities in the underlying approach. Applying additional techniques such as calibration, gain-enhancement, and output-swing enhancement may enable an opamp to function in nanoscale environments, but it won't grant it the ability to scale at the same pace as digital performance improvements. A truly scalable amplifier must operate natively in its environment, in a way that implicitly uses the characteristics of scaled CMOS to its advantage, transforming potential weaknesses into inherent strengths. Since technology scaling is deliberately designed to favor the time-domain world of high-speed digital, viable nanoscale analog techniques are likely to be found in the time-domain realm as well. In order to fully exploit the abilities of a transistor, the biasing and small-signal properties of the device must be viewed as highly coupled, time-dependent variables which can be applied as feedback to each other with respect to time.

Here we explore one such technique: ring amplification. A ring amplifier (ringamp, RAMP) is a small modular amplifier derived from a ring oscillator which naturally embodies all the essential elements of scalability. It can amplify with rail-to-rail output swing, efficiently charge large capacitive loads using slew-based charging, scale well in performance according to process trends, and is simple enough to be quickly constructed from only a handful of inverters, capacitors, and switches.

2. A small-signal, steady-state perspective

The transient, large-signal, and small-signal operation domains are often much more interdependent in ring amplifier toplogies than classical amplifier topolo-

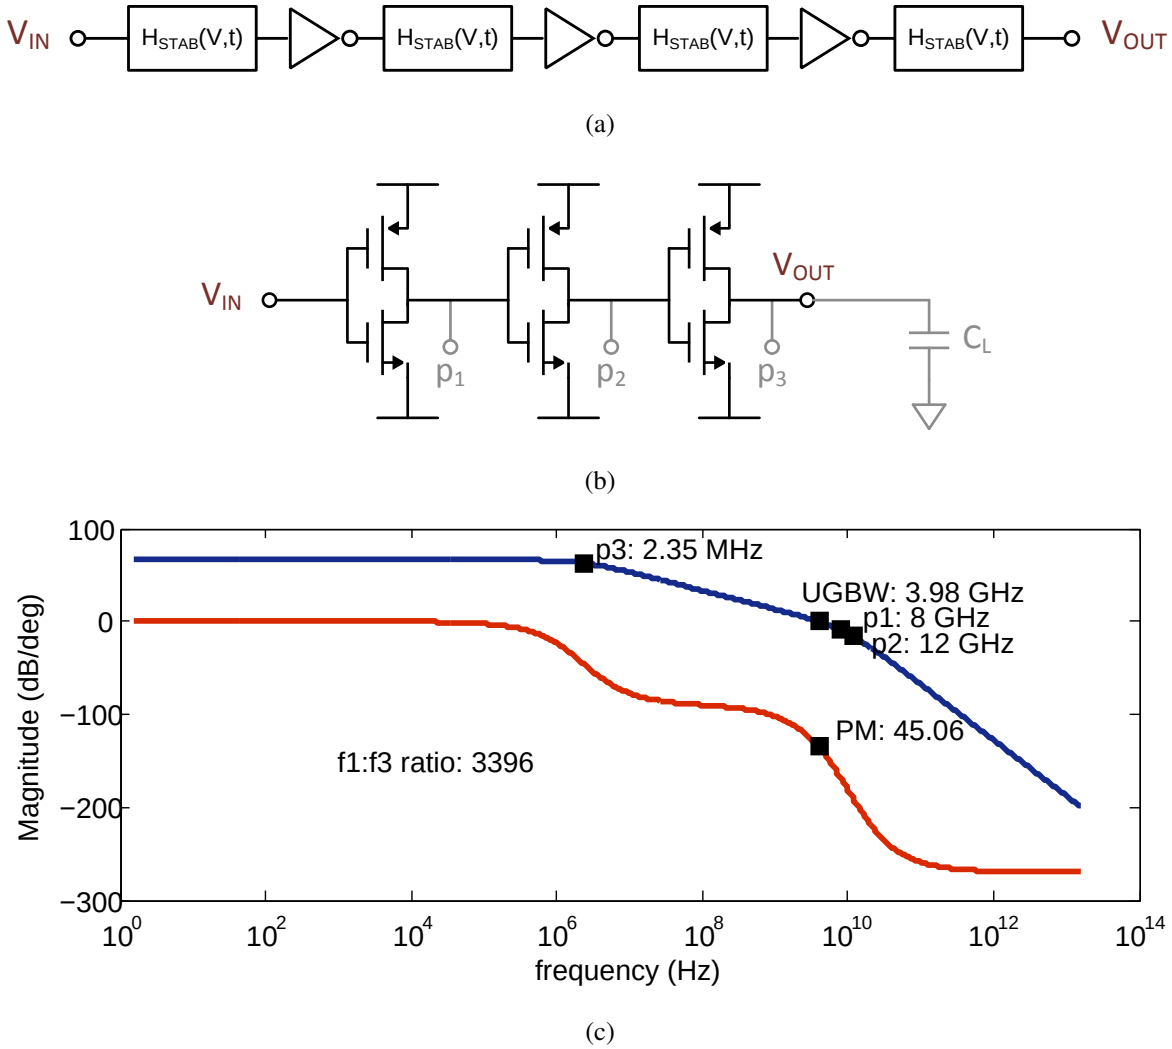

Fig. 1: A generalized three-stage ringamp in (a) is shown for the specific case of a three-inverter ring oscillator in (b). An example frequency response of (b) is given in (c).

gies. Despite this added complexity, as a starting point we can begin with the steady-state small-signal analysis of classical amplifier analysis.

Simply put, a ringamp is a multi-stage amplifier stabilized by a dominant output pole. A generalized three stage ringamp is shown in Fig. 1a, consisting of three inverting gain stages and optional stabilization networks. For now, we will consider the simple case where the stabilization networks have a steady-state, frequency-independent gain of 1 and each gain stage is a single-pole system. This results in Fig. 1b: a cascade of three inverting stages.

The open-loop gain of this three pole system is given by

$$H(s) = \frac{g_{m1}r_{o1} \cdot g_{m2}r_{o2} \cdot g_{m3}r_{o3}}{(1 + sr_{o1}C_{p1})(1 + sr_{o2}C_{p2})(1 + sr_{o3}(C_{p3} + C_L))} \quad (1)$$

where r_{ox} is the impedance seen at pole/node X, g_{mx} is the trans-conductance of inverter stage X, and C_{px} is the total capacitance seen at pole/node X.

Recalling basic stabilization theory, in order to transform this structure from an unstable ring oscillator into a stable ring amplifier, we must create a sufficiently large ratio between the lowest frequency pole in the system and the higher frequency poles. Given that an external load capacitance is required for any practical switched-capacitor design scenario, it is then relatively simple to prove that the strategy of stabilizing with a dominant output pole (p_3) will always yield the maximum amplifier bandwidth and highest efficiency.

The optimal output pole location can be created by first placing p_1 and p_2 at the highest frequencies possible and then adjusting the location of p_3 until it is at a sufficiently low, stabilizing frequency. Critically, for a three-stage opamp using conventional techniques and current biasing, this would not yield a very practical solution. The internal poles would still be relatively large, which would limit bandwidth and require an often impractically large explicit load capacitance at the output. Miller-compensation can be used to make other poles in the system dominant instead, but at high price in terms of bandwidth and efficiency. In the case of a ringamp with very small transistors used in gain stages 1 and 2 and dynamic biasing (i.e. just a basic inverter in this case), output pole stabilization becomes a realistic possibility. Poles p_1 and p_2 can be placed at very high frequencies, which allows us to place p_3 at a sufficiently stable location using a much more reasonably sized load capacitance.

This analysis describes the steady-state condition that a ringamp must reach in order to stabilize, but it does not tell us anything about *how* it does it. In many application scenarios, Fig. 1b is not the optimal implementation, and may not even be capable of meeting the required performance specs. At this point it is worth revisiting Fig. 1a and re-considering the "*how*" of ringamp stabilization. In the next section, we will take a closer look at a particular ringamp structure that uses time-domain feedback to dynamically adjust the output pole location and thereby relax stability constraints. Afterwards, in Section 4, we will look at other existing solutions, including techniques to extend ringamp operation into high accuracy applications.

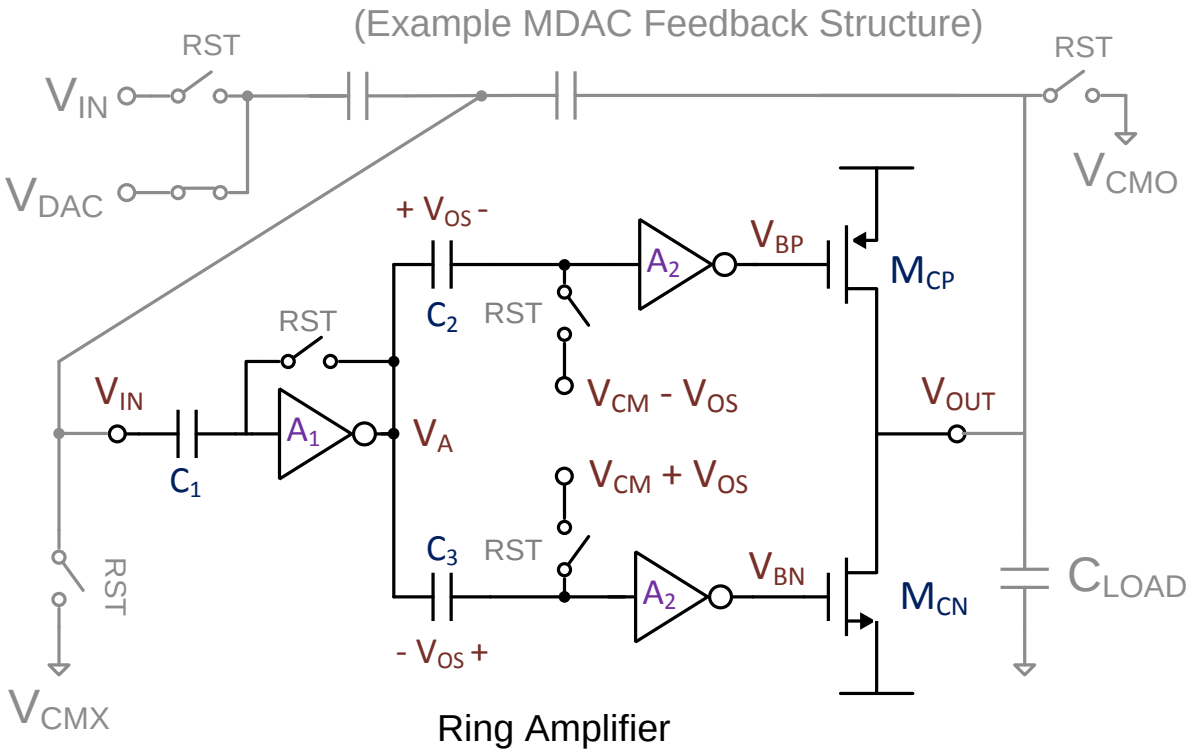

Fig. 2: The ringamp and basic switched-capacitor feedback network that we will primarily consider in this section. Devices and parameters that are referenced throughout the paper are labeled.

3. Utilizing time-domain feedback

We will now consider the specific case of the ringamp topology shown in Fig. 2 that utilizes time-domain feedback to enhance stabilization, and thereby improve efficiency. Fundamentally, the ringamp of Fig. 2 is a ring oscillator that has been split into two separate signal paths. A different offset is embedded into each signal path in order to create a range of input values for which neither output transistor M_{CN} nor M_{CP} of Fig. 2 will fully conduct. If this non-conduction "dead-zone" is sufficiently large, the ring amplifier will operate by slewing-to, stabilizing, and then locking into the dead-zone region. When placed in the example switched capacitor MDAC feedback structure also shown in Fig. 2, this charging and settling behavior results in the waveforms of Fig. 3b.

Before we examine how and why this occurs, it is useful to first understand some of the basic characteristics of the structure itself. To begin with, consider the capacitor C_1 of Fig. 2. C_1 is used to cancel the difference between the MDAC virtual-node sampling reference (V_{CMX}) and the trip-point of the first stage inverter. This ensures that the ideal settled value for V_{IN} will always be V_{CMX}, independent of the actual inverter threshold. Any sources of offset that are gener-

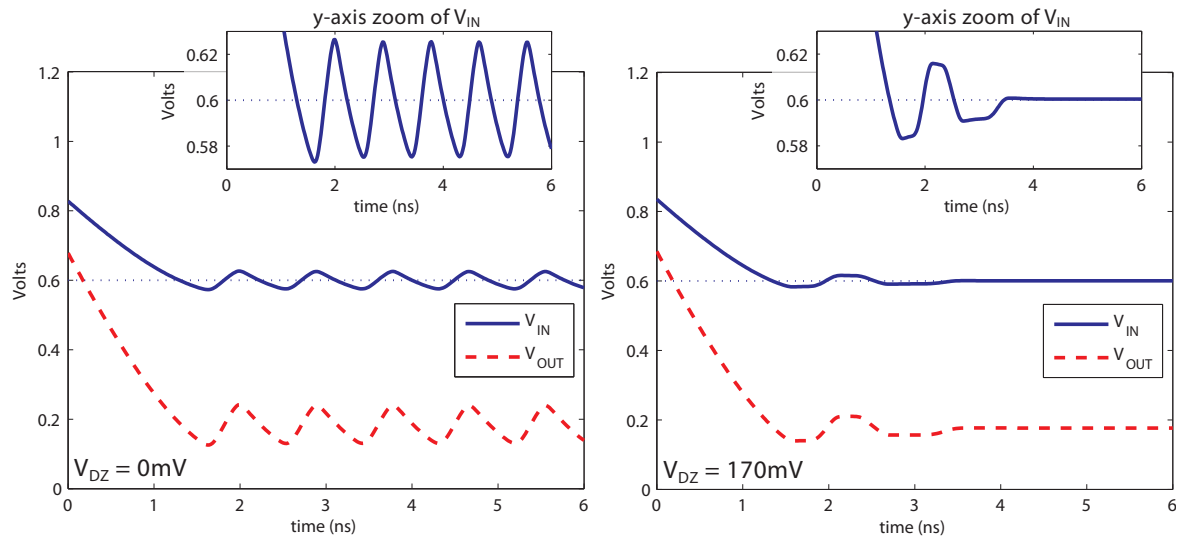

Fig. 3: Input and output charging waveforms of Fig. 2. In (a), when V_{DZ} = 0mV, the ringamp is functionally identical to a three-inverter ring oscillator. In (b), the dead-zone is set large enough to generate stability (V_{DZ} = 170mV) and the ringamp functions as an amplifier.

ated after the first stage inverter will not be removed by C_1, but the input-referred value of such offsets will typically be negligibly small.

The dead-zone of the ringamp in Fig. 2 is embedded prior to the second stage inverters by storing a voltage offset across capacitors C_2 and C_3. Any value for V_{IN} within the dead-zone region is a viable steady-state solution for the ring amplifier, and the input-referred value of the dead-zone will determine the overall accuracy of the amplifier for most practical cases. In other words, the error at V_{IN} when the ringamp has stabilized and locked will be

$$-\left|\frac{V_{DZ}}{2{\cdot}A_1}\right| \leq \epsilon_{V_{IN}} \leq \left|\frac{V_{DZ}}{2{\cdot}A_1}\right| \tag{2}$$

where $V_{DZ} = 2V_{OS}$, A_1 is the final settled small-signal gain of the first stage inverter, and finite gain effects of the latter stages are ignored (revisited later).

Considering all this, we will now examine the interplay of small-signal, large-signal, and time-domain operation in a ringamp. Its behavior can be subdivided with respect to different modes of operation in time: slewing, stabilization, and steady-state. To illustrate key concepts, we will use the exaggerated charging waveform of Fig. 4 (taken from the ring amplifier of Fig. 2) that has been designed with relatively low bandwidth, excessive drive current, and a dead-zone size that biases the ringamp right at the edge of stability. Although one would never wish to make a real design in this way, as a teaching example it is quite useful. V_{CMX}

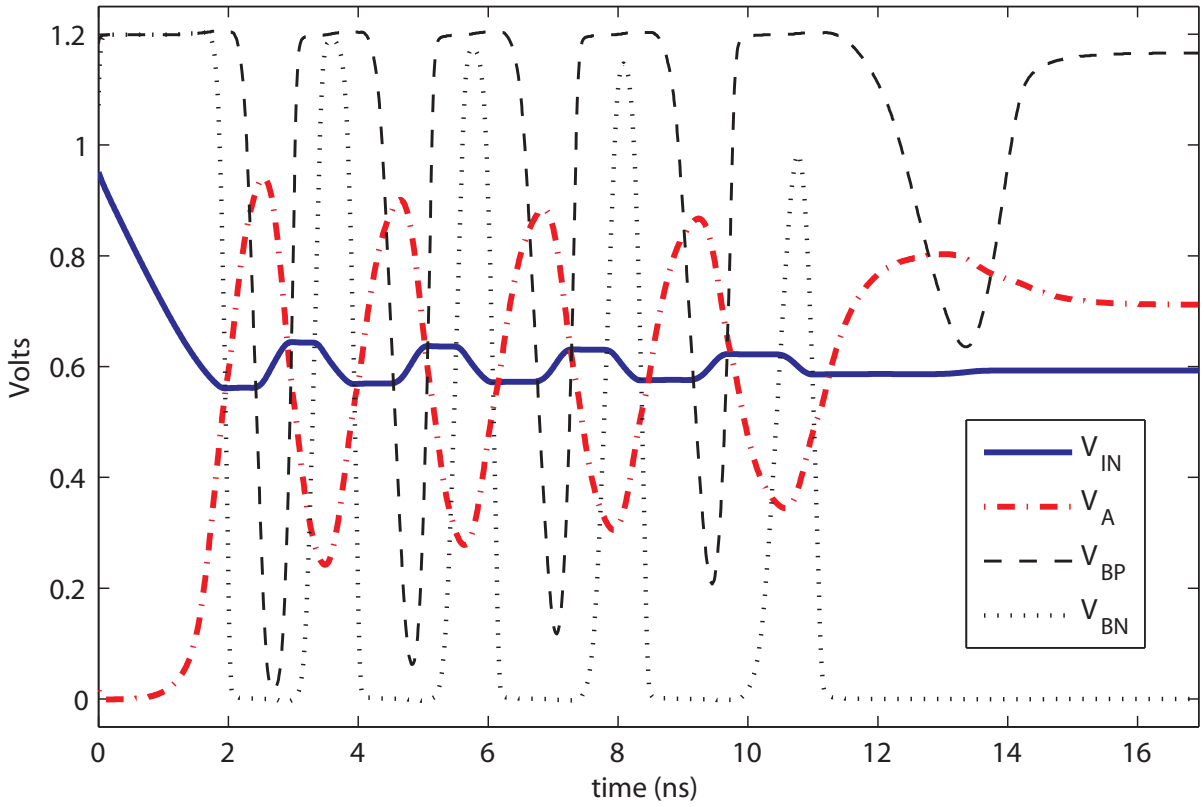

Fig. 4: Example ring amplifier operation for an exaggerated design biased at the edge of stability, showing the three key phases of operation: 1) initial ramping, 2) stabilization, and 3) steady-state.

is set to 0.6V, and thus the ideal settled value of V_{IN} will also be 0.6V. For the sake of simplicity and generality V_{OUT} is not shown (because it is simply a scaled, shifted, signal-dependent replica of V_{IN}). Unless otherwise stated, any mention of the amplitude of the fed-back signal will refer to the amplitude seen at V_{IN}.

In Fig. 4 we can clearly see three main phases of operation. Initially, from 0ns to 2ns, the ringamp rapidly charges toward the dead-zone. Then, from 2ns to about 14ns it oscillates around the dead-zone region as it attempts to stabilize. By 15ns, with the output transistors M_{CP} and M_{CN} both completely cutoff, the ring amplifier reaches a steady-state solution within the dead-zone, and remains locked.

Initial Ramping

In the initial slew-charging phase of operation, the ring amplifier is functionally equivalent to the circuit of Fig. 5. The first two stages of the ring amplifier act like a pair of bi-directional continuous-time comparators that correctly select which

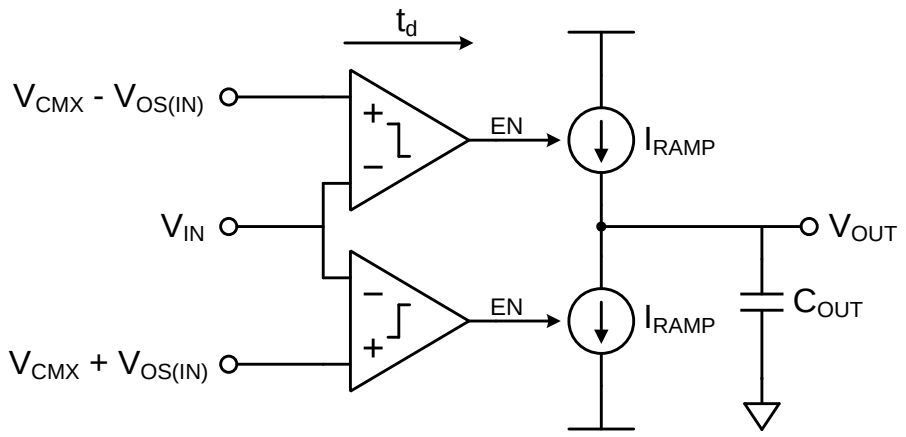

Fig. 5: Conceptual model of a ringamp during the initial slew-charging phase of operation. This model only applies to the initial charging phase and does not include the key ringamp stabilization mechanisms. $V_{OS(IN)}$ is the input-referred value of the dead-zone offset.

output transistor (M_{CN} or M_{CP}) to use depending on the value of the input signal. The selected output transistor then operates as a maximally-biased current source and charges the output load with a ramp. In this initial charging phase the ringamp behaves similar to a zero-crossing based circuit [4] [5].

The ramping phase ends when the input signal crosses the threshold of the comparator and the current source turns off. Due to the finite time delay of the comparator, there will be some amount of overshoot beyond the comparator threshold, which will be given by:

$$\Delta V_{overshoot} = \frac{t_d \cdot I_{RAMP}}{C_{OUT}} \tag{3}$$

where t_d is the time delay of the comparator decision, I_{RAMP} is the current supplied by the active current source, and C_{OUT} is the total loading capacitance seen at the output. This overshoot is with respect to the trip point, which will be on the boundary of the dead-zone. It will be more useful later on if we consider Eq. 2 as well, and express the input-referred overshoot with respect to the ideal settled value (the center of the dead-zone):

$$\Delta V_{init} = \frac{t_d \cdot I_{RAMP}}{\psi\, C_{OUT}} - \left|\frac{V_{DZ}}{2{\cdot}\hat{A}_1}\right| \tag{4}$$

where ψ is the scaling factor that refers the output overshoot to the ringamp's input (and depends on feedback factor, parasitics, and feedback structure) and $\hat{A}_1$ is the effective gain of the first stage inverter at the end of the ramping operation (explained later).

Stabilization

After the initial charging ramp, the ring amplifier will begin to oscillate around the target settled value with amplitude ΔV_{init}. With no dead-zone, the structure is functionally identical to a three-inverter ring oscillator, and will continue to oscillate indefinitely (Fig. 3a). However, as the size of the dead-zone is increased, the ringamp will eventually reach an operating condition where it is able to self-stabilize, such as in Fig. 4. If the dead-zone size is increased further still, the time required to stabilize decreases substantially, and for most practical designs, a ringamp will stabilize in only one or two periods of oscillation (i.e. Fig. 3b).

The most fundamental mechanism in the process of stabilization is the progressive reduction in the peak overdrive voltage applied to the output transistors M_{CN} and M_{CP} on each successive period of oscillation. This effect is illustrated in Fig. 4 by the progressive decrease in amplitude of the signals V_{BP} and V_{BN}. When the following relation is true, the trough (minimum value) of V_{BP} will be limited by the finite-gain of the first two stages, and begin to de-saturate from rail-to-rail operation:

$$\hat{A}_2[\hat{A}_1(min(\tilde{V}_{IN}) - V_{CMX}) - V_{OS}] \geq V_{SS} - V_{CM} \tag{5}$$

(where $\tilde{V}_{IN}$ is the peak-to-peak amplitude, and $\hat{A}_1, \hat{A}_2$ are the negative-valued effective instantaneous inverter gains). A similar relation can also be expressed for the lower signal path and V_{BN}:

$$\hat{A}_2[\hat{A}_1(max(\tilde{V}_{IN}) - V_{CMX}) + V_{OS}] \leq V_{DD} - V_{CM} \tag{6}$$

The key point to notice in these expressions is that each signal path is being fed a different shifted replica of the oscillatory waveform generated at V_A. The upper path is given a replica where the *peaks* of the wave are *lowered* closer to the second stage inverter's threshold, and the lower path is given a replica where the *troughs* of the wave are *raised* closer to the threshold of the second stage inverter. For a sufficiently large shift in each path (V_{OS}), this creates the possibility that even for relatively large values of $\tilde{V}_{IN}$, finite gain effects will simultaneously limit the overdrive voltage that is applied to both M_{CP} and M_{CN}. This stands in stark contrast to the behavior of a three-inverter ring oscillator, where the decrease in V_{OV} of one output transistor necessarily means an increase in V_{OV} applied to the other.

When Eqs. 5 and 6 are true, the resulting reduction in V_{OV} applied to the output transistors M_{CN} and M_{CP} will reduce the magnitude of the output current

I_{RAMP}. This decrease in output current will also cause a decrease in the amplitude of $\tilde{V}_{IN}$ by a proportional amount, due to Eq. 4. The left sides of Eqs. 5 and 6 are therefore reduced further, and the V_{OV}'s of M_{CN} and M_{CP} will decrease even more for the next oscillation cycle. This effect will continue to feedback until the input signal amplitude becomes smaller than the input-referred value of the dead-zone, at which point the ring amplifier will stabilize and lock into the dead-zone.

If we combine Eqs. 5 and 6 and rearrange, we see that in order to trigger this progressive overdrive reduction effect, the input signal must satisfy the following relation:

$$\tilde{V}_{IN} \leq \frac{1}{\hat{A}_1}\left(\frac{V_{DD}-V_{SS}}{\hat{A}_2} - V_{DZ}\right). \tag{7}$$

Furthermore, at the beginning of the stabilization phase:

$$\tilde{V}_{IN} = 2{\cdot}\Delta V_{init} \tag{8}$$

Finally, using Eqs. 3, 4, 7, and 8, we can express the stability criterion in terms of the dead-zone (i.e. settled accuracy) and the initial slew rate (i.e. speed):

$$\frac{t_d \cdot I_{RAMP}}{\psi\, C_{OUT}} \leq \frac{1}{2{\cdot}\hat{A}_1}\left(\frac{V_{DD}-V_{SS}}{\hat{A}_2} - 2{\cdot}V_{DZ}\right) \tag{9}$$

Recall once again that $\hat{A}_1$ and $\hat{A}_2$ are negative valued gains.

From this relation we see that there is a clear design tradeoff between accuracy, speed, and power. Let's assume for a moment that only t_d, I_{RAMP}, and V_{DZ} can be adjusted. To increase speed, one can either increase the initial ramp rate or decrease the time required to stabilize. Both options require sacrificing either accuracy (by increasing V_{DZ}) or power (by decreasing t_d). Likewise, to increase accuracy (by decreasing V_{DZ}), one must either decrease I_{RAMP} or decrease t_d accordingly. While these simple tradeoffs serve as a good starting point, as we will soon discover, every parameter in Eq. 9 is variable to some extent.

The discussion thus far is only a first-order model, and there are additional bandwidth, slewing, and device biasing dynamics which are not represented. Let's take a moment to evaluate this model in the form of a practical example. Consider a pseudo-differential ringamp where $A_1 = A_2 = -25\frac{V}{V}$, $V_{DZ} = 100mV$, $V_{DD} = 1.2V$, and $V_{SS} = 0V$. By Eq. 2, the input-referred size of the dead-zone will be about 4mV, which for a 2V pk-pk input signal would ideally be accurate enough to achieve an input-referred SNDR of 54dB. By Eq. 7, the maximum allowable peak-to-peak amplitude of $\tilde{V}_{IN}$ is approximately 6mV, and by Eq. 4, the maximum allowable input-referred overshoot at the end of the initial ramping phase must be less than 5mV.

This isn't a very encouraging result, since such a small overshoot will place a tight constraint on the parameters in Eq. 3. However, if one were to simulate this same scenario, it will turn out that the peak-to-peak amplitude of oscillation can be significantly larger than the predicted 6mV and still achieve stability. A closer look at Fig. 4 reveals an important contributor to this disparity between theory and practice. Although the AC small-signal gain of the first stage inverter, A_1, may be $-25\frac{V}{V}$, the effective instantaneous value

$$\hat{A}_1(t) = \frac{V_A(t)}{V_{IN}(t)} \tag{10}$$

in the actual transient waveform will be several times smaller at the beginning of stabilization. Thus, although the overall accuracy of the ringamp is determined by the final, settled, small-signal value of A_1, the stability criterion is determined by the initial, transient, large signal effective value of A_1. This reduction in A_1 occurs because the first stage inverter inherently operates around its trip point, where it will be slew limited. The maximum slewing current that the inverter can provide will be

$$I_{slew} = I_P - I_N \tag{11}$$

and for a square law MOSFET model, this will become:

$$I_{slew} = 2k'\frac{W}{L}\left(\frac{V_{DD} - V_{SS}}{2}\right)V_{IN} \tag{12}$$

Notice here that the slew current is linearly related (not quadratically) to the input voltage. Thus, for the first stage inverter, slew rate limiting (and finite bandwidth) has an important impact on determining the effective value of $\hat{A}_1$ during stabilization (and to a lesser extent, the value of $\hat{A}_2$). This dynamic adjustment of the effective inverter gain is a very attractive characteristic, and improves the design tradeoff between speed, accuracy, and power by a significant factor. Similar effects also influence the operation of the second stage inverters in an additional way: although Eqs. 5 and 6 assume rail-to-rail swing for the second stage inverters when $\tilde{V}_{IN}$ is large, in reality the output swing of the second stage inverters may never completely reach rail-to-rail, regardless of the value of $\tilde{V}_{IN}$ due to slew rate limiting, finite bandwidth, and triode device operation.

Relating the discussion of progressive overdrive reduction in this section back to the steady-state stability discussion of Section 2, this behavior can be conceptualized as a dynamic adjustment of the ringamp's output pole corner frequency. The decrease in output current due to V_{OV} reduction increases the output impedance (R_o) of the ringamp, and pushes the output pole (formed by R_o and C_{LOAD}) to lower frequency. As the V_{OV} reduction effect gains momentum on each successive oscillation half-period, the output pole progressively pushes to lower and

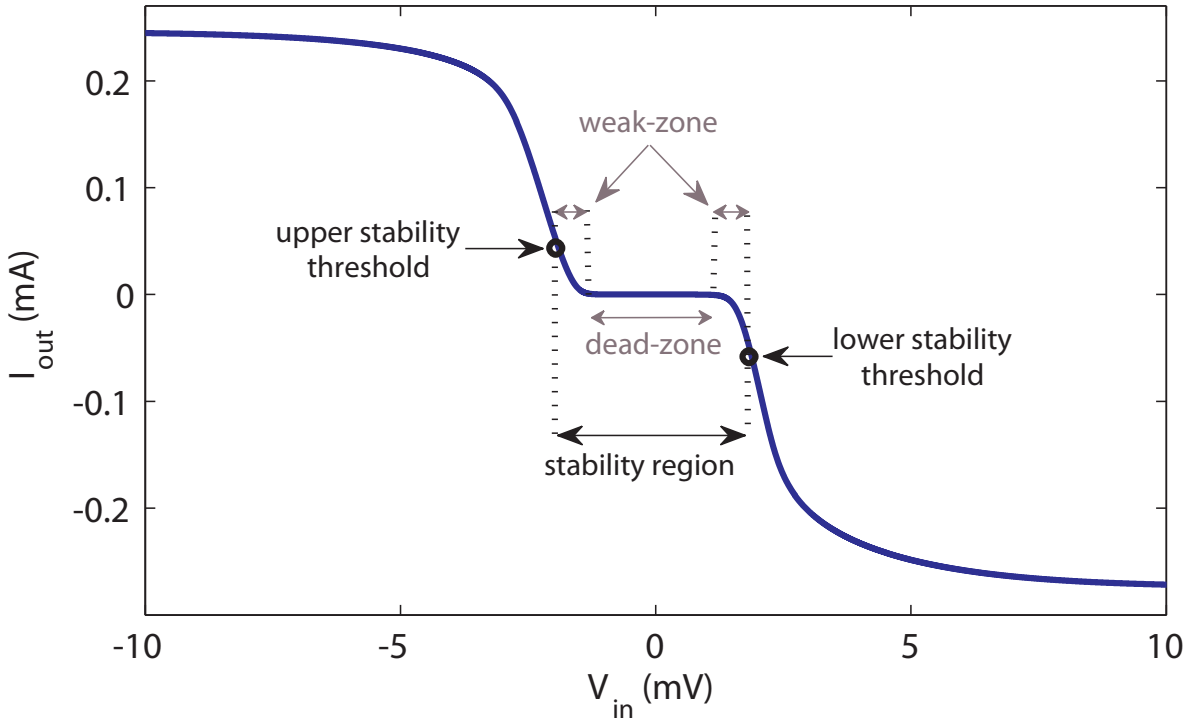

Fig. 6: DC sweep of V_{IN} vs. I_{OUT} for a typical configuration of Fig. 2, illustrating the full input-referred characteristic near the dead-zone region. In addition to a true "dead-zone" where both output transistors are in cutoff, there is also a small boundary region "weak-zone" where the output pole location is low enough to create stability.

lower frequency. By the time the ringamp is locked into the dead-zone and the output transistors are in cutoff, R_o is infinite and the output pole is at DC.

Steady State

Thus far, we have defined the steady-state condition for a ring amplifier as the complete cutoff of both output transistors, with the input signal lying solidly within the dead-zone, such as is the case in Fig. 4. However, considering the discussion about pole adjustment in the previous paragraph, it's clear that the ringamp can in fact be stable for a range of low frequency output pole locations down to DC. Such a situation will in practice occur often, even for a large dead-zone, since there is always a finite probability that the ring amplifier will happen to stabilize right at the edge of the dead-zone. If that happens, one of the output transistors will still conduct a small amount of current to the output, and may never fully shut off before the amplification period ends. The existence of this stable, boundary-region "weak-zone" is illustrated in the V_{IN} vs. I_{OUT} plot of Fig. 6.

The weak-zone isn't an inherent problem for ring amplification operation, since any low-bandwidth settling will only serve to further improve accuracy. However, there are sometimes higher-level structural considerations that make it advantageous to ensure that both output transistors are completely non-conducting once settled. We will see some designs where this is the case later on, in Section 4.

3.1. Key Advantages

Ring amplifiers are in many ways both structurally and functionally quite different from conventional opamps, and it is in these differences that the ringamp finds a unique advantage in the context of modern low-voltage CMOS process technologies. In this section, we will examine several of these important benefits in greater detail.

Output Compression Immunity

In low-voltage scaled environments, kT/C noise, SNR, and power constraints will typically be dictated by the usable signal range available [1], and any practical amplification solution for scaled CMOS must therefore utilize as much of the available voltage range as possible. As it turns out, ring amplifiers are almost entirely immune to output compression, and this enables them to amplify with rail-to-rail output swing.

To understand the basis of this output compression immunity, we must consider two scenarios. First, imagine a ringamp whose dead-zone is large enough that when the ringamp is locked into the center of the dead-zone, both M_{CN} and M_{CP} will be in cutoff. In other words, when:

$$V_{DZ} \geq \left| \frac{V_{DD} - V_{SS} - 2V_T}{A_2} \right| \tag{13}$$

As a rule of thumb, this relation will usually hold for low and medium accuracy ringamps up to about 60dB. Under this scenario, M_{CN} and M_{CP} function as current sources whose linearity and small-signal gain has no appreciable effect on settled accuracy. The internal condition of the ringamp depends only on the signal at the input, and it will continue to steer toward the dead-zone until M_{CN} and M_{CP} are completely cut-off, regardless of whether they are in saturation or triode. Final settled accuracy will be governed by Eq. 2, independent of the characteristics at the output.

Now let's consider the condition where Eq. 13 does not hold. This will occur when the dead-zone is very small, and accuracies in the 60dB to 90dB range are desired. Although other practical issues in the ringamp structure of Fig. 2 may hinder such design targets, we will consider a ringamp structure later in Section 4 where it applies. In this scenario, the stability region of Fig. 6 is so small that the two weak-zones touch, and M_{CN} and M_{CP} will still conduct a small amount once settled. The ringamp's steady state condition will essentially be that of a three stage opamp, and the open loop gain will be the product of the three stage gains. With no true dead-zone, the distortion term of Eq. 2 becomes zero, and finite loop gain will become the fundamental limitation on accuracy. At first glance, generating sufficient loop gain appears to be a problem, since the gain of M_{CN} and M_{CP} will depend on output swing (which must be as large as possible in nanoscale CMOS). Consider the case where all three stages have a gain of 25dB when operating in saturation. In the best case, the open loop gain will be 75dB, and in the worst case perhaps 50dB. Even in the best case, this seems to suggest that to build an 80dB accurate ringamp, an additional gain stage is required.

Luckily, there is another effect at play here. In the ideal square-law MOSFET model M_{CN} and M_{CP} will be in saturation when $V_{OV} < V_{DS}$. Furthermore, the small signal output impedance, r_o, is inversely proportional to the drain current, I_D. In the context of the progressive overdrive voltage reduction that occurs in ringamp stabilization, both V_{OV} and I_D will in fact trend towards zero. This implies that during steady-state, M_{CN} and M_{CP} will remain in saturation or weak-inversion even for very small values of V_{DS}, and moreover, that their gain will be enhanced by a dynamic boost in r_o. Thus, even for a nominal open loop gain of 75dB, with a wisely chosen topology it is possible to have an enhanced steady-state gain of at least 90dB, even when swinging close to the rails.

Although output swing has little effect on ringamp accuracy, it will indeed affect speed, both with respect to slewing and settling. In the initial ramping phase, the selected current source transistor will be biased with the maximum possible V_{OV}, and this guarantees that for much of the possible output range it will initially be operating in triode. As seen in Fig. 7, for settled output values near mid-rail, I_{RAMP} will be the highest and the initial ramping will be faster, but more time will be required to stabilize for the reasons discussed in Section 3. Likewise, for values close to the rails, I_{RAMP} will be smaller, so the initial ramping will be slower but the stabilization time will be shorter. For the most part, this works out quite nicely, since the total time required to reach steady state in each case turns out to be approximately the same. However, for extreme cases very close to the rails, the large RC time constant of the output transistor in triode operation will require a comparatively long time to reach its target value. Ultimately, it is

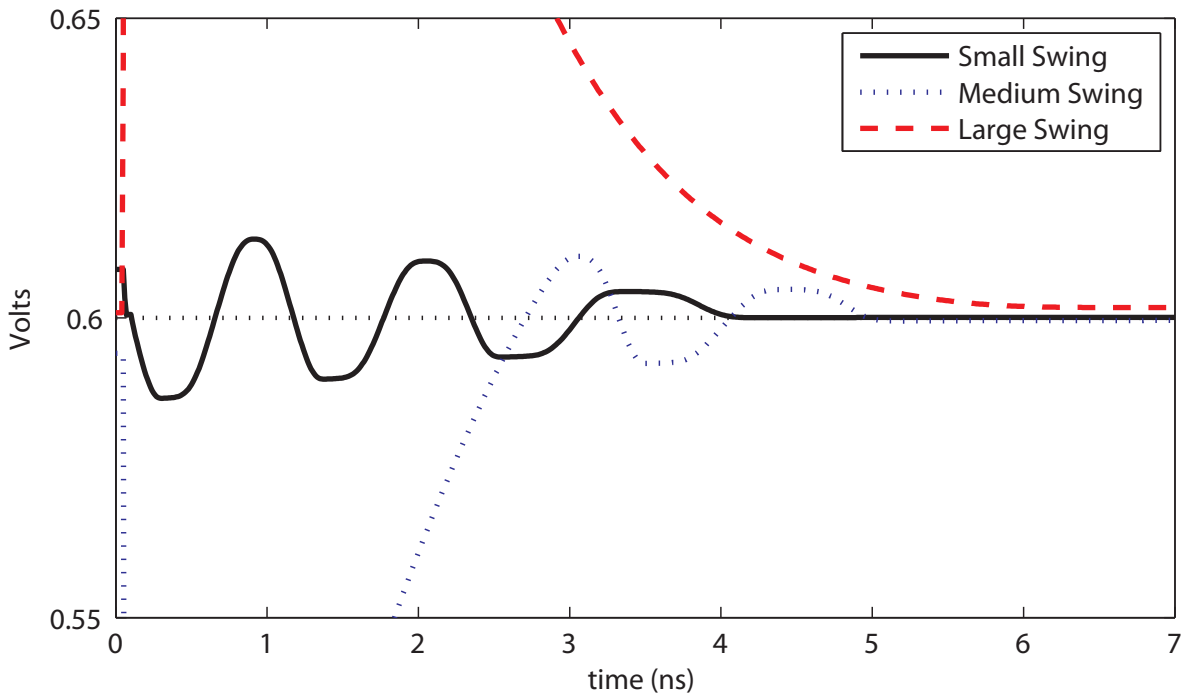

Fig. 7: Zoomed stabilization waveform of V_{IN} for three output swing cases: small (output near mid-rail), medium, and large (final output near the supply).

this RC settling limitation that will usually dictate the maximum output swing possible for a given speed of operation.

Slew-Based Charging

Whereas a conventional opamp charges its output load with some form of RC-based settling, the output transistors M_{CN} and M_{CP} in the ring amplifier behave like digitally switched current sources, and charge the output with slew-based ramping. This is a much more efficient way to charge, since only one of the current sources in Fig. 5 will be active at a time, and the only power dissipated will be dynamic. Furthermore, during the initial ramping operation, M_{CN} or M_{CP} (whichever is selected) will be biased with the maximum V_{OV} possible for the given supply voltage. This is a major benefit, because it means that even for large capacitive loads, small transistor sizes can still produce high slew rates, and with small output transistors, the second stage inverters will be negligibly loaded by M_{CN}/M_{CP}. This effectively decouples the internal power requirements from that of the output load size, and for typical load capacitances in the femto and pico-farad range, the internal power requirements are more-or-less independent of output capacitance. This unique property stands in stark contrast to the power-loading relationship for a conventional opamp, where settling speed is typically

proportional to g_m/C_{LOAD}. Even for large load capacitances, where the size of M_{CN}/M_{CP} does have an appreciable effect on the internal power requirements, the ratio of static-to-dynamic power will scale very favorably.

Performance Scaling with Process

In order for a technique to be truly scalable, it must meet two criteria. First, the given technique must operate efficiently *in* a scaled environment. This requirement has been our primary focus thus far. Second, the technique must inherently scale *with* advancing process technology, improving in performance simply by migrating into a newer technology. It is this second criteria that we will discuss now.

Intuitively, the ring amplifier seems like a prime candidate to benefit from process scaling, simply due to its structural similarity to a ring oscillator. The stability criterion of Eq. 9 suggests this to be true. As stated previously, the internal power consumption of a ringamp is governed much more by inverter power-delay product and internal parasitics than the size of the output load (in stark contrast to conventional opamps). Since the power-delay product of an inverter decreases approximately linearly in accordance with decreasing feature size [6], the ringamp's inverter chain propagation delay, t_d, can be expected to scale according to digital process performance as well. With the relationships in Eq. 9, this reduction in t_d can be directly traded for an improvement in any of the three main design specifications: speed, accuracy, and power.

The simple scaling experiment conducted in [7] suggests that this is indeed the case. The results of the test are shown in Fig. 8. The upper trend line represents predicted power efficiency with the power spent in charging the fixed (not scaled) load capacitance included. The lower trend line shows the power efficiency with the ideal power required to charge the load capacitance subtracted out. The result is a power efficiency trend that scales very well with advancing technology node. Although there are not yet enough measured ringamp designs to verify predicted trend, we do see a similar level of speed and efficiency scaling between the 0.18μm CMOS design of [8] and the 65nm CMOS design of [9]. Furthermore, recent investigations of ringamp structures in 28nm CMOS have produced results that also support this hypothesis.

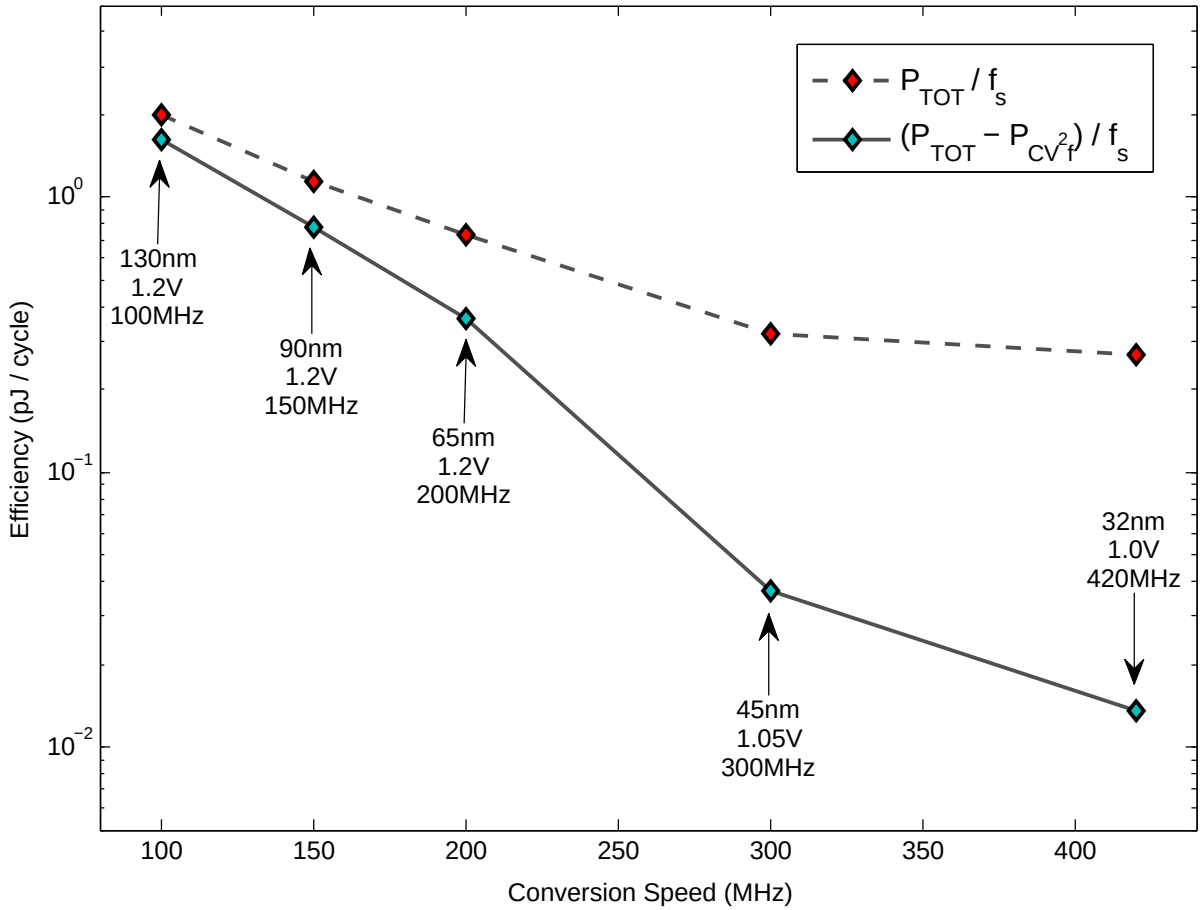

Fig. 8: Approximate power efficiency scaling trends predicted by [7].

4. A Diversity of Solutions

There are many different opamp structures available to designers to meet diverse needs. The same should also be true for ringamps. Accuracy, speed, power, and design effort are all important factors in choosing the best ringamp structure for the job. Although the exploration of possible ringamp techniques is still in its infancy, we can at least look at the comparative merits of those ringamp implementations which are already known and reported in the literature.

To begin with, for high-speed, medium-accuracy amplification, the ringamp circuits of Fig. 1b and Fig. 2 are both attractive candidates. In technologies where the gain of the ringamp is more than that required for accuracy specifications, the ringamp of Fig. 2 is a good choice. Technologies where this is the case include older processes such as 0.18μm, 0.13μm, and 90nm, and possibly even some newer processes such as 14nm FinFET, where inverter gains are improved [3]. The dead-zone embedding in the structure of Fig. 2 allows bandwidth to be decoupled from the small-signal gain of the inverters. Even if the small-signal gain

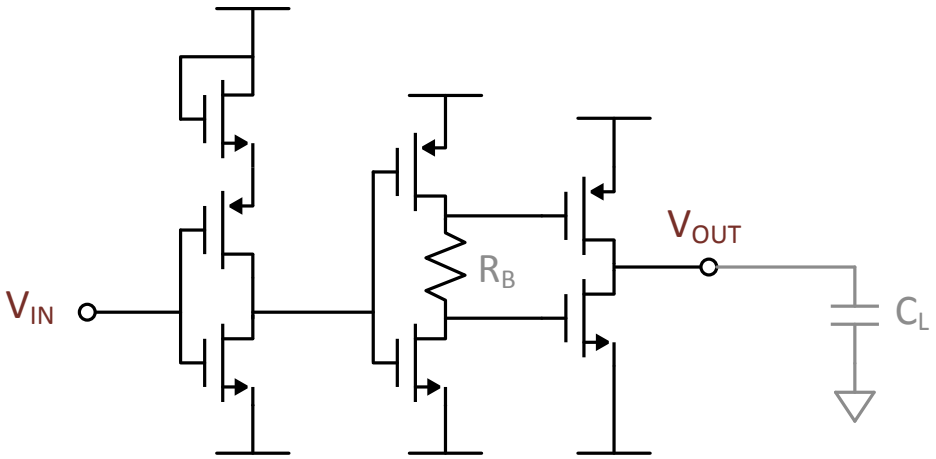

Fig. 9: Self-biased ringamp structure of [9]

of the inverter chain is very high, the dead-zone allows us to effectively reduce it and improve the gain-bandwidth product while still benefiting from the fast transient behavior associated with high gain inverters. This structure is used for the purposes of medium-accuracy amplification in the 10.5b pipelined ADC of [8], and as a sub-component of the 15b Split-CLS pipelined ADC of [10] and the Composite Ringamp Amplifier Block of [11].

Alternatively, if working in a technology with low intrinsic inverter gains such as 65nm, 45nm, 32nm, 28nm, and 22nm planar CMOS, it may not be possible to set a true dead-zone with the structure of Fig. 2. In this case, the simple inverter chain of Fig. 1b may actually be a realistic option to consider. The structure still scales excellently and retains many (although not all) of the benefits of a dead-zone stabilized ringamp. However, the ringamp structure introduced in [9], shown in Fig. 9, provides an even better solution for most scenarios. Like Fig. 1b it is also a single chain of inverters. However, it uses resistor R_B to embed a weak-zone offset for improved dynamic biasing, enhancing both efficiency and speed. The design of [9] also introduces the idea of using the time information contained in the internal nodes of a ringamp to perform quantization, and this time-domain information is used to build a 1.5b pipeline stage sub-ADC that results in significantly relaxed timing constraints.

For design scenarios requiring high accuracy amplification, a different approach to ringamp design is required. Fully differential operation is typically mandatory, and the amplifier must either produce very high effective gain or be assisted by calibration. In [10] the technique of Split-CLS [12] is used to combine a pseudo-differential ringamp structure with a fully differential telescopic opamp. Fig. 10a shows the basic principle of this technique. The ringamps provide an initial fast and coarse charge of the output. Then, the opamp is coupled into the output via the

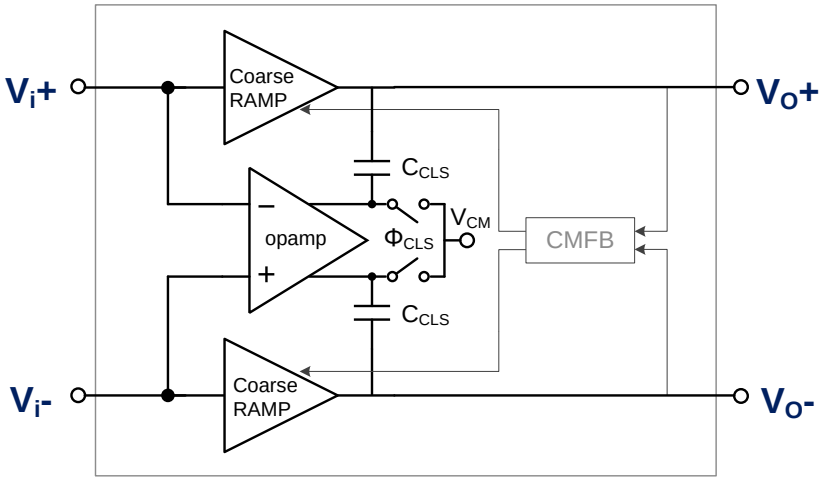

(a) Split-CLS Structure [10]

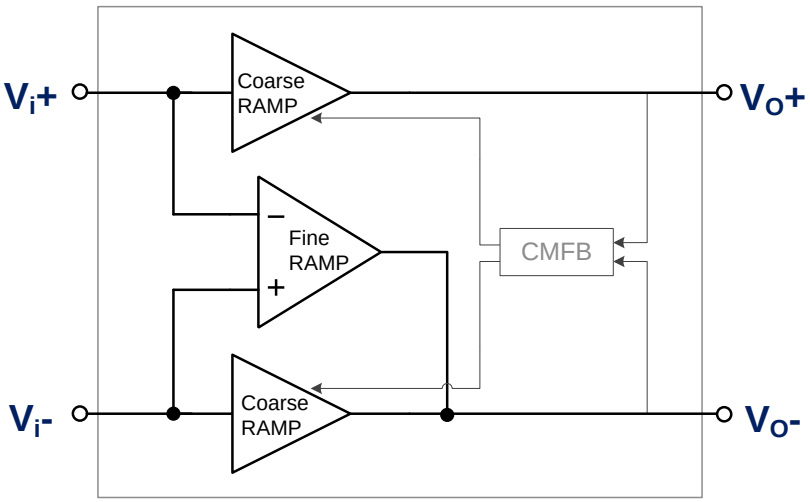

(b) Composite Ring Amplifier Block [11]

Fig. 10: Ringamp structures suitable for high-accuracy applications.

level shifting capacitors (C_{CLS}) and proceeds to fine settle the output with greatly relaxed slew and swing requirements. At the end of an amplification period, the total accuracy of the output is the combined accuracy of the ringamp *and* the telescopic opamp. This structure is capable of very high accuracy, and allows a robust differential opamp to be the final determinant of accuracy.

Another approach to precision amplification is the structure of Fig 10b. Introduced in [11], a Composite Ring Amplifier Block uses only ring amplifiers (no opamps) and consists of a coarse but fast and efficient pseudo-differential ringamp connected in parallel with a differential-input, single-output precision ringamp. When placed in this parallel configuration, the coarse ringamp will automatically and asynchronously cutoff and transfer control to the fine ringamp at the correct moment of operation. Initially, all ringamps are enabled, and contribute charge

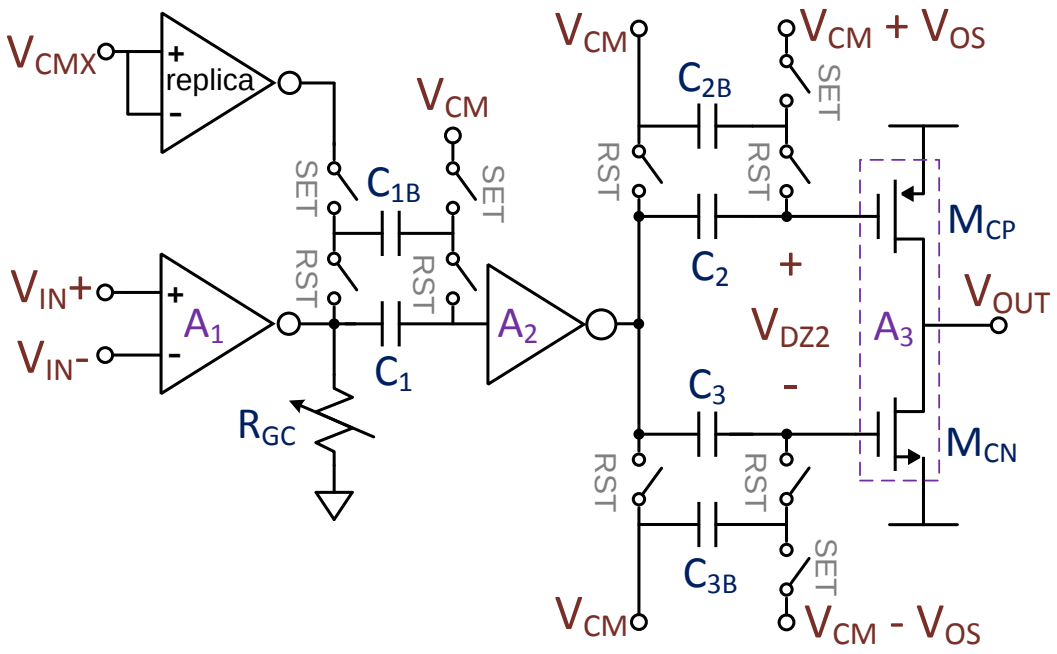

Fig. 11: Precision ringamp circuit used in the Composite Ring Amplifier Block (Fig. 10b) of [11].

to the output. However, the coarse ringamps have a larger slewing capability and dominate the initial charging behavior, quickly settling both the differential and common-mode levels close to their final target values. As the coarse ringamps enter their dead-zone, they automatically disconnect from the output. The fine ringamp remains active, however, since its stability region (weak-zone) is necessarily quite small and thus completely enclosed by the larger stability region (dead-zone) of the coarse ringamps. With the common-mode level already settled to sufficient accuracy by the coarse ringamp's CMFB, and V_{O+} floating, the fine ringamp simply settles V_{O-} differentially around a stationary V_{O+}. Thus, in addition to providing speed enhancement this scheme removes the need for common-mode feedback in the fine ringamp, permitting a single ended output to be used without loss of accuracy, thereby minimizing both complexity and power.

A simplified schematic of the fine ringamp is depicted in Fig. 11. Notably, the offset V_{DZ2} is embedded just prior to M_{CN}/M_{CP}, and allows the settled value of V_{OV} to be precisely set (and weak-zone operation to be guaranteed). However, the constraint that this places on the value of V_{DZ2} consequently limits its ability to tune stability. This is solved by observing that stability is actually determined by the input-referred value of V_{DZ2}, which can also be tuned by adjusting the gain of either the first or second stage inverter. Thus, a tunable gain-control resistor (R_{GC}) is used to set the size of the stability region. There is almost no linearity requirement for this tuned resistance, and it is implemented in [11] as a simple

Table 1: Summary of ringamp ADCs

	VLSI '12 [8]	**ISSCC '12 [10]**	**VLSI '13 [11]**	**ISSCC '14 [9]**
Technology	0.18μm	0.18μm	0.18μm	65nm
Supply Voltage	1.3 V	1.3 V	1.2 V	1.2 V
Resolution	10.5b	15b	15b	10.5b
Input Range	2.2 V pk-pk	2.5 V pk-pk	2.4 V pk-pk	2 V pk-pk
Sampling Rate	30 Msps	20 Msps	20 Msps	100 Msps
SNDR	61.5 dB	76.8 dB	75.9 dB	56.3 dB
SFDR	74.2 dB	95.4 dB	91.4 dB	67.6 dB
Total Power	2.6 mW	5.1 mW	2.96 mW	2.46 mW
FoM	90 fJ/c-step	45 fJ/c-step	29 fJ/c-step	46 fJ/c-step

3-bit DAC composed of tiny MOSFET resistive elements.

A summary of the measured performance of the ringamp-based pipelined ADCs discussed here is provided in Table 1.

5. Conclusion

Even in the initial design attempts listed in Table 1, we already see very promising performance numbers being achieved across a range of target accuracies and speeds. Performance will continue to improve in the future for a couple of reasons. First, simply scaling down into newer technology nodes should yield substantial benefits. Second, the potential for using time-domain properties to improve ringamp efficiency is by no means exhausted. For example, the idea to exploit the time-domain behavior of a ringamp to perform quantization in [9] is a solution with attractive benefits. Some of the insight gained in the research of other emerging time-domain techniques such as VCO-based quantizers [13] may prove relevant to ringamp research as well.

Just how many niches in the circuit ecosystem ring amplification can enhance diversity in remains to be determined. It is already evident that they are useful in pipelined ADCs. Many other ADC architectures can benefit from ringamps as well. For example, they are an enticing candidate for use in the integrator structures in discrete-time sigma delta modulators. Many of the tradeoffs associated with conventional opamps have influenced which $\Sigma\Delta$ topologies are ultimately the most successful. Now that we have an amplifier where swing, loading capacitance, and gain aren't nearly as constraining, we can discard old assumptions and reconsider the possibilities. This may even allow discrete-time $\Sigma\Delta$ ADCs to ex-

tend their accuracy and robustness benefits to bandwidths on the order of tens of megasamples that are currently achieved only by continuous-time $\Sigma\Delta$ ADCs [2].

There is also much to explore beyond the realm of ADCs. Anything with a capacitive load is a prime candidate for consideration. This includes switched-capacitor circuits such as discrete-time filters as well as a variety of sensing and imaging applications. In all these cases, it is once again useful to re-examine many of the assumptions about what constitutes an "optimal" structure for a given application with specific regard to the strengths and weaknesses of ringamps. In some cases, ringamps may provide the best solution. In other cases, a different technique in the circuit ecosystem will. This is the strength of circuit diversity.

References

[1] B. Jonsson, "On cmos scaling and a/d-converter performance," in *NORCHIP, 2010*, nov. 2010, pp. 1 –4.

[2] B. Murmann. (2014) Adc performance survey 1997-2014. [Online]. Available: http://www.stanford.edu/ murmann/adcsurvey.html

[3] M. Shrivastava, R. Mehta, S. Gupta, N. Agrawal, M. Baghini, D. Sharma, T. Schulz, K. Arnim, W. Molzer, H. Gossner, and V. Rao, "Toward system on chip (soc) development using finfet technology: Challenges, solutions, process co-development and optimization guidelines," *Electron Devices, IEEE Transactions on*, vol. 58, no. 6, pp. 1597 –1607, june 2011.

[4] J. Fiorenza, T. Sepke, P. Holloway, C. Sodini, and H.-S. Lee, "Comparator-based switched-capacitor circuits for scaled cmos technologies," *Solid-State Circuits, IEEE Journal of*, vol. 41, no. 12, pp. 2658 –2668, dec. 2006.

[5] L. Brooks and H.-S. Lee, "A 12b, 50 ms/s, fully differential zero-crossing based pipelined adc," *Solid-State Circuits, IEEE Journal of*, vol. 44, no. 12, pp. 3329 –3343, dec. 2009.

[6] H. Iwai, "Cmos scaling towards its limits," in *Solid-State and Integrated Circuit Technology, 1998. Proceedings. 1998 5th International Conference on*, 1998, pp. 31 –34.

[7] B. Hershberg, S. Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U. Moon, "Ring amplifiers for switched capacitor circuits," *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 12, pp. 2928–2942, Dec 2012.

[8] B. Hershberg, S. Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U. Moon, "A 61.5db sndr pipelined adc using simple highly-scalable ring amplifiers," in *VLSI Circuits (VLSIC), 2012 Symposium on*, June 2012, pp. 32–33.

[9] Y. Lim and F. P. Flynn, "A 100ms/s, 10.5-bit, 2.46mw comparator-less pipeline adc using self-biased ring amplifiers," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, feb. 2014.

[10] B. Hershberg, S. Weaver, K. Sobue, S. Takeuchi, K. Hamashita, and U. Moon, "Ring amplifiers for switched-capacitor circuits," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, Feb 2012, pp. 460–462.

[11] B. Hershberg and U. Moon, "A 75.9db-sndr 2.96mw 29fj/conv-step ringamp-only pipelined adc," in *VLSI Circuits (VLSIC), 2013 Symposium on*, June 2013, pp. C94–C95.

[12] B. Hershberg, S. Weaver, and U. Moon, "Design of a split-cls pipelined adc with full signal swing using an accurate but fractional signal swing opamp," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 12, pp. 2620 –2630, dec. 2010.

[13] M. Straayer and M. Perrott, "A 12-bit, 10-mhz bandwidth, continuous-time $\Sigma\Delta$ adc with a 5-bit, 950-ms/s vco-based quantizer," *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 4, pp. 805–814, April 2008.