A Low-Power Microcontroller in a 40-nm CMOS Using Charge Recycling

Blutman, Kristof; Kapoor, Ajay; Majumdar, Arjun; Martinez, Jacinto Garcia; Echeverri, Juan; Sevat, Leo; van der Wel, Arnoud P.; Fatemi, Hamed; Makinwa, Kofi A.A.

DOI
10.1109/JSSC.2016.2637003

Publication date
2017

Document Version
Accepted author manuscript

Published in
IEEE Journal of Solid State Circuits

Citation (APA)
Blutman, K., Kapoor, A., Majumdar, A., Martinez, J. G., Echeverri, J., Sevat, L., van der Wel, A. P., Fatemi, H., & Makinwa, K. A. A. (2017). A Low-Power Microcontroller in a 40-nm CMOS Using Charge Recycling. *IEEE Journal of Solid State Circuits*, *52*(4), 950-960. Article 7815353. https://doi.org/10.1109/JSSC.2016.2637003

A Low Power Microcontroller in 40nm CMOS using Charge Recycling

Kristof Blutman[1], Ajay Kapoor[1], Arjun Majumdar[1], Jacinto Garcia Martinez[1,2], Juan Echeverri[1], Leo Sevat[1], Arnoud van der Wel[1], Hamed Fatemi[1], Kofi Makinwa[3], José Pineda de Gyvez[1]

[1] NXP Semiconductors, Eindhoven, the Netherlands

[2] Fontys University of Applied Sciences, Eindhoven, the Netherlands

[3] Delft University of Technology, Delft, the Netherlands

Email: {kristof.blutman, ajay.kapoor, jose.pineda.de.gyvez} @ nxp.com

***Abstract* – A 40nm microcontroller featuring voltage-stacked memory and logic is presented. This involved connecting the power domains of the memory and logic in series, such that the ground of one power domain is connected to the positive supply rail of the other. In this work, an ARM Cortex-M0+ and its peripherals are powered from 0V and V_{DD}, while its 4kB ROM and the 16kB SRAM are powered from V_{DD} and $2V_{DD}$. Since the memory and logic will, in general, draw different supply currents, the mid-rail V_{DD} is provided by an on-chip switched-capacitor voltage regulator (SCVR). To allow direct comparison of voltage stacking with a conventional single supply, it can be turned off by configuring the SCVR to power both the memory and logic from 0V and V_{DD}. Turning on voltage stacking results in 96% power conversion efficiency, while the active converter area is reduced by 2.6x. Despite the use of a smaller SCVR, voltage stacking reduces the supply noise by 3.4dB and the output voltage drop from 58mV to 36mV.**

***Index Terms* – microcontrollers, switched capacitor regulators, balanced voltage islands, charge-recycling, power management, level shifter, voltage stacking**

I. INTRODUCTION

Power management plays a key role in an era of autonomously connected devices. Such devices form the backbone of the so-called 'Internet-of-Things' (IoT), and so must be as power efficient as possible in order to maximize their power autonomy and battery life. However, the steady scaling of CMOS technology has been accompanied by lower supply voltages, which translates into larger supply currents for a given power dissipation. At the same time, interconnect resistance has increased, leading to larger IR drops and making it increasingly difficult to realize efficient power management systems. Furthermore, battery voltages have not scaled at the same rate as supply voltages, adding to the costs.

To cope with this challenge and to deliver power with high efficiency and minimum area, various types of voltage regulators are used [1][2]. However, none of these are ideal, with each type having its own unique set of advantages and disadvantages. Although linear regulators can be quite area efficient and can provide high amounts of power, their efficiency is ultimately limited by the ratio of their input- and output voltages [3]. Inductive switched-mode power supplies have high efficiency but typically require off-chip discrete components such as inductors. Capacitive switched-mode power supplies can be integrated on-chip with moderate efficiency, at the cost of low area efficiency and output current. While some of these limitations can be mitigated by conventional techniques such as the use of trench capacitors [4] or in-package inductors [5], they cannot be completely circumvented. In this paper, we will introduce the voltage stacking technique [3], and show that it can significantly increase both the power- and area efficiency of on-chip power management systems.

Voltage stacking [3] is a technique that involves connecting power domains in series rather than in parallel. As shown in Fig. 1, disregarding the calculations for now, this is analogous to connecting resistors in series rather than in parallel. If a system has two power domains, each modeled as impedances between supply and ground rails, then a conventional implementation would see these impedances connected in parallel, i.e. connecting ground rail to ground rail and supply rail to supply rail. In contrast, a voltage stacking scenario

would involve connecting the supply rail of the first (bottom) power domain to the ground rail of the second (top) power domain. Assuming that each power domain uses the same supply voltage and current, this means that the total supply voltage doubles, while the total supply current halves. In other words, an implicit 2:1 conversion has been realized without the power conversion losses associated with voltage converters, and with no area overhead.

In practical implementations, the top and bottom power domains will not consume the same current and so a voltage regulator will be needed to stabilize the mid-rail between the series connected power domains. If the supply current mismatch is small, the power supplied by this regulator will be much less than that supplied by the regulator of a conventional, parallel-connected, system. In consequence the voltage regulator of a voltage stacked system can be quite compact. Furthermore, since it delivers less current, its losses will not be significant, and so the overall system power efficiency will be higher than that of a conventional system. Therefore, the implicit conversion step associated with voltage stacking relaxes the requirements on the explicit conversion step (e.g. voltage regulator).

The benefits of voltage stacking have been demonstrated before [6], but its use in realistic applications has only been described in [7]. Earlier work featured either simple circuit blocks, or larger but disconnected systems. In [6] multipliers were used to read out operands from on-chip SRAM. The circuit infrastructure like level shifters between the power domains and voltage regulator that can control the mid-node were well studied and implemented, but the system lacks the complexity of a real application. Similarly, in [8] the concept of voltage regulation is further developed into several linear regulator blocks providing the current in the design, and the application space has similarly been well established in the form of a lock-step MCU system. The final silicon implementation, however, only featured PLLs stacked on top of each other, which is still not a complete system that could demonstrate the feasibility of voltage stacking. In [9] the complexity of the implemented system was much higher with stacked memory blocks and processor cores, however, the separate MCUs were disconnected from each other and did not function as one system that is needed in

realistic application. The other novelty was the introduction of switched capacitor converters into voltage stacking implementation, which employed an adaptive regulation scheme further enhanced by per-core frequency scaling that reduced the current imbalance even further. The final work that treated the topic was in [10] where a stacked IO driver was proposed with thin oxide transistors for low power and high speed. The benefits compared to thick oxide IO pad implementations were due to voltage stacking – although the delivered supply voltage was $2V_{DD}$, the thin oxide devices observed only V_{DD} over their respective voltage range.

In this paper, the realized system addresses the shortcomings of the previous implementations. Firstly, a full microcontroller IC is implemented featuring an ARM Cortex-M0+ processor, 16kB SRAM, 4kB ROM and an on-chip SCVR in 40nm technology and uses only thin-oxide, single threshold voltage transistors. The important part to realize that, while this system features sufficient complexity (20k standard cells), it also functions on its own without the use of duplicates or multiple systems stacked on top of each other. The system is stacked at the IP level where level shifters interface between the power domains. To the best of our knowledge, this is the first work where such scheme is implemented on heterogeneous components in a scalable manner, fully compatible with the conventional digital design flow. Furthermore, while previous voltage stacking implementations were fixed and could not be configured into an alternative operating mode with voltage stacking off, this work features a reconfigurable chip where the system can work both with voltage stacking turned on or off. This way, the benefits of voltage stacking can explicitly be measured on the same chip, which was not done in previous works. Thirdly, previous level shifter implementation were typically impractical for high complexity systems and were not standard cell compatible, while here a novel level shifter is presented that does not rely on exotic elements like capacitors, thick oxide devices and can be densely laid out along with the standard cells.

The test chip not just fulfills the mentioned features, the benefits from voltage stacking are also significant. While voltage stacking was off, the maximum power efficiency was 81%, which improved to 96% for the case

when voltage stacking was enabled. Similarly, the converter area has also reduced and the supply quality has improved.

II. SYSTEM DESIGN

The primary goal of voltage stacking is to reduce the amount of power processed by the voltage regulators. This comes with power- and area benefits. While the area of the converter can be estimated from its type and its output current, the power benefit has to be calculated. If a 100% efficient power delivery system would exist, voltage stacking would not confer power benefits since it does not influence the power consumption of the load circuitry itself, only reduces the losses of the power conversion step. The power efficiency for stacked system can be calculated as the ratio of the output- and input power, illustrated in Fig. 1 b. At the input, I current flows into the upper domain supply rail, while $I_{in,VR}$ flows into the VR input, both at $2V_{DD}$. At the output, I current is the same as at the input, while the VR output observes $I_{out,VR}$ flowing out at V_{DD} voltage. This adds up to a combined system efficiency η_{sys}, which is also derived in Fig. 1:

$$\eta_{sys} = P_{out}/P_{in} = (2I+I_{out,VR}) / (2I+ I_{out,VR}/\eta_{VR}) \quad (1)$$

Where $I_{out,VR}=\Delta I$ is the difference current between the top- and bottom domains that needs to be supplied by the VR. In the limiting case when I=0, voltage stacking is off since there is no charge recycling between the power domains, while when $\Delta I=0$, there is perfect voltage stacking where there is no current mismatch between the power domains. Since the converter size is also proportional to its output current ΔI, the ultimate goal while designing a voltage stacked system is to minimize ΔI. Illustration of (1) for various VR efficiency can be seen in Fig. 2.

While (1) seems simple enough to minimize, there are several secondary effects that are needed to be taken into account. The first effect is that the efficiency of the voltage regulator depends on the output current ΔI. On one hand, according to (1), higher the current imbalance the more ΔI current is sourced from the VR,

increasing the proportion of the current that suffers from conversion losses. This is limited by the output current dependence of the VR efficiency: $\eta_{VR}(\Delta I)$ itself is a function rather than a constant in relation to ΔI. The second mechanism that influences the efficiency is the additional power consumption of the level shifters, which is not represented in (1). Finally, the output impedance of the VR drops the output voltage by $R_{out} \cdot \Delta I$, which scales the voltages applied to the bottom and top power domains, and influences their power consumption.

Based on these considerations, a practical implementation of voltage stacking should, on one hand, benefit in power and area, while on the other hand, should be simple enough to be implemented with minimum overhead in design effort. In the following, we present our voltage-stacked microcontroller system where we focused on maximizing gains with only small modifications to the otherwise complex, standard design flow.

A. Power delivery

Fig. 1 a) shows a conventional power delivery scheme where the power domains are connected in parallel and are supplied by the same voltage regulator. The first power domain in our case incorporates most of the logic of the system, while the second power domain contains the memory. The voltage stacking scenario which has beneficial power efficiency over the traditional approach, is shown in Fig. 1 b). The second (top) power domain is now stacked on top of the first (bottom) power domain. Since in prior works either a) or b) is implemented on-chip, quantifying the benefits of voltage stacking has been largely implicit and relied on many assumptions. Therefore, we transform the system into a reconfigurable architecture where the first power domain containing the logic is fixed between 0V and V_{DD}, while the second power domain with the memory is reconfigurable to be either between 0V and V_{DD} or between V_{DD} and $2V_{DD}$. The same system is compared in two power modes, 'Stacking Off' as shown in Fig. 1 a) and 'Stacking On' as in Fig. 1 b).

In 'Stacking On' mode, the memory becomes the top power domain and the logic the bottom power domain. The current that is drawn by the memory from the external $2V_{DD}$ supply is routed from the V_{DD} ground of the

top power domain to the V_{DD} supply rail of the bottom power domain, which means that this current is directly sourced from the $2V_{DD}$ supply without conversion losses. The SCVR regulating the V_{DD} supply now has to be capable of supplying the system in both power modes.

The transition between power modes should be as seamless as possible. The system should not require a power-up or even a reset after stepping from stacking off to stacking on and back. The whole procedure should only require the clock of the system to be stopped and then resumed.

B. System Architecture

Fig. 3 depicts the architecture of the complete microcontroller system. The bottom power domain incorporates the core and the peripherals, and the top power domain the memory. The top power domain is placed in a deep n-well so that its supply- and ground voltages can be arbitrarily set. With either stacking on or off, the system requires an external supply, $V_{IN}=2V_{DD}$, with $V_{OUT}=V_{DD}$ generated by the SCVR. The two power domains communicate via level shifters in Stacking On mode and through regular buffers acting as a bypass circuitry in Stacking Off mode.

The Cortex-M0+ core accesses the ROM and SRAMs through Advanced High-performance Bus (AHB), which also can be controlled through Serial Wire (SW) interface. Next to general purpose IO (GPIO), the AHB bus is connected to Advanced Peripheral Bus (APB) that provides access to various peripherals like UART module, 32-bit timer and Clock Generation Unit (CGU). The ISRAM can be programmed through the SW interface.

The partitioning choice of placing memory on top of logic was motivated, on one hand, by the system power composition for the typical testcase of Fig. 4. As can be seen from the figure, memory- and logic power consumption is well matched, furthermore, there is not a scenario where one is completely off – in active mode, the core needs to read instructions from the instruction SRAM every clock cycle, toggling several nodes in both power domains and securing a well-balanced power consumption scenario. While more refined

partitioning approaches that exploit high granularity partitioning are possible, our heuristic approach produced good results with minimum design effort which is important in typical applications. Next to this, the system's overhead needs to be taken into account. Basically, the number of level shifters and the layout partitioning cost limit the partitioning efforts. From this perspective simplicity is important, e.g. a minimum number of level shifters and power domains. When counting the number of level shifters, the connections to the IO pads also need to be included. This motivated that the memory would be placed in the top domain and the logic in the bottom one, and not the other way around. No IO pad is directly connected to the memories, this way the order of stacking ensured that no level shifter was needed to interface to the IO pads, which reduced the number of level shifters.

C. Level shifter

As shown before, the level shifter is responsible for the effectiveness of voltage stacking – it must offer minimum timing, power and area overhead. The levels shifter needs to be standard cell compatible and laid out in a dense manner. These requirements necessitate the use of thin oxide devices, which in turn raises questions about reliability. The $2V_{DD}$ voltage that is the maximum voltage drop across the level shifter is usually harmful for thin oxide devices, so special care has to be taken to ensure that there devices are not subject of voltage overstress. Furthermore, the extreme voltage shift between the input and output is a challenge since there is barely any overlap between the voltages of the input and the output. The signal voltages within the level shifter have to be high enough to toggle a change and low enough to avoid the problem of voltage overstress.

Meeting the aforementioned requirements is not possible with most conventional level shifters. Commonly, the application that requires these kind of extreme level shifters is floating voltage HV drivers for e.g. boost converter control circuitry as in [12]. In these applications, however, thick-oxide devices are used. Thin oxide

devices are used in [6] for voltage stacking purposes, however these cells were still heavily dependent on the ratio of the input- and output devices and need a large capacitor to operate.

The proposed level shifters compatible with standard digital cells and entirely made of thin oxide devices, are depicted in Fig. 5. Fig. 5 a (b) shows an up (down) level shifter instance. They are fully-static and can be densely laid out, enabling standard cell-based design. Illustrating the operation on the up level shifterthe input signal is buffered by two inverters that, shielded by four clamp transistors, control a PMOS latch, whose output swings from rail-to-rail without floating nodes, due to the PMOS pull-down transistors. To protect these devices from overvoltage, the PMOS devices have been placed in a so-called “hot” n-well which limits the voltages between two arbitrary terminals of the device to V_{DD}.

The operation of the up level shifter in Fig. 5 a, is as follows (the down level shifter in Fig. 5 b works in a similar way). The input signal propagates through I1 and I2 inverters that buffer it and convert the single-ended input into a differential signal. Being differential, either I1 or I2 assumes a low level value, and the corresponding transistors (either M1-M5 or M2-M6 pair) opens and pulls down the appropriate node in the PMOS latch. Having pulled down one node of the latch, as a results, either M8 or M7 opens and pulls up the other node, producing a differential signal which is buffered by I3 at the output.

From possible voltage overstress point of view, M1-M8 devices need to be carefully analyzed. If the output of I2 is at 0V, M2 is open and pulls down the drains of M4 and M6 to 0V. In turn, the M4 and M6 pull down the gate of M7 to VSS TOP voltage (1.1V). The reason why it is not pulled further is because M6 connects the gate of M7 to VSS TOP, and this turns off M4, which has also VSS TOP at its gate, therefore V_{GS}(M4) becomes 0V. We can now observe that M2 has V_{DSGB}=(0,0,1.1,0)V while M7 has V_{DSGB}=(2.2,2.2,1.1,2.2)V, therefore no overvoltage occurs. For M4, V_{DSGB}=(0,1.1,1.1,1.1)V holds and for M6, V_{DSGB}=(1.1,1.1,0,1.1)V. The last concern is on transistor M1, since M5 and M3 pull up its drain to 2.2V. However, the lowest voltage over M1 is still 1.1V since I1 keeps its source at VDD level: V_{DSGB}=(2.2,1.1,1.1,1.1)V. The illustrated behavior of the up level shifter cell holds similarly for the down level shifter cell, and the degradation mechanisms for

both cells have been simulated with industry standard aging simulation tools. It has been found that the lifetime of the level shifters did not exceed the standard requirements.

To enhance the operation, hot wells are used for transistors M1-M6. I3 and M7-M8 are placed in a triple well so that the NMOS body bias voltage equals V_{DD}. During voltage stacking off, the level shifter is in off state and the transistors M3-M8 receive 1x V_{DD} lower supply, so that M7 source is connected to V_{DD} while M5 drain to 0V, etc. The sizing of the transistors in the design is important, the I1N, I2N and M1-M2 need to be about 5x larger than the rest of the devices. The up- and down level shifters were implemented with high threshold voltage devices due to process limitations that did not allow multiple threshold voltage design.

D. Bypass

To ensure correct operation with voltage stacking on or off, the level shifters in Fig. 5 must be bypassed in the conventional power mode when voltage stacking is off. To achieve this, the scheme in Fig. 6 is proposed, which shows the implementation for up level shifters (a) as well as down level shifters (b). The signal path is split into two paths; the first path with a level shifter, and the second bypassing the level shifter. The two paths are selected with demultiplexers and multiplexers. The demultiplexers are realized with isolation cells that select and drive the path that is active in the given power mode (stacking on and stacking off). The multiplexer at the output of the level shifter then selects the active path and forwards it to the output.

The bypass circuit operates without external control signal. Staying at the example shown in Fig. 6 a, when voltage stacking is on, memory ground VSS TOP node is at V_{DD} voltage, and memory supply VDD TOP at $2V_{DD}$. This enables the AND isolation gate and disables the OR isolation gate, since one of their input is connected to VSS TOP. The OR gate output settles at 1.1V, not causing voltage overstress in the multiplexer. Therefore the level shifter is activated and produces an output signal between V_{DD} and $2V_{DD}$. The multiplexer receives the VDD node at its select input, which corresponds to a logic 0 since the ground node VSS TOP is at V_{DD}, the same voltage as VDD node. Therefore the 'SEL=0' input is selected which happens to be that of

the level shifter. When voltage stacking is turned off, memory ground VSS TOP node is at 0V, and memory supply VDD TOP at V_{DD}. This disables the AND isolation gate and enables the OR isolation gate, activating the bypass path and disabling the level shifter. The multiplexer selects the 'SEL=1' input since its ground voltage is 0V and the select input is hooked up to the VDD node. The output of the multiplexer copies the bypass signal.

E. Switched Capacitor Voltage Regulator

The task of the voltage regulator is to regulate the mid node that serves as the supply rail of the top power domain as well as the ground of the bottom power domain. It has to be sized for the worst-case current consumption so that the mid node is always kept in the desired voltage level.

Just like in [9], we chose a switched-capacitor voltage regulator (SCVR) since it provides superior efficiency compared to a linear regulator, and higher regulator efficiency means that there is less room for power saving for voltage stacking, making the comparison with conventional power delivery more realistic. Since the system is composed of two power domains stacked on top of each other, the voltage halving architecture was chosen for its simplicity, high efficiency and low area, shown in Fig. 7. The concept of the voltage halving SCVR is simple; a so-called flying capacitor swaps places between the input and the output. There are three main nodes present in the converter, the input (V_{IN}), output (V_{OUT}) and ground (V_{SS}). In phase 1, which is active corresponding to Φ_1, the capacitor is connected between V_{IN} and V_{OUT}, while in phase 2, it is connected between V_{OUT} and V_{SS}. Since the charge on the capacitor stored in each phase corresponds to the voltage present, over time there is a net charge transfer between V_{IN} and V_{OUT} which reaches equilibrium once the voltage on V_{OUT} is half that of V_{IN}. This charge transfer works both ways, thus positive and negative excess current both can be handled by the SCVR, making it suitable for voltage stacking.

The exact circuit implementation of the voltage halver is shown in Fig. 7.There is a clock signal needed to control phase 1 and 2, which needs to be level shifted for the switches that are connecting the flying capacitor

to V_{IN} and V_{OUT}. For this purpose, the level shifter discussed in Section II.C is used. Furthermore, the switches cannot have overlap since that would result in undesired short-circuit current, therefore the switch control signals that are derived from the clock signal need to be non-overlapping. To achieve this, the clock signal is converted to non-overlapping clock signal pair through a separate circuit block. The control signals in the top and the bottom part always have some latency compared to each other. Due to the synchronicity of the signals, the overlap margin was chosen to be high since the level shifter delay is hard to account for over several PVT corners. This high timing margin meant that the clock signal had a maximum frequency of 20MHz that could not be exceeded.

To benefit from high efficiency, the flying capacitor and the switches have to be sized so that the switching and the conduction losses are in balance and none of them limit the peak efficiency considerably. During the design stage, a total of 656pF flying capacitance was available in the form of 2x20.5pF accumulation PMOS capacitors per SCVR instance, and the maximum switching frequency was set to 25MHz. According to the Seeman model [14], this meant a minimum output impedance with ideal switches (slow switching limit) of about $R_{SSL}=1/4f_{sw}C\approx19.1\Omega$. Considering a maximum of 7.5% allowed average output voltage drop (82.5mV) and 3mA maximum output current, the maximum allowed output impedance is 27.5 Ω. Adding to this the effect of finite switch conductance, the goal was to stay under 25Ω output impedance, which indicated about 8.1Ω switch impedance based on [14]:

$$R_{FSL}=2R_{SW};\ R_{out}=\sqrt{R_{SSL}^2+R_{FSL}^2} \tag{2}$$

There were in total 16 SCVR instances implemented. To adjust for the output current requirements between stacking on and -off mode, 10 instances could be turned off for 6 active instances in stacked mode, while all the 16 instances were active when voltage stacking was off. Further, the efficiency is enhanced by not interleaving the regulator instances, similarly to the approach in [15]. This meant connecting the SCVR instances in a daisy chain configuration, one 6x and one 10x chain, as shown in Fig. 7. The reduction in

charge sharing by connecting the capacitor instances in series increased the SCVR efficiency at the cost of higher voltage ripple.

Just like for the level shifter cells, overvoltage has to be avoided for the devices, since they are implemented with traditional thin-oxide devices. To ensure proper operation upon startup, parallel to the power switches, startup helper circuits were implemented to bring the flying capacitor's voltages to a known state. This ensures that the thin-oxide flying capacitor is not exposed to higher voltage than 1.21V (V_{DD}+10%) which could cause overstress and degrade the device. To ensure proper operation, the power-up of the SCVR is performed in a stepwise fashion. First V_{IN} is brought to V_{DD} voltage, then V_{OUT} also receives V_{DD} voltage, finally V_{IN} is ramped up to $2V_{DD}$ voltage.

III. EXPERIMENTAL RESULTS

The testchip has been fabricated in a CMOS 40nm process. The micrograph can be seen in Fig. 8. The total area accounts for 1.49mm^2, of which the area of the bottom power domain is 0.077mm^2, the top power domain is 0.113mm^2, while the level shifter instances occupy 0.028mm^2.

Since voltage stacking claims to improve the voltage regulator efficiency, first it is important to look at the SCVR itself to characterize its behavior under various load conditions. Therefore, the SCVR efficiency was measured separately under different load currents. The results can be seen in Fig. 9. It is important to note that while with voltage stacking off, the output current I_{OUT} is always positive, thus it is sourced from the SCVR, in stacked mode the current can be either positive or negative. The sign is determined by the current consumption of the power domains. If the bottom power domain consumes more, then I_{OUT} will be positive, meaning that the useful power is consumed between V_{OUT} and V_{SS}. However if the top power domain is consuming more, I_{OUT} will take a negative sign as the current has to be sank by the converter. In this latter case the useful power is consumed between V_{IN} and V_{OUT}. As can be seen from the figure, the SCVR has similar efficiency profile for both current sink and –source scenarios, which is expected since the capacitors

are mostly symmetric for the charge transfer sign. The peak efficiency achieved in both cases was 81%. At low load current I_{OUT}, the switching losses of the converter start to dominate since they represent a load current independent portion of the losses. The switching losses depend on the switching frequency of the converter, which was fixed in this experiment. The measurements were at room temperature for a typical sample.

After characterizing the SCVR, the next step was to measure the current consumption that can be expected from the two power domains. Over 400 testcases were examined, where the current consumption of the top- and the bottom power domains were compared. The results in Fig. 10 show that the logic in extreme cases consumes 1.75x more current than the memory, while on average the mismatch is much smaller. The power consumption of the two power domains is well correlated since the MCU core has to read instructions from the memory every clock cycle, and there is no scenario where one of the power domains is inactive.

The timing and power overhead from the level shifting must be minimized to keep voltage stacking beneficial for various systems. In the current testchip the timing impact from the level shifting was 1.5ns added delay to the critical path between memory and logic, which permitted only 80MHz operation in stacked mode instead of the 100MHz that was used with stacking off. Further, the power penalty from the level shifters can be captured with the difference in the power drawn by the system from the voltage conversion stage with stacking on and off. It has been found that on average the output power of the conversion stage is 8% higher in stacked mode, which is partially the power overhead of the level shifting, however, the total power consumed by the entire system was still lower, as can be seen in the following, due to the higher efficiency of the conversion stage.

Characterizing the SCVR and the digital MCU part separately, the next step is to operate them simultaneously with voltage stacking on and off and quantify the benefits of voltage stacking. The system power efficiency was measured across the 400 testcases in Fig. 11. In stacked mode, on one hand, the SCVR needs to provide 4x less maximum output current, as well as the system achieves a 96% efficiency, a 15%

improvement over the SCVR peak efficiency of 81%. All this despite that the SCVR becomes less efficient at low load currents. The number of SCVR instances and their switching frequency was determined separately for stacking on and stacking off modes, keeping the criterion that their output voltage drop and efficiency has to be as much as possible aligned. After calibration, it was found that when stacking is off, using 16 SCVR instances at 10MHz resulted in similar output voltage and SCVR standalone efficiency as 6 SCVR instances at 5MHz in stacked mode.

Fig. 12 shows a reconfiguration scenario where the system is in transition between the two power modes. The supply and ground rails switch simultaneously so that the memory content is preserved and after re-enabling the clock, the program execution can be resumed. The clock has to be disabled before the supply rails change, and similarly, some minimum time is required before re-enabling the clock after the supply transition. The transition did not take longer than 0.5μs.

Zooming in to the supply waveforms in Fig. 13, the waveforms were compared for the case of voltage stacking on or off, under active load. With stacking off, 16 instances of the SCVR, switched at 10MHz, are used to process current $I_{BOT}+I_{TOP}$, while with stacking on, only 6 SCVR instances are turned on at 5MHz to process current $I_{BOT}-I_{TOP}$. Turning on voltage stacking resulted in a 3.4dB supply ripple reduction, and the voltage drop of the SCVR reduced from 58mV to 36mV, in accordance with [16]. Due to the smaller current that is processed in stacked mode, even weaker VR can produce cleaner supply, lifting the tough constraints on the power delivery system that is due to the scaled CMOS process.

Summarizing the results of this work, Table I shows a comparison with the previously reported works on voltage stacking. In the comparison this is the first IC that implements voltage stacking within a practical MCU system. Despite using bulk CMOS process, higher efficiency can be obtained using voltage stacking, compared to converters implemented with deep trench capacitors. Another important point to note is that no previous work features comparison of both conventional and voltage stacking power delivery for the same

system on the same die. Here we present a comparison for voltage stacking on and off modes. The efficiency has improved from 81% to 96%, while the power density increased threefold.

The benefit from voltage stacking can be further understood by comparing the efficiency and power density to other converters. In Fig. 14, the X axis shows the power density while the Y axis shows the efficiency of various SCVR implementation form bulk through SOI to processes that allow deep trench capacitors. It can be observed that, for the same technology, there is a trade-off between efficiency and power density; high efficiency converters typically have poor power density, and high power density converters have poor efficiency. In this work, while the same limitation applies for the implemented SCVR, by using voltage stacking, the trade-off can be bypassed and an improvement for both efficiency and power density can be achieved. In the comparison, only advanced processes can achieve the same kind of benefit that voltage stacking delivers in standard bulk CMOS.

IV. CONCLUSION

A microcontroller system with voltage stacking has been presented. For the first time, voltage stacking has been applied on a practical system. Unlike previous works, the implementation did not focus on the stacking of simple circuit blocks or larger, but independent systems, instead, a realistic system has been chosen as a demonstration vehicle for the benefits of voltage stacking. Furthermore, for the first time, voltage stacking could be turned off on the same system, making the benefits of voltage stacking directly quantifiable.

The benefits of voltage stacking were as follows. The system power efficiency improved from 81% with voltage stacking off to 96% with voltage stacking on, while using 5 times less effective power converter area. All this was achieved using bulk CMOS process with fully on-chip solution. Next to all this, though using a much weaker converter, the supply ripple has reduced by 3.4dB and the voltage drop from 58mV to 36mV. As future work, the technique can be extended to arbitrary system provided the partitioning of the design is automated.

ACKNOWLEDGEMENTS

The authors would like to thank Rina Lim for her help in realizing the level shifters.

REFERENCES

[1] P. Hazucha *et al*., "High Voltage Tolerant Linear Regulator With Fast Digital Control for Biasing of Integrated DC-DC Converters," in *IEEE Journal of Solid-State Circuits*, vol. 42, no. 1, pp. 66-73, Jan. 2007.

[2] W. Kim, D. Brooks and G. Y. Wei, "A Fully-Integrated 3-Level DC-DC Converter for Nanosecond-Scale DVFS," in *IEEE Journal of Solid-State Circuits*, vol. 47, no. 1, pp. 206-219, Jan. 2012.

[3] S. Rajapandian, Zheng Xu and K. L. Shepard, "Implicit DC-DC downconversion through charge-recycling," in *IEEE Journal of Solid-State Circuits*, vol. 40, no. 4, pp. 846-852, April 2005.

[4] T. M. Andersen *et al*., "A 4.6W/mm^2 power density 86% efficiency on-chip switched capacitor DC-DC converter in 32 nm SOI CMOS," *Applied Power Electronics Conference and Exposition (APEC), 2013 Twenty-Eighth Annual IEEE*, Long Beach, CA, 2013, pp. 692-699.

[5] W. J. Lambert, M. J. Hill, K. Radhakrishnan, L. Wojewoda and A. E. Augustine, "Package Inductors for Intel Fully Integrated Voltage Regulators," in *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 6, no. 1, pp. 3-11, Jan. 2016.

[6] S. Rajapandian, K. L. Shepard, P. Hazucha and T. Karnik, "High-voltage power delivery through charge recycling," in *IEEE Journal of Solid-State Circuits*, vol. 41, no. 6, pp. 1400-1410, June 2006.

[7] K. Blutman, A. Kapoor, A. Majumdar, J. Garcia Martinez, J. Echeverri, L. Sevat, A. van der Wel, H. Fatemi, J. Pineda de Gyvez and K. Makinwa, “A Microcontroller with 96% Power-Conversion Efficiency using Stacked Voltage Domains”, Proc. IEEE Symposium on VLSI Circuits, 2016

[8] K. Ueda, et al., "Low-Power On-Chip Charge-Recycling DC-DC Conversion Circuit and System," IEEE J. Solid-State Circuits, vol. 48, pp. 2608-2617, Nov. 2013.

[9] S. Lee, et al., "A 16-core voltage-stacked system with an integrated switched-capacitor DC-DC converter," IEEE Symposium on VLSI Circuits, pp. 318-319, June, 2015.

[10] Y. Liu et al., "A 0.1pJ/b 5-10Gb/s Charge-Recycling Stacked Low-Power I/O for On-Chip Signaling in 45nm CMOS SOI," ISSCC Dig. Tech. Papers, pp. 400-401, Feb., 2013.

[11] K. Blutman, A. Kapoor, J. G. Martinez, H. Fatemi and J. Pineda de Gyvez, “Lower Power by Voltage Stacking: A Fine-grained System Design Approach”, Proc. DAC, 2016, pp. 78:1-78:5.

[12] Z. Liu, L. Cong and H. Lee, "Design of On-Chip Gate Drivers With Power-Efficient High-Speed Level Shifting and Dynamic Timing Control for High-Voltage Synchronous Switching Power Converters," in *IEEE Journal of Solid-State Circuits*, vol. 50, no. 6, pp. 1463-1477, June 2015.

[13] Y. Moghe, T. Lehmann and T. Piessens, "Nanosecond Delay Floating High Voltage Level Shifters in a 0.35µm HV-CMOS Technology," in *IEEE Journal of Solid-State Circuits*, vol. 46, no. 2, pp. 485-497, Feb. 2011.

[14] H. Meyvaert, T. Van Breussegem and M. Steyaert, "A 1.65W fully integrated 90nm Bulk CMOS Intrinsic Charge Recycling capacitive DC-DC converter: Design & techniques for high power density,"*2011 IEEE*

Energy Conversion Congress and Exposition, Phoenix, AZ, 2011, pp. 3234-3241.

[15] B. Zimmer, et al., "A RISC-V Vector Processor with Tightly-Integrated Switched-Capacitor DC-DC Converters in 28nm FDSOI," IEEE Symposium on VLSI Circuits, pp. 316-317, June, 2015.

[16] R. Zhang, K. Mazumdar, B. H. Meyer, K. Wang, K. Skadron and M. R. Stan, "Transient voltage noise in charge-recycled power delivery networks for many-layer 3D-IC,"*Low Power Electronics and Design (ISLPED), 2015 IEEE/ACM International Symposium on*, Rome, 2015, pp. 152-158.

[17] L. Chang, et al., "A fully-integrated switched-capacitor 2:1 voltage converter with regulation capability and 90% efficiency at 2.3A/mm2," IEEE Symposium on VLSI Circuits, pp. 55-56, June, 2010.

[18] M. Steyaert, N. Butzen, H. Meyvaert, A. Sarafianos, P.Callemeyn, T. Van Breussegem and M. Wens, "DCDC performance Survey", [Online]. Available: http://homes.esat.kuleuven.be/~steyaert/DCDC_Survey/DCDC_PS.html

[19] Cabe, A.C.; Zhenyu Qi; Stan, M.R., "Stacking SRAM banks for ultra low power standby mode operation," in Design Automation Conference (DAC), 2010 47th ACM/IEEE , vol., no., pp.699-704, 13-18 June 2010

[20] Kesarwani, K.; Schaef, C.; Sullivan, C.R.; Stauth, J.T., "A multi-level ladder converter supporting vertically-stacked digital voltage domains," in Applied Power Electronics Conference and Exposition (APEC), 2013 Twenty-Eighth Annual IEEE , vol., no., pp.429-434, 17-21 March 2013

[21] Shenoy, P.S.; Sai Zhang; Abdallah, R.A.; Krein, P.T.; Shanbhag, N.R., "Overcoming the power wall: Connecting voltage domains in series," in Energy Aware Computing (ICEAC), 2011 International Conference on , vol., no., pp.1-6, Nov. 30 2011-Dec. 2 2011

[22] Mazumdar, K.; Stan, M., "Breaking the 3D IC power delivery wall," in Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on , vol., no., pp.741-746, 4-7 Nov. 2012

[23] Ehsan K. Ardestani, Rafael Trapani Possignolo, José Luis Briz, Jose Renau, " Managing Mismatches in Voltage Stacking with CoreUnfolding," TACO 12(4): 43 (2016)

[24] L. G. Salem, J. G. Louie and P. P. Mercier, "A flying-domain DC-DC converter powering a Cortex-M0 processor with 90.8% efficiency," *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, 2016, pp. 234-236.

[25] J. M. Wilson *et al*., "A 6.5-to-23.3fJ/b/mm balanced charge-recycling bus in 16nm FinFET CMOS at 1.7-to-2.6Gb/s/wire with clock forwarding and low-crosstalk contraflow wiring," *2016 IEEE International Solid-State Circuits Conference (ISSCC)*, San Francisco, CA, 2016, pp. 156-157.

[26] S. Zhang, J. S. Tu, N. R. Shanbhag and P. T. Krein, "A 0.79 pJ/K-Gate, 83% Efficient Unified Core and Voltage Regulator Architecture for Sub/Near-Threshold Operation in 130 nm CMOS," in *IEEE Journal of Solid-State Circuits*, vol. 49, no. 11, pp. 2644-2657, Nov. 2014.

[27] J. McClurg, Y. Zhang, J. Wheeler and R. Pilawa-Podgurski, "Re-thinking data center power delivery: Regulating series-connected voltage domains in software,"*Power and Energy Conference at Illinois (PECI), 2013 IEEE*, Champaign, IL, 2013, pp. 147-154.

[28] K. Onizuka, K. Inagaki, H. Kawaguchi, M. Takamiya and T. Sakurai, "Stacked-Chip Implementation of On-Chip Buck Converter for Distributed Power Supply System in SiPs," in *IEEE Journal of Solid-State Circuits*, vol. 42, no. 11, pp. 2404-2410, Nov. 2007.

Table Caption

Table I. Comparison of voltage stacking and other SCVR implementations

Figure Caption

Fig. 1. Power modes of the testchip. a) 'Voltage Stacking Off' mode b) 'Voltage Stacking On' mode

Fig. 2. Efficiency improvement calculation for voltage stacking scheme compared to conventional power delivery approach

Fig. 3. System level block diagram with the power domains, SCVR and IO pads

Fig. 4. Power composition of the digital part of the system. Memory and logic power well balanced

Fig. 5. Level shifter cells used to translate the signals between memory and logic

Fig. 6. Level shifting scheme with bypass. Signal path is selected according to the power mode.

Fig. 7. Block diagram of the 2:1 SCVR used to regulate the mid node between memory and logic

Fig. 8. Chip micrograph of the voltage stacked microcontroller

Fig. 9. SCVR efficiency for positive and negative load current

Fig. 10. Characterization of current profiles for memory and logic demonstrating low mismatch

Fig. 11. Efficiency improvement through voltage stacking compared to conventional power delivery

Fig. 12. Clock and memory supply waveforms while turning voltage stacking on and off

Fig. 13. Supply quality comparison with voltage stacking on and off

Fig. 14. Power conversion figure of merit [18] improvement through voltage stacking

Table I Comparison of voltage stacking and other SCVR implementations

	[6]*	[9]**	[8]*	[17]	[15]	This Work (flat)	This Work (stacked)
Circuit	Stacked Multipliers	Stacked MCU cores + mem	Stacked PLLs	Switched-Capacitor VR	MCU	**MCU**	**Stacked MCU**
Technology	180nm	40nm	90nm	45nm SOI	28nm FDSOI	**40nm**	**40nm**
Area [mm^2] (System + VR)	0.47 + 0.044	1.79 + 1.08	? + 0.032	0.0012 (VR)	1.19 + 0.19	**0.23 + 0.27**	**0.23 + 0.10**
Converter	Linear Regulator	Switched Capacitor	Linear Regulator	Switched Capacitor	Switched Capacitor	**Switched Capacitor**	**Switched Capacitor**
Stacked?	Yes	Yes	Yes	No	No	**No**	**Yes**
Conversion Voltage [V]	3.6 to 1.8	3.6 to {2.7 1.8 0.9}	Several	2 to 0.95	1.8 to 0.9, {0.67 1} to 0.5	**2.2 to 1.1**	**2.2V to 1.1**
Frequency [MHz]	-	adaptive	-	100	55-498 (adaptive)	**10-20**	**5-10**
Converter Efficiency	? (<50%)	75%	44%	90%	70-80%	**81%**	**81%**
Peak System Efficiency	93%	99%	87%	-	86%	**81%**	**96%**
Power Density [mW/mm^2]	1636	147	1500	2185	350	**7**	**21**

*** Stacked circuit blocks like PLLs, multipliers only**
**** Stacked, independent MCUs only**

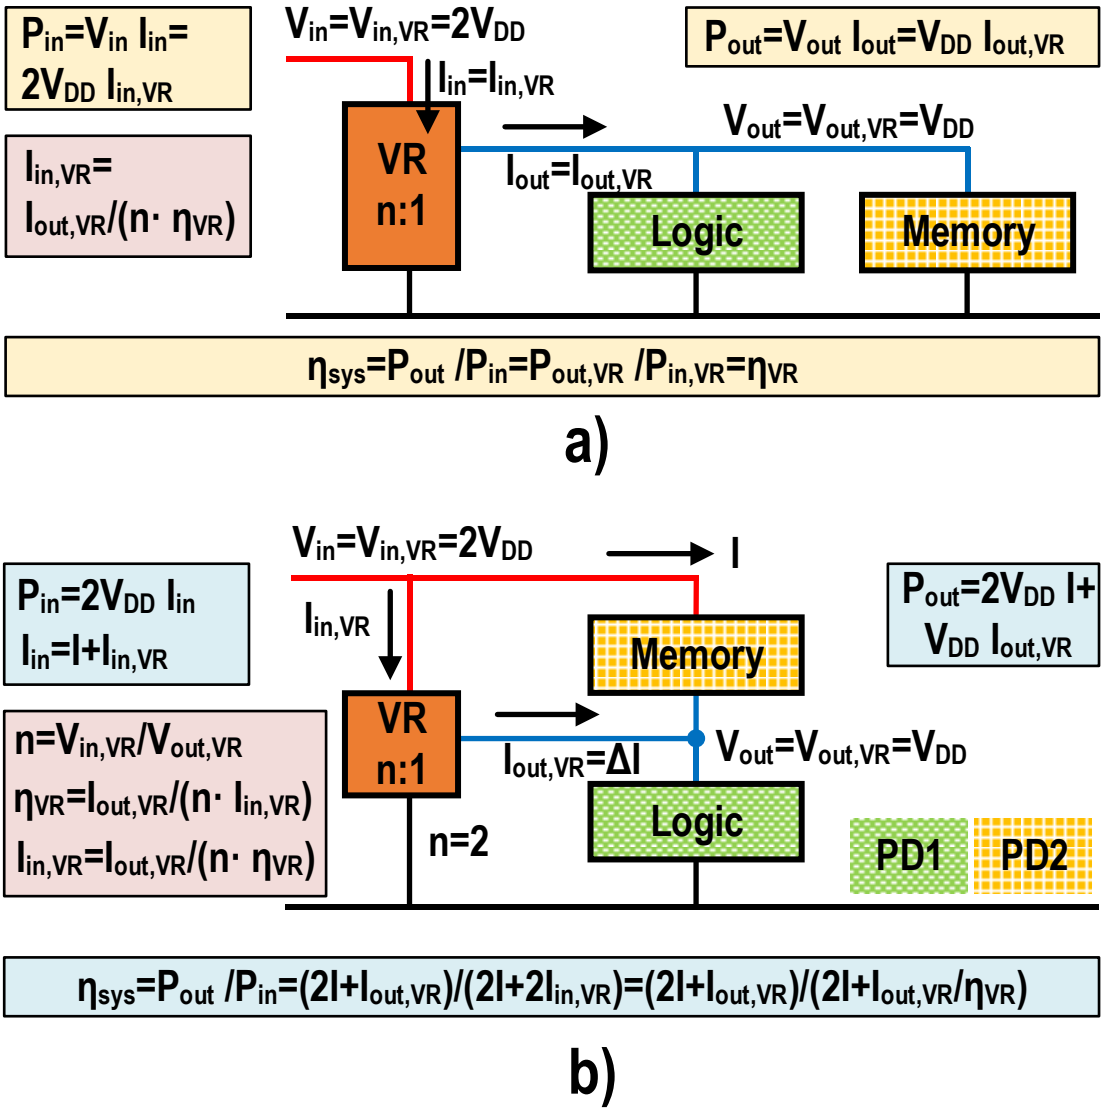

Fig. 1 Power modes of the testchip. a) 'Voltage Stacking Off' mode b) 'Voltage Stacking On' mode

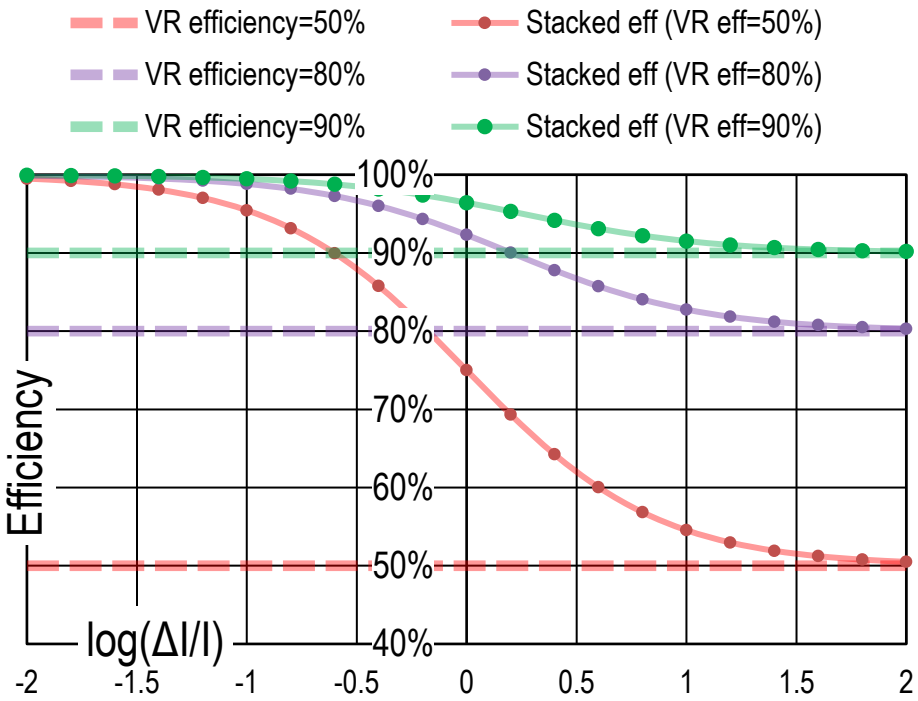

Fig. 2 Efficiency improvement calculation for voltage stacking scheme compared to conventional power delivery approach

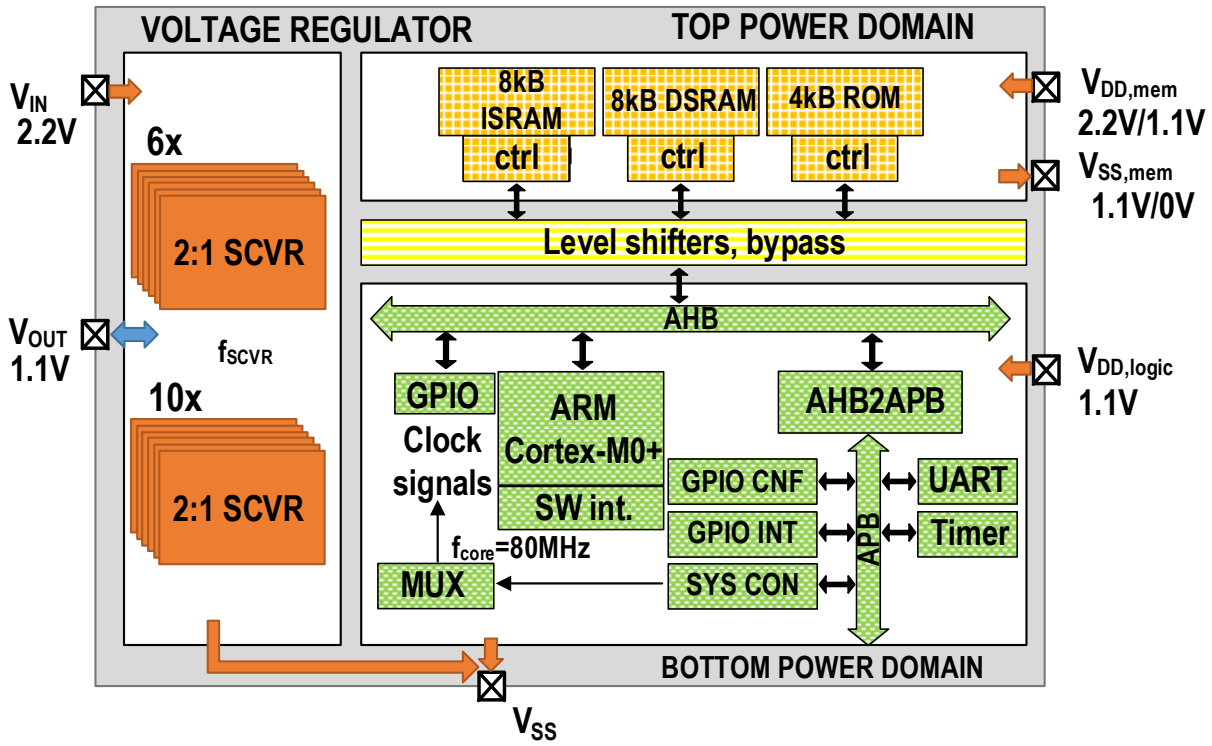

Fig. 3 System level block diagram with the power domains, SCVR and IO pads

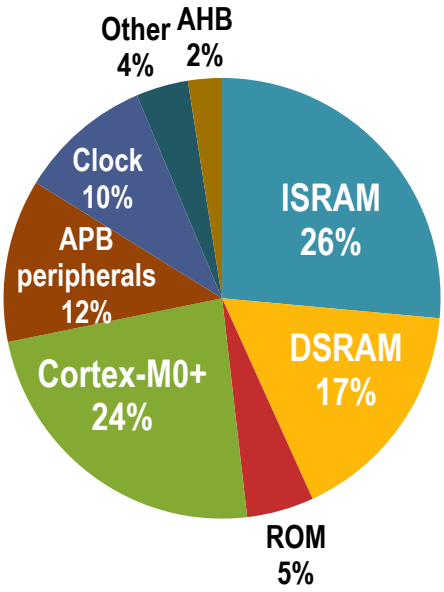

Fig. 4 Power composition of the digital part of the system. Memory and logic power well balanced

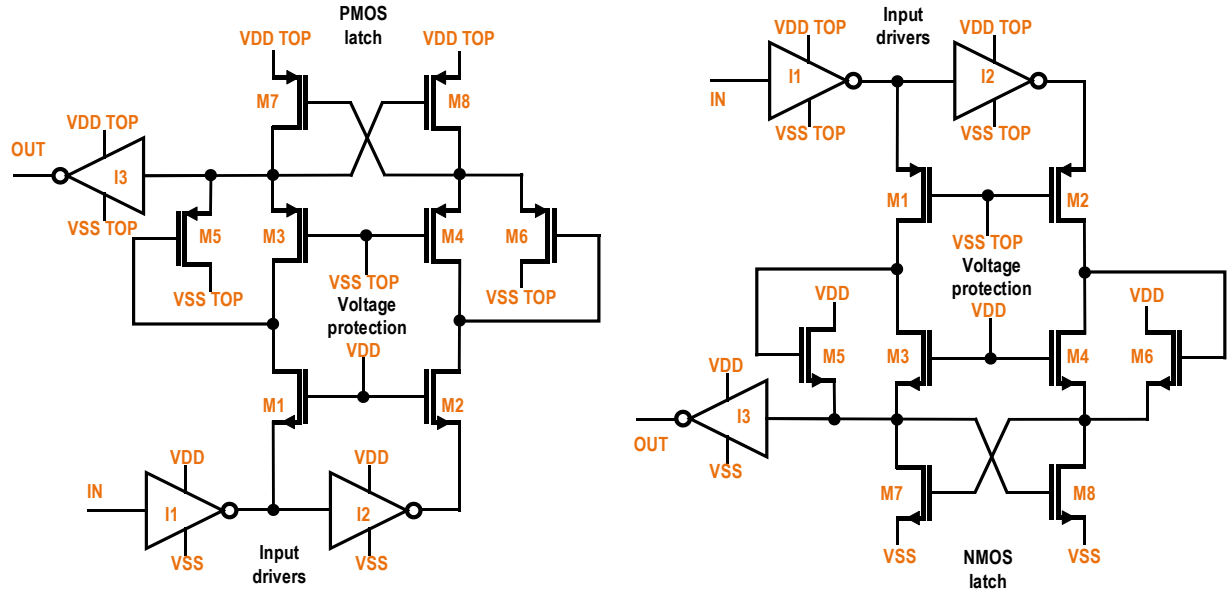

a) Up level shifter

b) Down level shifter

Fig. 5 Level shifter cells used to translate the signals between memory and logic

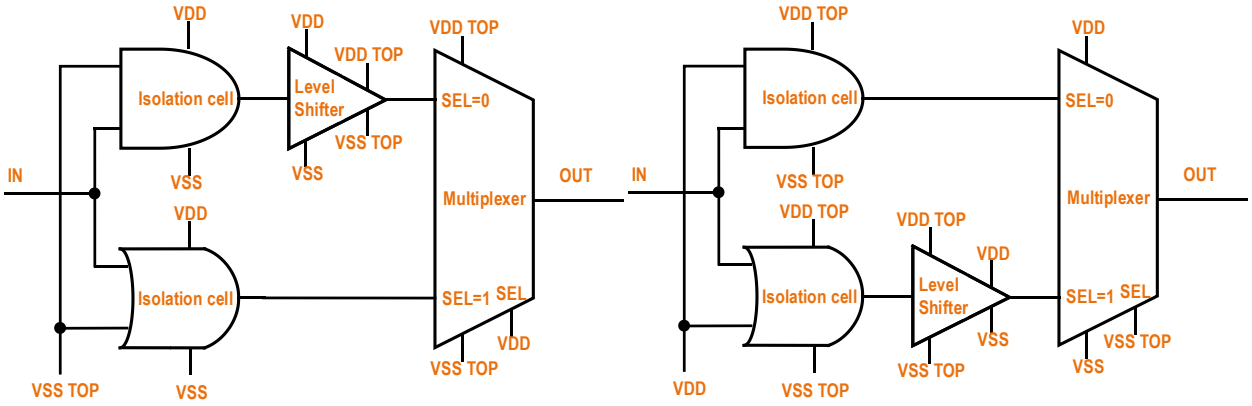

a) Up level shifting scheme

b) Down level shifting scheme

Fig. 6 Level shifting scheme with bypass. Signal path is selected according to the power mode.

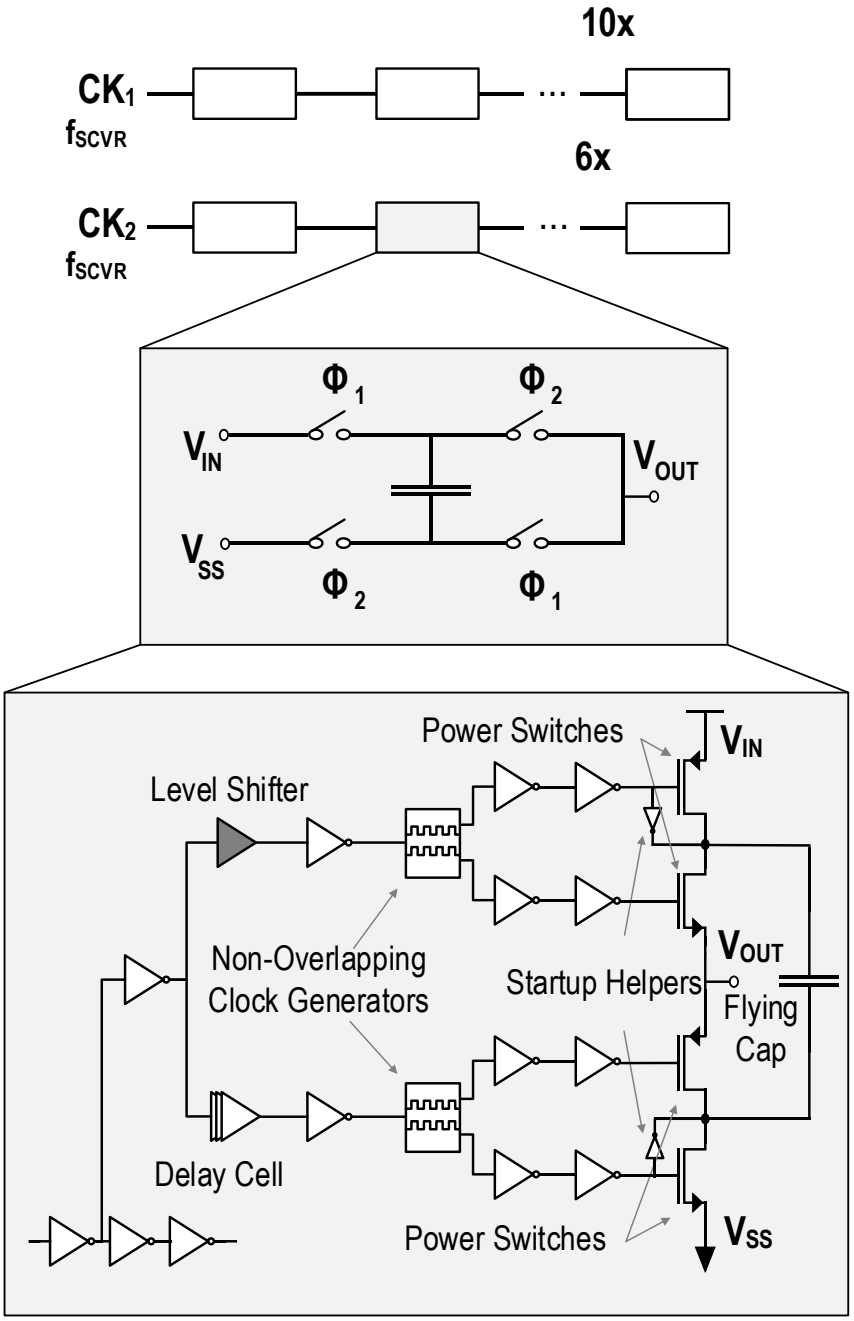

Fig. 7 Block diagram of the 2:1 SCVR used to regulate the mid node between memory and logic

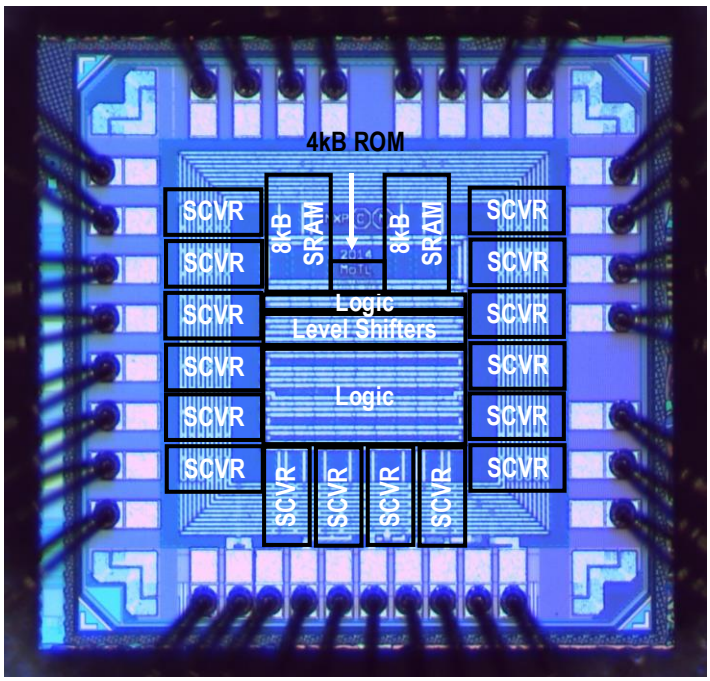

Fig. 8 Chip micrograph of the voltage stacked microcontroller

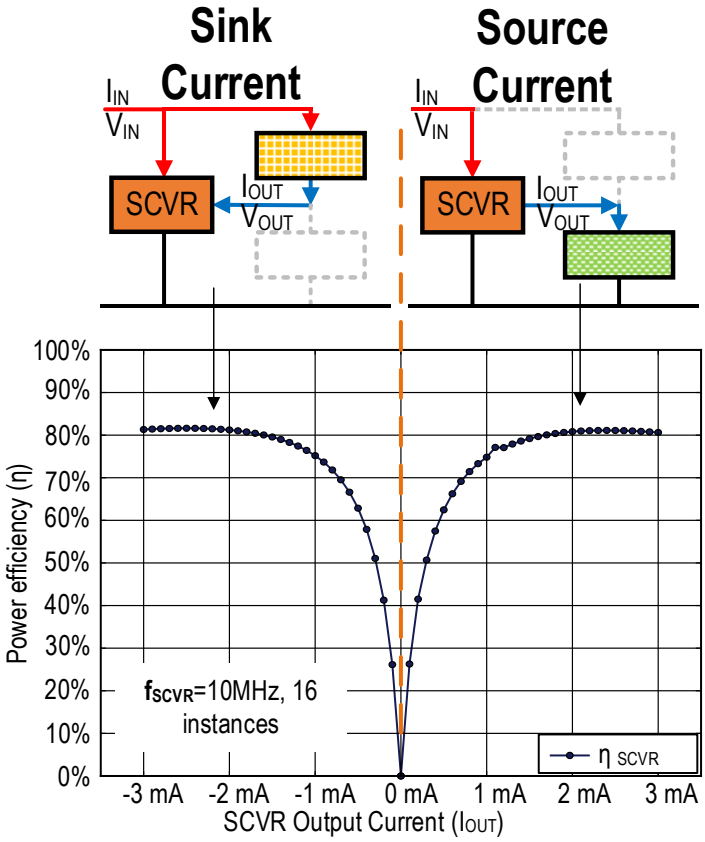

Fig. 9 SCVR efficiency for positive and negative load current

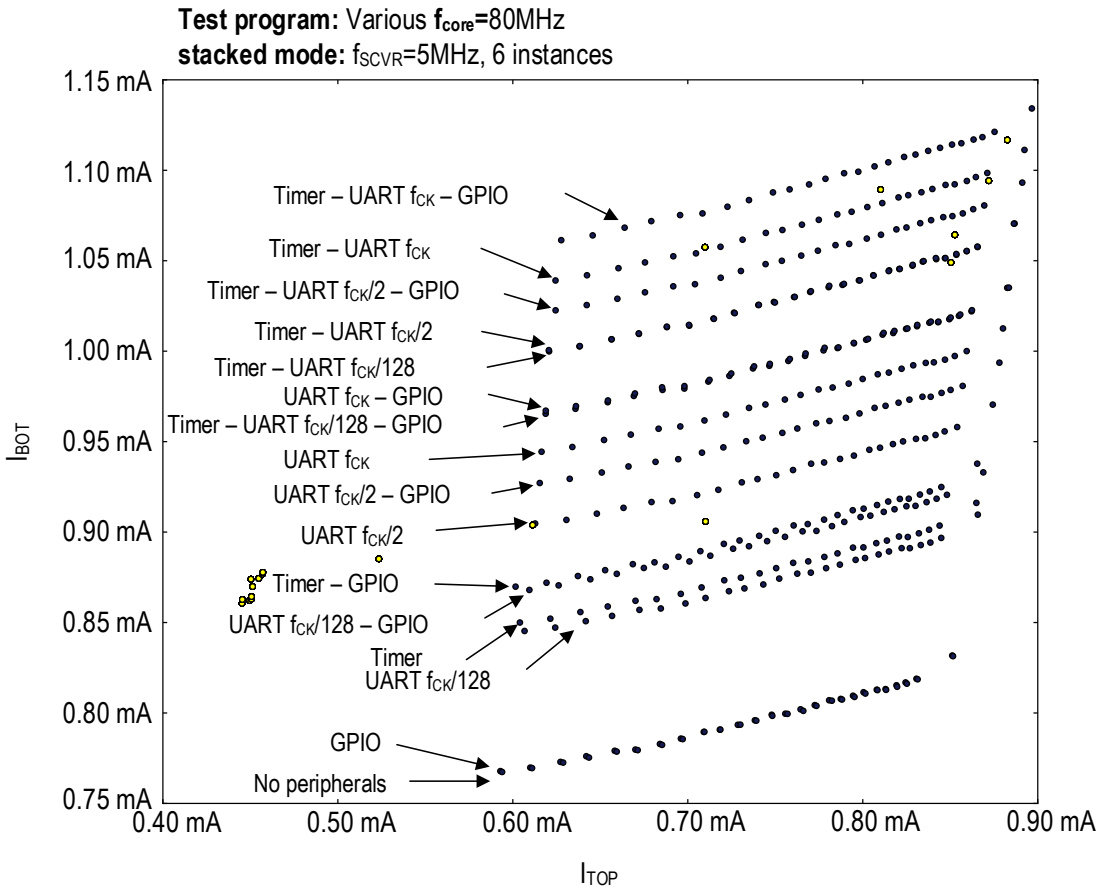

Fig. 10 Characterization of current profiles for memory and logic demonstrating low mismatch

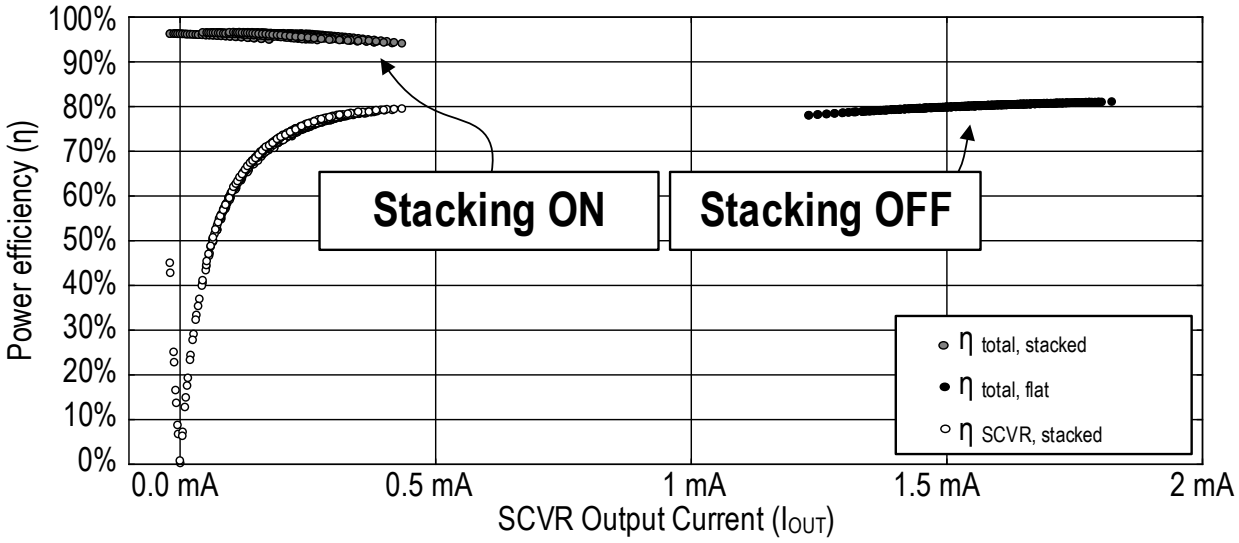

Fig. 11 Efficiency improvement through voltage stacking compared to conventional power delivery

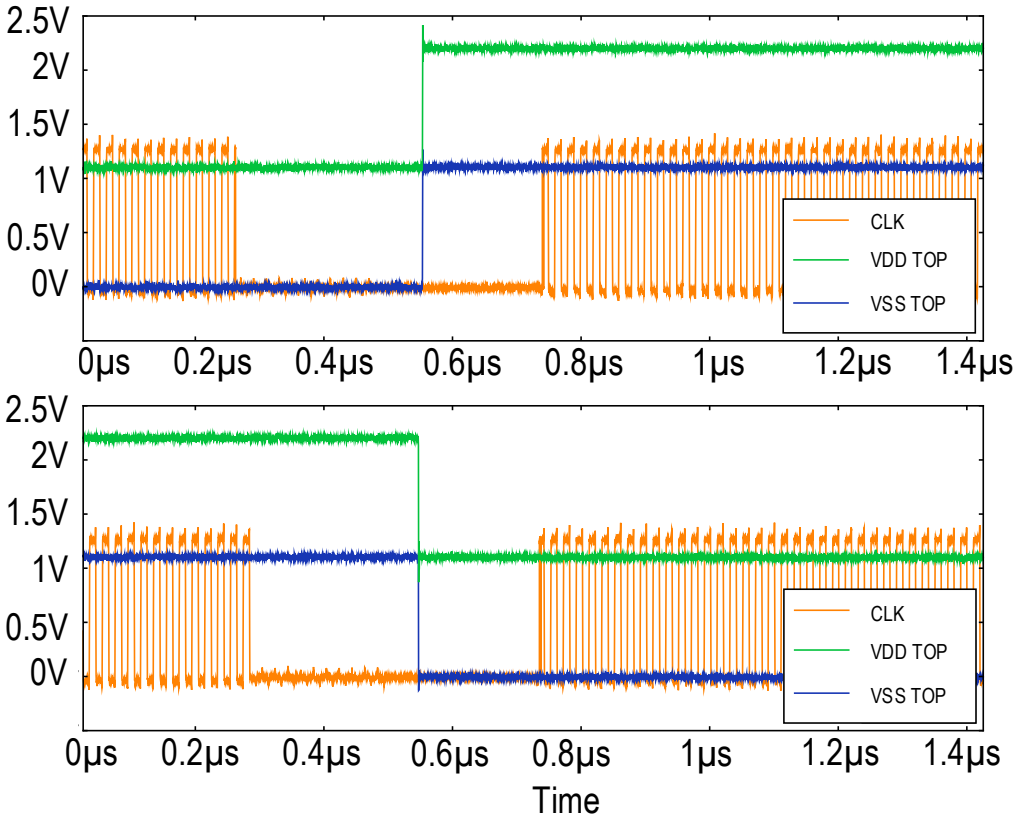

Fig. 12 Clock and memory supply waveforms while turning voltage stacking on and off

Test program: FFT
$\mathbf{f_{core}}$=80MHz

Stacking OFF: f_{SCVR}=10MHz, 16 instances
Stacking ON: f_{SCVR}=5MHz, 6 instances

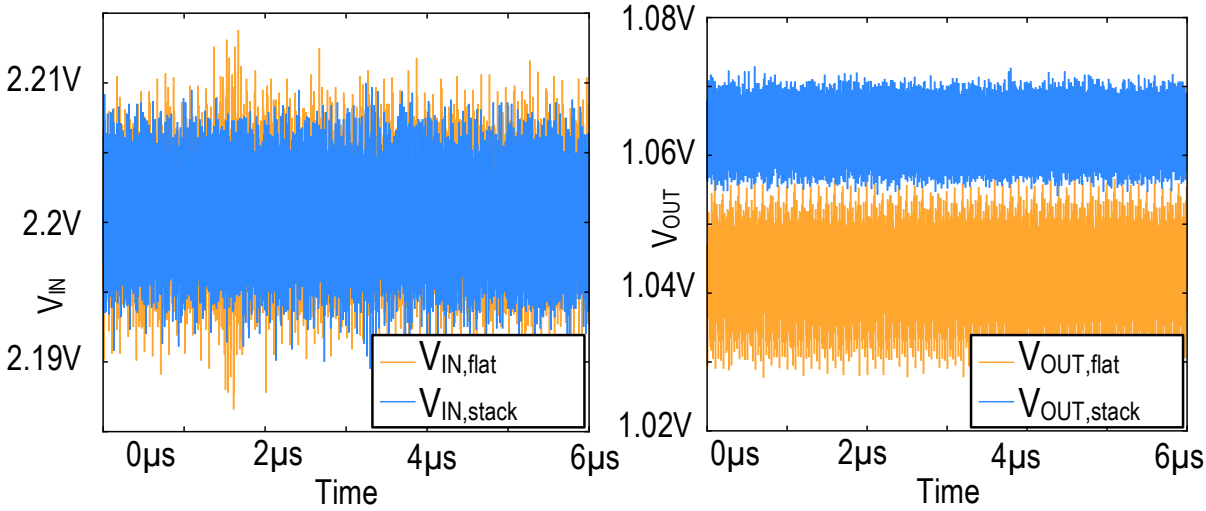

Fig. 13 Supply quality comparison with voltage stacking on and off

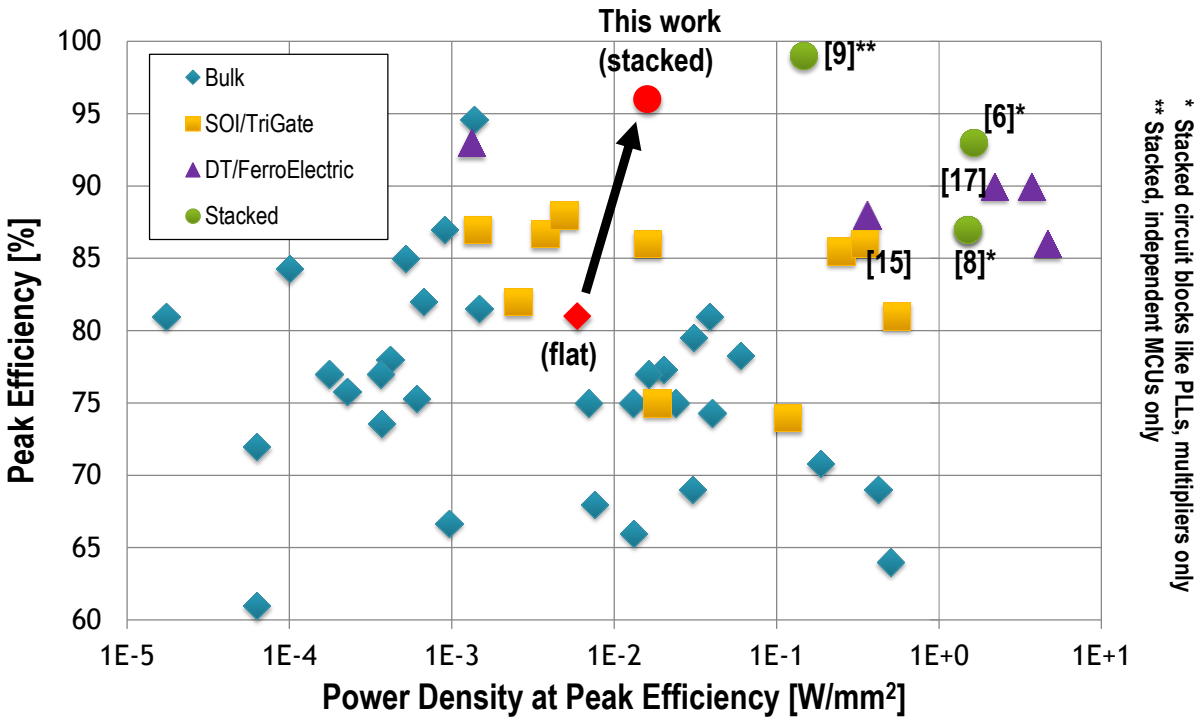

Fig. 14 Power conversion figure of merit [18] improvement through voltage stacking