# Proactive Power Regulation with Real-time Prediction and Fast Response Guardband for Fine-grained Dynamic Voltage Droop Mitigation on Digital SoCs

Xi Chen*, Jiaxiang Feng*, Aly Shoukry*, Xin Zhang[#], Raveesh Magod[+], Nachiket Desai**, Jie Gu*
*Northwestern University, [#]IBM, [+]Texas Instruments, **Intel,

**Abstract** - A proactive power regulation scheme for mitigating dynamic supply droop is proposed with a fully-integrated buck converter, a CPU core and a real-time machine learning (ML) engine. Combined with droop guardband circuits, the proactive scheme demonstrates up to 9.9% higher CPU frequency or 9.2% higher power efficiency compared with prior fast LDO scheme or conventional converters using a 65nm test chip.

## Introduction

Power integrity has become a major challenge in advanced CMOS with sub-1V supply voltage and drastically varying workload in a digital SoC. The conventional power management approaches engaging regular control loop of power converters are too slow for highly dynamic events such as sudden change of workload, resonant droop, or instruction specific power surges, causing excessive voltage margin [1]. For highly dynamic supply droops, fast digital LDOs or event-based LDOs were designed to react promptly to the workload variation and provide on-demand regulation [2, 3]. As shown in Fig. 1, a speed and efficiency tradeoff are observed where digital LDO offers sub-ns regulation with low efficiency while the DC-DC converters suffer from slow response. Recently, "proactive" clock throttling techniques were developed on both DSP and high-end processors where instruction-based prediction led to 10% frequency improvement [4, 5]. However, clock throttling incurs performance penalty and does not fundamentally remove the adversarial effects of supply droop. This work demonstrates proactive power management for modern SoC with fully integrated power converters and digital microprocessors. As in Fig.1, the current consumption of a microprocessor observes significant variation at cycle-by-cycle bases. Three dimensions of dependency need to be captured including (1) instructions being executed, (2) instruction sequence, (3) real-time traces of supply voltage to predict supply events leading to complex computational models. To deal with these challenges, this work demonstrates a comprehensive solution for proactive power regulation using a real-time ML engine for droop prediction and a fast power converter with "safety" guardband. The proactive techniques allow actions to be taken before the happening of real events reducing the fundamental speed limitation of power converters as demonstrated by a 65nm test chip.

## Proactive Power Regulation

Fig. 2 shows the overall chip architecture of the proposed scheme. A RISC-V CPU at about 800MHz from a nominal supply of 1.0V to 1.2V is powered by a fully integrated buck converter which includes a feedback control loop at 10~50MHz and a fast PWM generator with PWM signals at 600~1GHz for two-phase interleaved power cores. The power cores integrate on-chip spiral inductors and convert an IO input voltage of 1.8V to CPU's supply voltage between 0.6V and 1.4V. A special ML core is used to generate prediction of supply droop and issue power regulation through a fast PWM modulation (FPWMM) module. Short-term guardband circuits are included as "safety net" for ML's false prediction while long-term guardband deals with larger and slower droop. For testing, a cycle-by-cycle voltage recorder (VR) with 2k depth is implemented to record real-time voltage droop.

Fig. 3 shows the detailed design of the real-time ML core. The fetched instructions from RISC-V CPU are sent into the ML core for prediction providing 3~4 cycles lead time before the high-power stage at the execution or memory write back stages. For power dependency on instruction and sequence, the past three decoded instructions are kept in an instruction queue. The instruction information is further grouped into clusters of features, e.g. ALU activities, memory load/store activities, etc. with 52 1-bit real-time "feature" signals as input to a linear regression core to predict the upcoming current of the CPU. The regression result is then sent into a decision tree realized by a lookup table for combining with voltage information. The control command based on current prediction and real-time voltage is sent to buck converter for power regulation. The use of 1-bit feature and 4-bit stationary weight allows a multiplier-less regression with 16X reduction of power and minor accuracy loss compared with a 2-layer neural network. As shown in Fig. 3, a combined true positive rate (TPR) of 92.2% or false positive rate (FPR) of 4.7% for an undershot is achieved for an example benchmark program. The ML core incurs 2.4~3.5% power overhead to the CPU but achieves power saving due to higher efficiency from regulators as discussed later. The ML model is trained offline by a joint analog and digital simulation framework and is further adjusted based on silicon characterization data as described in Fig. 4.

Fig. 4 shows the main circuits inside the buck converter supporting ML operation. The FPWMM regulates the inductor current by modulating the original PWM signals within one PWM clock cycle (1~2ns) using time-domain programmable pulse "stretcher" and "trimmer" based on commands issued from ML core. To deal with the occasional misprediction of ML, a short-term droop guardband (SGB) is implemented near the processor to bring the voltage back immediately with a total transient delay of less than 300ps. As SGB behaves like digital LDO, it should be activated as little as possible to reduce efficiency loss demanding high accuracy of ML prediction. Long-term droop guardband (LGB) controlled by event-based scheme similar to [3] is deployed to deal with the slower but larger power change, e.g., sudden CPU's wakeup, which is beyond the fine-grain regulation from ML predictions. Additional PWM offset inversely proportional to threshold crossing time is issued in LGB through a fast feedforward path when supply droop is not recovered by ML/SGB within 2 clock cycles. LGB reacts within ~5 CPU clock cycles which is much faster than the main regulator loop at 10~50MHz.

## Measurement Results

A test chip was fabricated in a 65nm process. The CPU functionality under dynamic supply droops were verified by running benchmark programs and scanning out all internal register files and caches. Multiple benchmark programs have been run on the SoC with performance and ML accuracy evaluated. Fig. 5 shows examples of the measured supply waveforms captured from voltage recorder in comparison with simulation results on StringSearch and CoreMark programs. Three modes of operations were recorded including (1) baseline without ML or SGB, (2) operation with ML core but without SGB, (3) operation with both ML core and SGB.

Measurements from voltage recorder show a droop reduction up to 100mV when both ML and SGB was used. The SGB was triggered only at 1% of the time (about 10 out of 1k cycles) providing safety guardband for ML with negligible power loss. Measurement on LGB shows a 5-cycle recovering from a power surge of 5mW to 150mW. Fig. 5 also shows the measured converter efficiency under the CPU workload and the frequency improvement from droop reduction. The proactive power regulation results in 6.2~9.9% higher CPU frequency due to droop mitigation from 0.6~1.2V. The proactive approach enabled the use of slower buck converters with 5.1~9.2% higher power efficiency compared with digital LDO for providing similar cycle-level droop mitigation using prior fast LDO methods [2]. Fig. 5 also shows a comparison with conventional buck converters by simulating this design w/o ML prediction at 50MHz (w/ larger L) matching that with common buck converters. A ~100mV supply noise (average 52mV VDD) reduction or equivalent 9.1% power saving is observed showing the benefit of the proactive scheme over conventional converters. Fig. 6 shows the impact of prediction accuracy on power efficiency. When TPR is lower than ~60%, the efficiency dropped by 6~9% due to frequent activation of SGB approaching conventional "reactive approach". The five benchmark programs all achieve more than 80% TPR with negligible power loss from SGB. Fig. 6 also shows the die micrograph, efficiency of the buck converter and comparison with prior droop mitigation schemes. Compared with clock-based schemes [4, 6], this work removes the impact of supply droop. This is the first work utilizing proactive technique for droop mitigation overcoming the efficiency and speed tradeoff of conventional LDO and power converters.

**Acknowledgements** This work is supported in part by SRC task 2810.039 and NSF CCF-1846424.

**References:**

[1] Bo-Jr Huang, et al. *ISSCC*, 2021 [2] Z. K. Ahmed et al. *VLSISymp*, 2019 [3] S. J. Kim et al. *VLSISymp*, 2019 [4] V. K. Kalyanam et al. *VLSISymp*, 2020 [5] A. Agrawal et al. *ISSCC*, 2021 [6] A. Nayak et al. *ISSCC*, 2022.

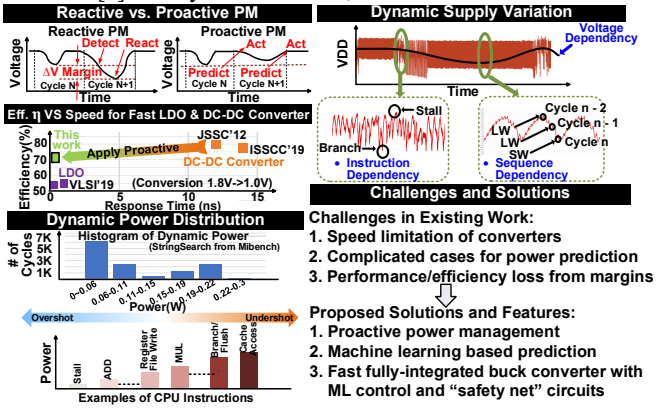


Fig. 1: Concept of proactive power regulation and challenges.

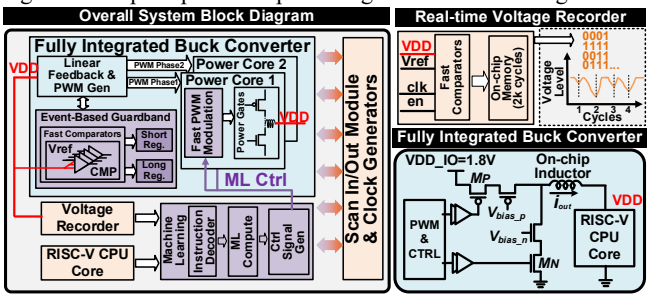


Fig. 2: Top-level block diagram of the test chip.

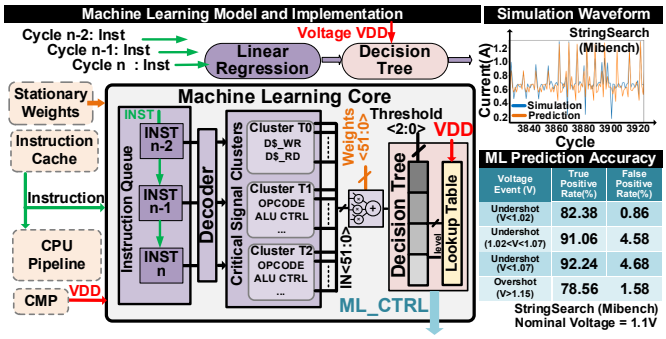


Fig. 3: Architecture of the ML core and accuracy.

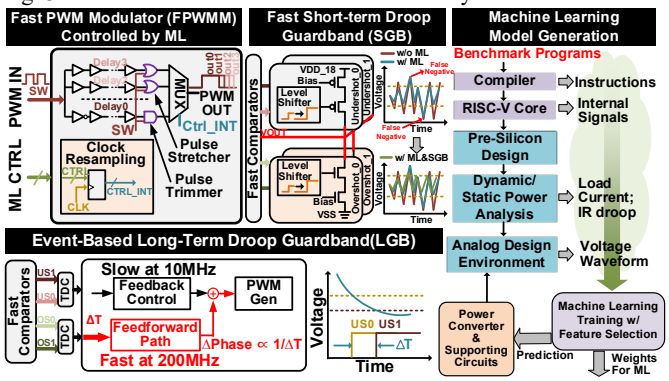


Fig. 4: Fast PWM modulation, short-term droop guardband and long-term droop guardband (left) and ML design flow (right).

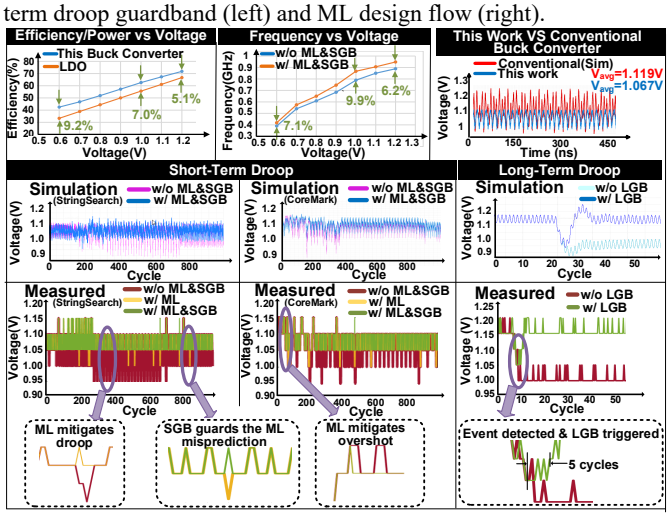


Fig.5: Measured supply droop performance, frequency, and efficiency benefits across VDD.

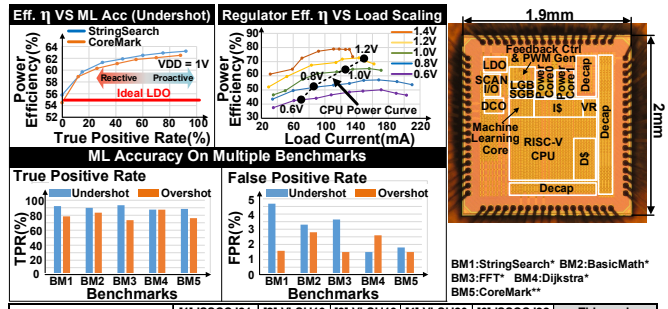


| | [1] ISSCC '21 | [2] VLSI '19 | [3] VLSI '19 | [4] VLSI '20 | [6] ISSCC '22 | This work |
|---|---|---|---|---|---|---|
| Process | 7nm | 22nm | 65nm | 7nm | 5nm | 65nm |
| Processor | CPU/GPU/APU | N/A | N/A | DSP | ARM CPU | RISC-V CPU |
| Mitigation Scheme | DVFS w/ PVT Sensors | Compute LDO | Event LDO | Proactive Clock | DVFS w/ FLL/PLL | Proactive Power Management |
| Vin (V) | N/A | 0.55-1.2 | 0.5-1.0 | N/A | N/A | 1.8 |
| Vout (V) | N/A | 0.5-1.15 | 0.45-0.9 | N/A | N/A | 0.6-1.4 |
| Response Time | 1ms | 1ns | 0.2ns | A few Clk Cycles | μs to ms | 3~4 Clk Cycles for ML *** |
| Computing Method | Software Machine | Droop Speed | Event Detection | Equation Based | N/A | On-chip Machine Learning |
| Performance or Power Improvement | 13% | N/A | N/A | 10% | 4% | 9.9% |

*Mibench **EEMBC ***3~4 cycles after prediction is made; 0.3ns for SGB

Fig. 6: Die micrograph, more benchmark accuracy, efficiency vs. prediction accuracy, regulator efficiency and comparison table.