

MEMTI: Optimizing On-Chip Nonvolatile Storage for Visual Multitask Inference at the Edge

Marco Donato, Lillian Pentecost,
David Brooks, and Gu-Yeon Wei
Harvard University

Abstract—The combination of specialized hardware and embedded nonvolatile memories (eNVM) holds promise for energy-efficient deep neural network (DNN) inference at the edge. However, integrating DNN hardware accelerators with eNVMs still presents several challenges. Multilevel programming is desirable for achieving maximal storage density on chip, but the stochastic nature of eNVM writes makes them prone to errors and further increases the write energy and latency. In this article, we present MEMTI, a memory architecture that leverages a multitask learning technique for maximal reuse of DNN parameters across multiple visual tasks. We show that by retraining and updating only 10% of all DNN parameters, we can achieve efficient model adaptation across a variety of visual inference tasks. The system performance is evaluated by integrating the memory with the open-source NVIDIA deep learning architecture.

■ **IN RECENT YEARS**, deep neural networks (DNNs) have become essential to tasks across application domains, including image recognition

and detection, language processing, and translation. This increase in popularity, together with the continued proliferation of low-power embedded devices, has motivated the design of DNN-specific hardware accelerators.¹ While many energy-efficient DNN hardware implementations have been proposed, a major challenge remains: the large memory requirement to store DNN

Digital Object Identifier 10.1109/MM.2019.2944782

Date of publication 4 October 2019; date of current version 8 November 2019.

parameters. Although entirely on-chip storage would guarantee better inference performance, limited on-chip SRAM capacity inevitably leads to reliance on costly off-chip memory accesses to DRAM.

Embedded nonvolatile memories (eNVMs) provide higher density than SRAM and can ameliorate the need for power-hungry DRAM storage. However, the benefits of eNVMs come at the cost of larger write energy and write latency. Moreover, limited eNVM write endurance is an obstacle to the adoption of certain technologies if DNN parameter values require frequent updates. For instance, embedded devices for robotics or augmented reality applications often required a combination of multiple inference tasks, including image classification, object detection, and action recognition. These cases highlight the need for scalable solutions that can flexibly accommodate DNN parameters for multiple tasks.

We present a DNN model and memory code-sign solution that leverages a multitask learning technique to reduce eNVM writes, while enabling systems to efficiently perform multiple inference tasks. Maximizing the reuse of the learned parameters across different DNN-dependent vision tasks without retraining enforces the assumption of infrequent writes: parameters shared by multiple tasks are trained and written only once, and therefore are highly suitable for eNVM storage; in contrast, the remaining parameters can be retrained to accommodate new inference tasks, and stored in SRAM. In addition to the storage density benefits, we evaluate how the process of retraining specific parameters can be used to recover from accuracy loss due to the adoption of denser, fault-prone multilevel eNVM storage. This article provides the following contributions.

- Leverage residual adapters to optimize parameter storage in dense eNVMs.
- Evaluate application accuracy with quantization and multilevel cell (MLC) RRAM faults when a majority of DNN parameters is shared across inference tasks.
- Quantify the system-level performance and energy advantages of a multitask-enabled

deep learning architecture (NVDLA) integrated with optimized eNVM solutions.

DNN AND MEMORY CODESIGN

Generalizing deep learning architectures to enable different application domains and more varied inference tasks serves as a way of supporting more powerful and versatile models. For example, the work by Kaiser *et al.*² combines several building blocks for translation, speech, and visual inference that can be trained on all desired tasks simultaneously or on each task separately. In either case, however, introducing new inference tasks would require updating the entire set of model parameters. Other works have leveraged the concept of transfer learning to improve the performance of a single DNN on different data sets. These approaches are based on the observation that many visual inference tasks share low-level features, such as edge and shape detection, in the front-end layers, and become more task-specific as the computation moves closer to the classification layers. However, in order to preserve inference accuracy, transfer learning approaches either share only a limited number of front-end layers or fine-tune parameters by retraining the transferred features from one inference task to another.³ A recent proposal applies transfer learning to create a synthesizable fixed-parameter feature extractor.⁴ However, hardwiring the feature extractor in logic prevents from fine-tuning the parameters, limiting the amount of cross-task weight sharing.

While all these techniques enable a single DNN model to perform different inference tasks, they still require updating a considerable portion of parameters to achieve maximum adaptation. We pursue a specific transfer learning technique for which the learned parameters can be generalized across multiple vision inference tasks by maximizing DNN parameter reuse and enabling efficient inference on embedded devices. The high degree of DNN parameters reuse reduces memory traffic requirements, which makes nonvolatile memories a compelling solution for retaining shared parameters on-chip without incurring costs associated with frequent memory writes.

Multitask Learning Model

Our design is based on the DNN architecture presented by Rebuffi *et al.*,⁵ which uses residual adapter modules as a way to parameterize a generic ResNet network. These parametric modules are themselves residual blocks which use 1×1 filters and skip connection. In this setting, the number of domain-specific parameters, which comprises adapter filters, batch normalization, and fully connected classifier parameters, can be reduced to roughly 10% of the total model size. For our experiments, we integrate the residual adapter modules in a ResNet26 network.

The baseline network is pretrained on ImageNet, which is standard practice in transfer learning and model fine-tuning techniques. The pretrained version for ImageNet achieves top-1 accuracy of 67.65%. The ResNet26 weight parameters obtained during pretraining are the backbone of this multitask inference system as they are reused for running inference on any additional visual task. The degree of adaptation is tested against five data sets, which have been selected to be representative of popular image processing tasks including classification (cifar100, aircraft), object detection (German Traffic Signs, Daimler pedestrian classification), and action recognition (UCF101 Dynamic Images).

Table 1 summarizes the best accuracy in the case of the model being either trained entirely from scratch or only for the task-specific parameters. As anticipated, for all data sets, the adapters overhead is around 10%. The accuracy of the network trained using adapters is always better than or comparable to training the entire network independently for each data set. In addition, we observe that the modified model converges to the best accuracy in fewer training epochs, which results in training speedup reported in Table 1.

Nonvolatile Memory Technologies

The landscape of nonvolatile memories includes a wide range of emerging technologies.⁶ These memories are generally characterized by high energy efficiency and high storage density, which can be further increased by programming multiple levels in a single cell. We label this

storage solution as MLC storage, in contrast to single-level cell (SLC) storage, for which each eNVM cell stores a single binary value. In this article, we focus on a specific eNVM implementation, namely RRAM. Various implementations such as phase-change memories, embedded flash, or ferroelectric memories can also be used for MLC storage. On the other hand, STT magnetic memories (STT-MRAM), while having the best write and read performance,⁶ are not a suitable candidate because compelling MLC implementations with comparable density have not been demonstrated to date.

There are alternative implementations with varying advantages and limitations. For example, the storage requirements for the DNN architecture we are leveraging could be met by read-only memories (ROMs) as well. ROMs ensure the best density for storing the shared parameters, however, they also require configuring the network at fabrication time, which makes the design less scalable and cost-effective. One-time programmable memories such as antifuse, while being amenable to postfabrication configuration, are far less dense than other memory solutions, even when compared to SRAM.⁷ Previous work has investigated how threshold voltage shifts induced by hot-carrier injection in standard high- k transistors could be used as nonvolatile memories.⁸ Moreover, recent work has shown how eNVMS implemented with this approach can be used for on-chip MLC weight storage for DNN accelerators.⁹ The same behavior has been demonstrated on a variety of technologies, including bulk, silicon-on-insulator, and FinFET devices. A major limitation for these eNVMS is the long write latency, which falls in the range of milliseconds.

Memory System for Energy-Efficient Multitask Inference (MEMTI)

In order to complement the properties of residual adapter networks and dense MLC RRAM storage, we propose MEMTI. A large fraction of the parameters in a residual adapter network is shared across multiple applications, and can be efficiently stored in MLC RRAM, while application-specific parameters can be stored in SRAM. By partitioning on-chip

Table 1. Summary of data set characteristics, and maximum training accuracy for the model trained entirely from scratch on each data set or using residual adapters on a pretrained network.

| | Data Set | | | | |
|---------------------|----------|----------|---------------|--------|--------|
| | cifar100 | aircraft | daimlerpedcls | gtsrb | ucf101 |
| # images | 50K | 7K | 30K | 40K | 9K |
| # classes | 100 | 100 | 2 | 43 | 101 |
| Full model | 72.78% | 40.98% | 99.88% | 99.97% | 73.77% |
| Only adapters | 79.61% | 43.8% | 99.51% | 99.94% | 73.16% |
| Parameters overhead | 10.4% | 10.4% | 10.1% | 10.2% | 10.4% |
| Training speed-up | 4× | 2× | 1.35× | 3.23× | 4.74× |

Pretrained shared parameters on ImageNet with 67.65% accuracy.

memory area between RRAM and SRAM, we achieve the best tradeoff between storage density for the shared parameters and fast and energy-efficient updates for the task-specific parameters (SRAM). Off-chip DRAM stores multiple sets of task-specific parameters. The resulting memory hierarchy takes advantage of RRAM nonvolatility for intermittent operation by powering down the system between inferences. In this scenario, only a small portion of the model parameters must be written to SRAM during power up or task switching. Moreover, storing task-specific parameters in more robust memory allows us to mask MLC RRAM faults via retraining, as shown in the “Model Compression and Training Techniques” section.

EVALUATION FRAMEWORK

To evaluate the proposed memory architecture, we quantify the impact of RRAM fault characteristics and MLC encoding on inference accuracy, memory architecture and array properties, and system-level performance. The fault model is derived from previous work integrating eNVM device and circuit-level fault characteristics with DNNs evaluation frameworks to allow for extensive memory and DNN codesign space exploration.⁹ We model MLC RRAM faults based on stochastic level distribution which arise from the random nature of memristors programming. When multiple levels are programmed in a single RRAM cell, the distributions overlap can be used to extrapolate the read fault probabilities for each level.

The resulting error map is then integrated in a DNN evaluation framework to simulate the

impact MLC RRAM faults on inference accuracy. The level distributions are extrapolated from measured MLC RRAM characteristics.¹⁰ We use a version of the residual adapter architecture implemented in PyTorch to evaluate the DNN accuracy under different storage schemes. The existing implementation is modified by adding transform functions that manipulate the weight parameters value according to different multilevel encoding and compression techniques.

Based on the MLC RRAM fault probabilities, we sample the value of the stored weight matrix based on a predefined multilevel encoding configuration to evaluate the impact on the model accuracy. In addition, we improve the fault model by including the effects of the sensing circuitry on the read error probability.

The corresponding framework is used to drive the design toward a solution that would minimize the on-chip memory footprint without increasing the inference error. After identifying the best MLC encoding without loss in accuracy, we perform a memory design space exploration using a modified version of NVSim.¹¹ Once again, we consider the contribution of sense amplifiers to area, energy, and performance of the memory array. Reading back the stored value requires converting the programmed analog level to a binary word, and can be done using parallel sensing or sequential sensing schemes. Parallel sensing is similar to using a flash ADC, and requires each bitline to have dedicated sense amplifiers for each possible stored level. Sequential sensing uses a single sense amplifier and recovers the stored binary word iteratively for each bit. While sequential sensing reduces the

overall number of sense amplifiers, we noticed that implementing parallel sensing with small sense amplifiers does not incur an excessive area penalty. The impact of off-chip memory accesses is quantified using a model of LPDDR4 DRAM. Power and performance estimates are derived assuming a power consumption of 200 mW at a 1 GHz operating frequency. Finally, we integrate the resulting memory hierarchy with a proven CNN accelerator architecture developed by NVIDIA (NVDLA), which, combined with NVSim results and DRAM estimates, allows us to evaluate the system energy and performance for different application scenarios.

MODEL COMPRESSION AND TRAINING TECHNIQUES

In this section, we explore the tradeoffs between different storage techniques and model accuracy. In particular, we look at the combined effect of circuit-level optimization (RRAM MLC encoding) and reducing DNN model size (quantization and pruning). Quantizing the whole model using a fixed point encoding with 2 b for sign and integer and 6 b for the fractional part achieves 80.3% accuracy on cifar100. We highlight some key insights by considering three examples. For all three cases, we take advantage of the residual adapter ability to compensate for accuracy loss associated with the faults in MLC RRAM and due to reduced DNN weight precision. We demonstrate that by fine-tuning only the task-specific parameters, the DNN learns variability-induced errors affecting the parameters stored in RRAM, and this helps maintain inference accuracy closer to the baseline value. Other approaches could be used to mitigate the impact of MLC faults on accuracy. For instance, error correction codes (ECCs) are an established solution for improving the reliability of fault-prone MLC eNVMs and can be implemented with reasonable overhead in terms of additional circuitry. ECC-based solutions offer considerable benefits when applied to eNVM storage.¹² It is worth noting that ECC protection could be used to mask errors in the shared parameters, thereby removing the need for retraining. However, ECC alone does not solve the issues related to eNVM write

endurance. For this reason, it represents a possible extension to working with residual adapters for multitask scenarios, rather than a concrete alternative.

Our first example shows the implications of storing both shared and task-specific parameters on 3 b/cell MLC RRAM. For 8-b weights, we can reduce the effective fault rate for the sign and integer values by spacing the levels as described by the nonuniform encoding technique presented by Donato *et al.*⁹ In addition, we also prune the shared parameters, which has also been shown to help mask MLC errors in nonvolatile memories.⁹ We observe that storing all weights for cifar100 in MLC RRAM results in an average accuracy over 100 trials of 28.34%, and retraining the residual parameters raises the accuracy to just 56.98%. These results highlight the importance of protecting the value of the task-specific parameters from errors.

MEMTI proposes a hybrid memory architecture in which shared and task-specific parameters are split between RRAM and SRAM. Starting from the same quantization and MLC configuration as in the previous example, we show an average accuracy of 79.72% after retraining the residual parameters if residual parameters are stored securely in SRAM. In fact, this retraining strategy can completely mask the impact of storing weights in MLC RRAM; even for the worst case accuracy degradation when storing shared parameters in RRAM (36.91%), retraining and securely storing the residual adapter parameters results in an accuracy of 79.24%, consistent with baseline accuracy. Motivated by this result, we propose leveraging this technique to achieve even more aggressively dense storage by reducing the number of bits for the shared weights to 6. This configuration uses just 2 RRAM cells per weight. In this case, training the residual adapters results in an average accuracy of 79.05%. The worst-case accuracy before training the adapters is 2.04%, and can be recovered up to 77.57% after training. Figure 1 shows the same trend in terms of the nominal baseline accuracy, accuracy before retraining, and accuracy after retraining of the task-specific parameters for the data sets considered in this work.

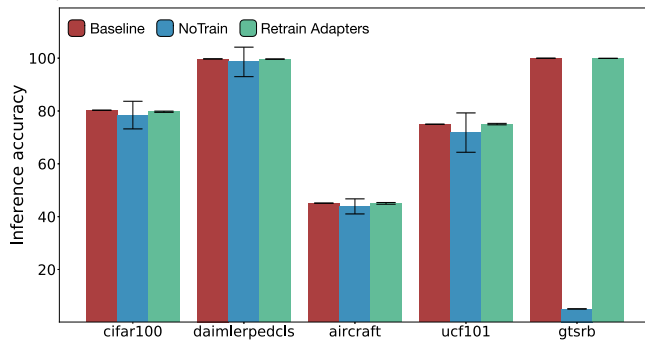


Figure 1. Accuracy for different data sets when the weights are stored using MEMTI. The results before and after task-specific parameters retraining are compared to the baseline model accuracy for the compressed, error-free model. Each bar shows the mean accuracy and standard deviation over 100 random trials.

SYSTEM-LEVEL CHARACTERIZATION

As shown in Figure 2, the baseline NVDLA system comprises a convolutional core with 1024 MAC units fed by a convolutional buffer and supplemented by several additional computational units for pooling and data transformation operations. NVDLA also supports a memory interface block and DMA that fetches model weights per layer from off-chip DRAM and leverages on-chip SRAM (2 MB) to buffer inputs and intermediate results of computation between layers in the DNN. We flexibly integrate MEMTI with the NVDLA performance model as an additional memory interface to leverage for model weights, either in addition

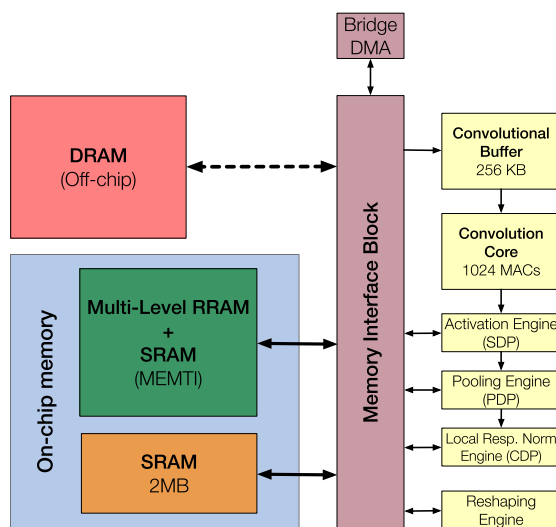


Figure 2. NVDLA system diagram, with additional optional interface to MLC RRAM for on-chip weight storage.

to or in place of fetching parameters from off-chip DRAM.

For a competitive multitask inference application, we evaluate the performance, energy, and area for the NVDLA system when executing three inferences per input frame using three representative visual tasks, namely image classification (cifar100), object detection (gtsrb), and action recognition (UCF101). This series of tasks computed per input frame would be appropriate, for example, for an autonomous vehicle or a drone processing sensory data to understand and interact with the surrounding environment. For this application, we set the target operating frequency to 30 frames per second (FPS), or 90 inference tasks per second, which satisfies a breadth of applications. Both NVSim and NVDLA results are extrapolated for a system manufactured using a 22-nm technology node.

DRAM-based design: As a baseline case, we assume that the accelerator is continuously processing input frames and fetching both shared and task-specific parameters from off-chip DRAM for each layer's computation. This DRAM-only, always-on operation consumes a total power of 493 mW and a peak performance of 749 FPS. The estimated power includes datapath, DRAM refresh, and on-chip SRAM leakage. At this stage, the on-chip SRAM is exclusively used for storing the input features and intermediate values. We compute the energy per frame at peak performance to be 1.17 mJ.

Provisioning for on-chip SRAM: We first show the case in which we allocate enough on-chip SRAM to store the entire set of parameters for a single task. For a system designed to run a single inference task, having the option of storing all the network parameters on chip allows to reduce the memory access energy by 40 \times . This result demonstrates the strong impact of off-chip memory access on the entire system energy. These high energy savings are however impractical to realize with SRAM since for a 22-nm technology node, we estimate a total area of 6.55 mm². Moreover, when we consider the full system energy in the multitask scenario described above, the periodic parameter updates and SRAM leakage power reduce the energy savings to 0.64 \times .

Improving storage density with eNVM: As a first step toward reducing both power consumption and memory footprint, we consider storing the weights on chip using MLC RRAM. Without applying any DNN-level optimization, a multitask operation would still require updating all the model parameters stored in RRAM when switching to different inference tasks. While this type of operation has a clear downside dictated by the RRAM write endurance, our system level evaluation exposes other limitations. Although the overall leakage power and on-chip memory area can be reduced to 298 mW and 0.347 mm², respectively, the energy per inference increases to 14.58 mJ. This is caused by the combined effect of RRAM write energy and latency. These examples highlight the need for a solution capable of balancing on-chip memory density and write performance.

MEMTI for intermittent operation: MEMTI removes the memory write costs by replacing the RRAM portion storing the task-specific parameters with SRAM. This is possible thanks to the adoption of residual adapters in the DNN network. For the resulting design, the total system power is 362 mW and the energy per inference is 1.56 mJ, which is comparable with the baseline result. Isolating the costs associated with weight storage emphasizes the benefits introduced by MEMTI: power consumption is reduced by 3.9 \times , with an area overhead of 1.16 mm². The resulting peak performance is 429 FPS, well above the application requirements. Intermittent operation is where MEMTI truly stands out by taking advantage of the nonvolatility of RRAM. In this scenario, we fix the operation at 30 FPS and power down the system between frames, which reduces the energy per inference by 10.65 \times .

RRAM-based specialized design: Alternatively, we consider the case specifically tailored for the three chosen tasks. Using residual adapters still reduces the weights storage requirements by a factor of 2.3 \times . In addition, we store the whole set of task-specific parameters for the three tasks in SLC RRAM. Relaxing the density requirement for task-specific parameters allows to preserve inference accuracy while removing the need for additional SRAM. This design choice reduces the overall power to 343 mW, and the area to 1 mm².




Figure 3. Energy versus FPS for the different design configurations normalized to the DRAM baseline. The power savings of the RRAM-based design for higher frame rates is exacerbated by the frequent RRAM writes, making the design less efficient than the DRAM-based baseline. On the other hand, the energy per inference for MEMTI and RRAM adapters is strictly better than the baseline.

Allocating enough memory for storing all the parameters on chip increases the energy savings compared to the baseline by 13.6 \times , making this design the most area and energy efficient. Nonetheless, MEMTI maintains the advantage in terms of flexibility and robustness to RRAM errors thanks to ease of reprogrammability for the task-specific parameters, for which the memory capacity is determined only by the network structure and therefore is independent from the breadth of tasks considered in a specific application.

Figure 3 shows the relationship between FPS and energy per inference normalized to the DRAM case. The all SRAM and all RRAM configurations are heavily penalized by the inability of efficiently implement a multitask inference system. On the other hand, a codesign of the memory and DNN model using residual adapters shows much higher energy savings compared to the baseline. Table 2 summarizes the results at 30 FPS for the different configuration cases.

CONCLUSION

With the increasing adoption of DNN hardware accelerators for edge devices, there is a growing need for scalable design approaches that provide flexible and cost-effective implementations. In evaluating the performance of different memory solutions integrated with a DNN hardware accelerator, we show that technological improvements alone do not always lead to

Table 2. Summary of power, performance, and area for the four design configurations considered in this article.

| | Power [mW] | Max FPS | WMem Area [mm ²] | Saved energy | On-chip SRAM | On-chip RRAM |
|---------------|---------------|---------|---------------------------------|-----------------|-----------------|-----------------|
| all DRAM | 493 | 749 | — | 1× | 2MB | — |
| all SRAM | 634 | 485 | 6.55 | 0.64× | 8.5MB | — |
| all RRAM | 298 | 47 | 0.347 | 1.62× | 2MB | 2.2MB |
| MEMTI | 344 | 429 | 1.16 | 10.65× | 2.7MB | 2MB |
| RRAM adapters | 301 | 396 | 1 | 13.6× | 2MB | 4MB |

The energy savings are normalized to the all DRAM configuration for the intermittent multitask operation over three tasks running at 30 FPS. On-chip RRAM shows the physical memory capacity (i.e., number of cells).

the most optimized design, especially in the context of multitask inference. A codesign approach that leverages the properties of emerging memory technologies and DNN models allows to achieve both energy efficiency and flexibility. With this in mind, we present MEMTI as a methodology for enabling energy-efficient multitask inference on edge devices, while reducing the cost of nonvolatile memory writes. In addition, training the task-specific parameters based on the memory characteristics allows recovery of accuracy lost due to fault-prone eNVM storage.

ACKNOWLEDGMENT

This work was supported by the Applications Driving Architectures (ADA) Research Center, a JUMP Center cosponsored by SRC and DARPA.

REFERENCES


1. B. Reagen, R. Adolf, P. Whatmough, G.-Y. Wei, and D. Brooks, *Deep Learning for Computer Architects (Synthesis Lectures on Computer Architecture)*, vol. 12, no. 4, pp. 1–123, Aug. 2017. [Online]. Available: <http://www.morganclaypool.com/doi/10.2200/S00783ED1V01Y201706CAC041>
2. L. Kaiser *et al.*, “One model to learn them all,” 2017. [Online]. Available: <http://arxiv.org/abs/1706.05137>
3. J. Yosinski *et al.*, “How transferable are features in deep neural networks?” in *Proc. 27th Int. Conf. Neural Inf. Proc. Syst.*, 2014, vol. 2, pp. 3320–3328. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2969033.2969197>
4. P. N. Whatmough, C. Zhou, P. Hansen, S. K. Venkataramanaiah, J. Sun Seo, and M. Mattina, “FixyNN: Efficient hardware for mobile computer vision via transfer learning,” in *Proc. SysML Conf.*, 2019.
5. S.-A. Rebuffi, H. Bilen, and A. Vedaldi, “Efficient parametrization of multi-domain deep neural networks,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2018.
6. D. C. Daly, L. C. Fujino, and K. C. Smith, “Through the looking glass—The 2018 edition: Trends in solid-state circuits from the 65th ISSCC,” *IEEE Solid-State Circuits Mag.*, vol. 10, no. 1, pp. 30–46, Jan. 2018.
7. S. H. Kulkarni *et al.*, “A 32nm high-k and metal-gate anti-fuse array featuring a 1.01m21t1c bit cell,” in *Proc. Symp. VLSI Technol.*, Jun. 2012, pp. 79–80.
8. F. Khan, E. Cartier, C. Kothandaraman, J. C. Scott, J. C. S. Woo, and S. S. Iyer, “The impact of self-heating on charge trapping in high-*k*-metal-gate nFETs,” *IEEE Electron Device Lett.*, vol. 37, no. 1, pp. 88–91, Jan. 2016.
9. M. Donato *et al.*, “On-chip deep neural network storage with multi-level eNVM,” in *Proc. 55th Annu. Design Autom. Conf.*, 2018, pp. 169:1–169:6. [Online]. Available: <http://doi.acm.org/10.1145/3195970.3196083>
10. L. Zhao *et al.*, “Improved multi-level control of RRAM using pulse-train programming,” in *Proc. Techn. Program—Int. Symp. VLSI Technol., Syst. Appl.*, Apr. 2014, pp. 1–2.
11. X. Dong, C. Xu, Y. Xie, and N. P. Jouppi, “NVSIM: A circuit-level performance, energy, and area model for emerging nonvolatile memory,” *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 31, no. 7, pp. 994–1007, Jul. 2012.
12. L. Pentecost, M. Donato, B. Reagen, U. Gupta, S. Ma, G.-Y. Wei, and D. Brooks, “MaxNVM: Maximizing DNN storage density and inference efficiency with sparse encoding and error mitigation,” in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 769–781. [Online]. Available: <https://doi.org/10.1145/3352460.3358258>

Marco Donato is currently a Postdoctoral Fellow with Harvard University, Cambridge, MA, USA. His research interests include novel design methodologies targeting energy-efficient and reliable circuits and architectures for emerging computing paradigms. He received the PhD degree in electrical engineering from Brown University, Providence, RI, USA. Contact him at: mdonato@seas.harvard.edu.

Lillian Pentecost is currently working toward the PhD degree with Harvard University, Cambridge, MA, USA, studying specialized computer architecture and memory systems for machine learning. She received the bachelor's degree in physics and computer science from Colgate University, Hamilton, NY, USA, in 2016, and the master's degree in computer science from Harvard University in 2019. Contact her at: lillian_pentecost@seas.harvard.edu.

David Brooks is currently the Haley Family Professor of computer science with Harvard University, Cambridge, MA, USA. His research interests include architectural and software approaches to address power, thermal, and reliability issues for embedded and high-performance computing systems. He received the PhD degree in electrical engineering from Princeton University, Princeton, NJ, USA. Contact him at: dbrooks@seas.harvard.edu.

Gu-Yeon Wei is currently a Gordon McKay Professor of electrical engineering and computer science with Harvard University, Cambridge, MA, USA. His research interests include mixed-signal integrated circuits, computer architecture, and runtime software, looking for cross-layer opportunities to develop energy-efficient systems. He received the PhD degree in electrical engineering from Stanford University, Stanford, CA, USA. Contact him at: guyeon@seas.harvard.edu.



IEEE TRANSACTIONS ON

COMPUTERS

Call for Papers: *IEEE Transactions on Computers*

Publish your work in the IEEE Computer Society's flagship journal, *IEEE Transactions on Computers*. The journal seeks papers on everything from computer architecture and software systems to machine learning and quantum computing.

Learn about calls for papers
and submission details at
www.computer.org/tc.

