MEXT

WHITEPAPER | Q1'26

The Cost Benefits of MEXT for EDA Workloads

An overview of MEXT AI-Powered Predictive Memory™ software performance for EDA

Executive Summary

The Electronic Design Automation (EDA) landscape is characterized by ever-increasing chip complexity, driving an unprecedented demand for compute resources with large memory configurations. The recent surge in AI-driven EDA flows has amplified this demand by an order of magnitude. This white paper highlights the significant cost benefits and performance improvements realized by integrating MEXT Predictive Memory™ into IT infrastructure targeting EDA workloads. MEXT delivers 2X more EDA capacity on existing infrastructure and within the same budget. By leveraging MEXT, chip developers can maximize the return on their IT and EDA investments, leading to superior Quality of Results (QoR) and faster Time to Results (TtR). The value proposition of MEXT is further strengthened by the current high prices of DRAM, which challenge the acquisition of the large DRAM configurations necessary for cutting-edge chip design.

The EDA Memory Challenge

The complexity of modern chip designs, encompassing billions of transistors, necessitates EDA tool flows that operate on the most powerful computers with substantial memory capacity. EDA workflows, including design, verification, and implementation, frequently become memory-bound, leading to stalled processes, inefficient resource utilization, and extended design cycles.

The rise of AI-driven EDA methodologies, which involve complex models and large datasets, has escalated the memory requirements dramatically. Traditional computing infrastructure struggles to keep pace with this demand without substantial and costly upgrades.

Furthermore, the recent volatility and high cost of DRAM have made provisioning state-of-the-art computers with the required large DRAM configurations economically challenging for many semiconductor companies. This creates a critical need for memory solutions that can deliver high performance and capacity at a lower total cost of ownership.

MEXT Predictive Memory™: A Cost-Effective Solution

MEXT Predictive Memory™ provides an innovative solution to the EDA memory challenge. By intelligently managing DRAM memory resources and predicting the memory access patterns of EDA tools, MEXT significantly improves memory utilization and access efficiency. This optimization allows existing IT infrastructure to handle larger, more complex EDA workloads than previously possible, effectively increasing the "usable" memory size without a physical DRAM upgrade or any application modification.

Benefits Across the EDA Flow

The entire chip development lifecycle benefits from MEXT integration:

- **Design:** Faster loading and manipulation of large design databases.
- **Verification:** Accelerated simulation and verification runs by managing vast state spaces more efficiently.
- **Implementation (Placement and Routing):** Improved performance on memory-intensive physical design stages, reducing runtime and potential congestion.
- **Physical Verification and Timing Analysis:** Ability to run larger designs and full-chip runs at critical points in a project schedule.

Maximizing IT and EDA Investments

The primary cost benefit of MEXT is its ability to extract more performance from existing hardware investments. Instead of purchasing new, expensive servers with large DRAM configurations to meet memory demands, MEXT optimizes the use of currently-deployed memory. MEXT also does not require any modifications to the workloads which enables rapid deployment with no risk to the project.

This approach translates directly into:

1. **Reduced Capital Expenditure (CapEx):** Lower immediate spending on new hardware.
2. **Lower Total Cost of Ownership (TCO):** Extended lifespan of current compute clusters and lower power consumption per workload.
3. **Improved Resource Utilization:** Higher throughput of EDA jobs on the same server farm thus enabling faster TtR.

Performance Characteristics on Real EDA Workloads

We have engaged with leading semiconductor companies to rigorously test and benchmark MEXT Predictive Memory™ against standard DRAM configurations across a variety of real-world EDA workloads. The data unequivocally demonstrates the performance and TtR benefits.

Across large synthesis runs, full-chip static timing analysis, and complex physical verification workloads, MEXT delivers nearly equivalent time-to-results as 100% DRAM configurations, with 50% of the DRAM removed from the system. If we estimate performance as equivalent (1X) and costs as halved (½ X) with the 50% DRAM + MEXT configuration, this translates to a **2X performance-per-dollar improvement** with MEXT.

MEXT can also deliver nearly-equivalent performance with 62.5% of the DRAM removed from the system, and only minimal performance degradation with 75% of the DRAM removed from the system. The results of a customer's testing of these scenarios are visualized below:

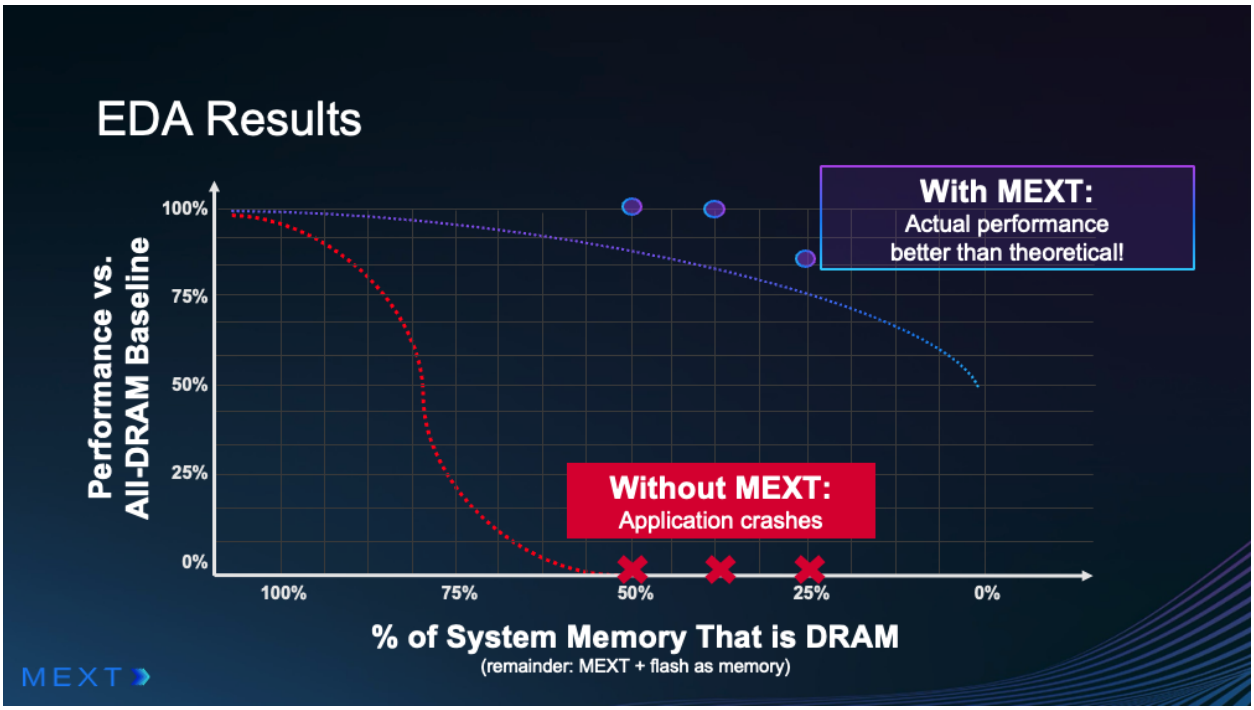

Quality of Results (QoR) Enhancement

Beyond TtR, MEXT contributes to improved Quality of Results. Running larger partitions as well as more AI-driven simulation, verification, and Place and Route runs, which MEXT enables, is key to improving the resulting QoR of the chip.

Conclusion

The convergence of increasing chip complexity, the emergence of AI-driven EDA flows, and the challenging high-cost environment for DRAM makes MEXT Predictive Memory™ an essential and timely solution for the semiconductor industry. MEXT offers a clear pathway for chip developers to significantly reduce their CapEx and TCO while achieving superior performance metrics. By allowing developers to get more out of their IT and EDA investments, MEXT directly contributes to faster Time to Results and improved Quality of Results, ensuring companies remain competitive in the rapidly evolving chip design landscape.

Contact eric.fitzpatrick@mext.ai for a personalized consultation.