Memory Controller Policies for DRAM Power Management

Xiaobo Fan, Carla S. Ellis, Alvin R. Lebeck

- Sriram Aananthakrishnan

Ideas

- Investigates Memory Controller policies for DRAM power management
- Claim : Transitioning to low power state when DRAM chip becomes idle
- Analytic Model : idle time of DRAM chips as an exponential distribution
- Model Validation by Trace driven Simulation

Key Terms

- Gap – interval between clustered accesses
  - Large gap values
  - Exponential distribution observed
- Threshold (Th) – Time spent in current power state before transitioning to low power state

RDRAM

- Power States
  - Active, Standby, Nap, Powerdown
- Active state to perform transactions
- Nap mode – Maximum power savings
- Additional delay for Clock resynchronization

RDRAM Power State and Transition Values

Power State Transition	**Power** (mW)	**Time** (nS)
Active	$P_a = 300$	t_{acc}=60
Standby	$P_s = 180$	-
Nap	$P_n = 30$	-
Powerdwn	$P_p = 3$	-
Stby $\rightarrow$ Act	$P_{s \rightarrow a} = 240$	$T_{s \rightarrow a} = +6$
Nap $\rightarrow$ Act	$P_{n \rightarrow a} = 165$	$T_{n \rightarrow a} = +60$
Pdn $\rightarrow$ Act	$P_{p \rightarrow a} = 152$	$T_{p \rightarrow a} = +6000$

DRAM Power Management

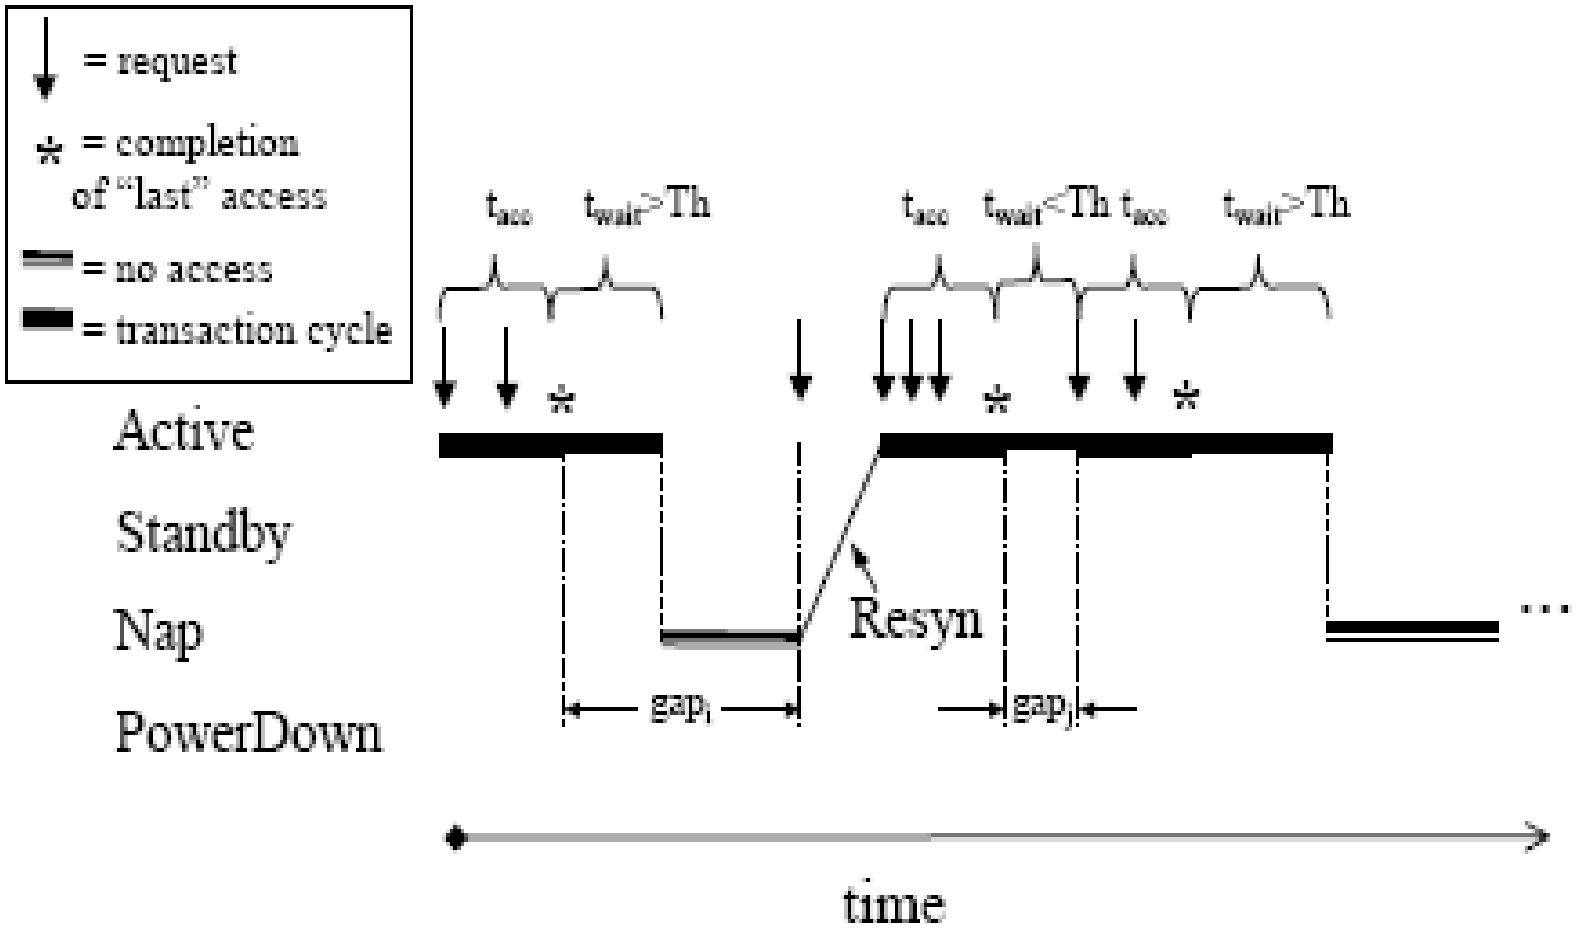

Methodology

- Metric : Energy – Delay product
- Analytic Evaluation
  - Evaluates the energy delay product with transition policies
- Trace Driven Simulation
  - Characterizes the memory access patterns
  - Validate the analytic model

Memory Access Pattern

- Time Measurement between clustered misses
- Approximation of Gap distribution

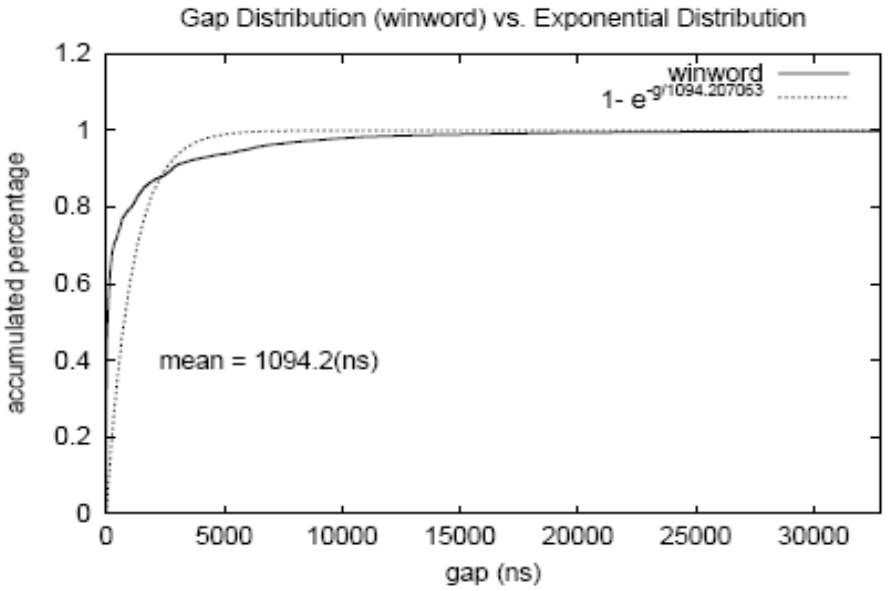

Benchmark	Result
Compress95	Pass
Go	Pass
Netscape	Pass
Acroread	Fail
PowerPoint	Fail
Winword	Fail

Analytic Model

1. Memory Access gap as Exponential Distribution

$$p(g) = \frac{1}{\mu} e^{-\frac{g}{\mu}}$$

2. Mean time of staying in low power

$$t_{nap} = \int_{Th}^{\infty} p(g)(g - Th)dg = \mu e^{-\frac{Th}{\mu}}$$

3. Mean Energy Savings

$$\Delta e_1 = (P_a - P_n)t_{nap} = \mu(P_a - P_n)e^{-\frac{Th}{\mu}}$$

Analytic Model

4. Mean Energy Cost for resynchronization

$$\begin{aligned}\Delta e_2 &= \frac{1}{2}(P_a + P_n)T_{n\to a}\int_{Th}^{\infty} p(g)dg \\ &= \frac{1}{2}(P_a + P_n)T_{n\to a}e^{-\frac{Th}{\mu}}\end{aligned}$$

5. Mean Energy Cost

$$\begin{aligned}\Delta e &= \Delta e_2 - \Delta e_1 \\ &= [\frac{1}{2}(P_a + P_n)T_{n\to a} - (P_a - P_n)\mu]e^{\frac{-Th}{\mu}}\end{aligned}$$

Analytic Model

6. Mean Increased Delay

$$\Delta d = \int_{Th}^{\infty} T_{n \to a} p(g) dg = T_{n \to a} e^{\frac{-Th}{\mu}}$$

7. With Power Transition Policy

$$e = e_0 + \Delta e \qquad d = d_0 + \Delta d$$

Analytic Model

8. Change in Energy Delay Product

$$\begin{aligned}\Delta(E \bullet D) &= E \bullet D - E_0 \bullet D_0 \\ &= n^2(d_0\Delta e + \Delta d e_0 + \Delta d \Delta e)\end{aligned}$$

9. Change in Energy delay product per gap

$$\Delta(e \bullet d) = d_0\Delta e + \Delta d e_0 + \Delta d \Delta e$$

Analytic Model – Mean Gap Variation

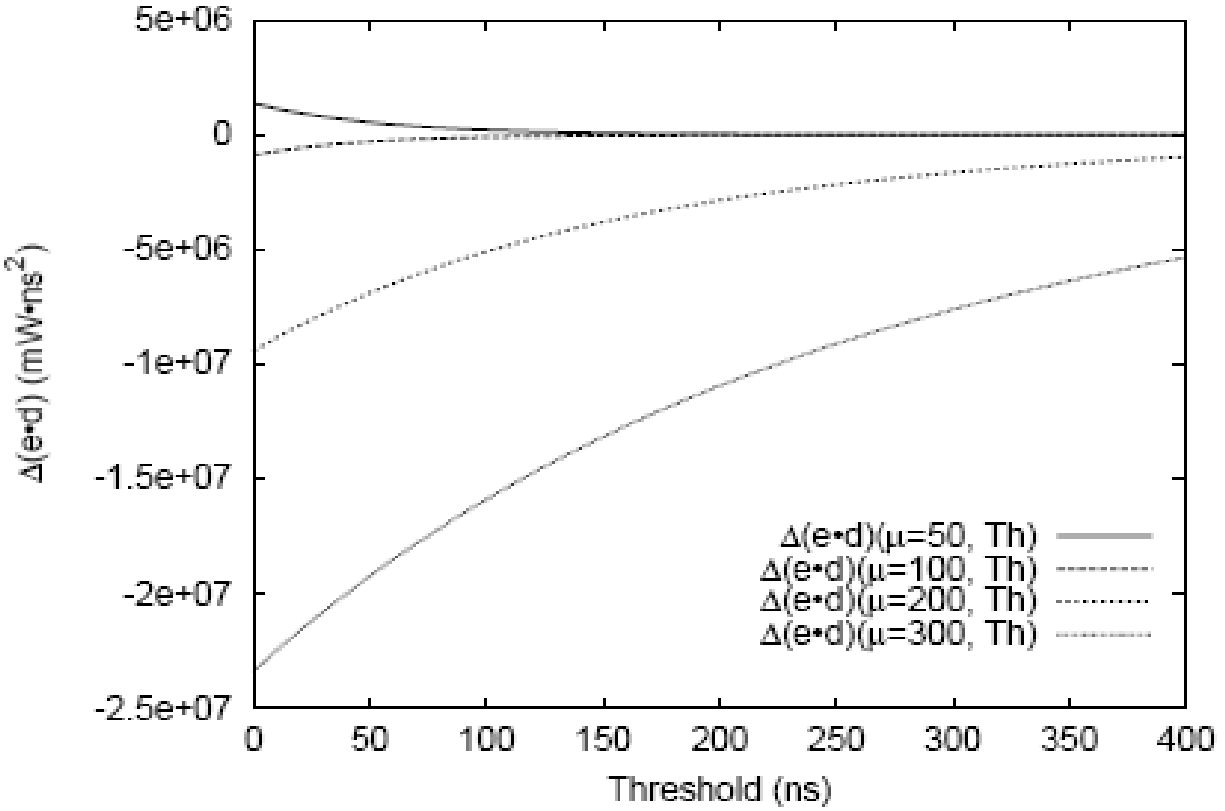

Figure 4: $\Delta(e \bullet d) = f(\mu, Th), \mu = 50, 100, 200, 300ns$

Analytic Model – Th Variation

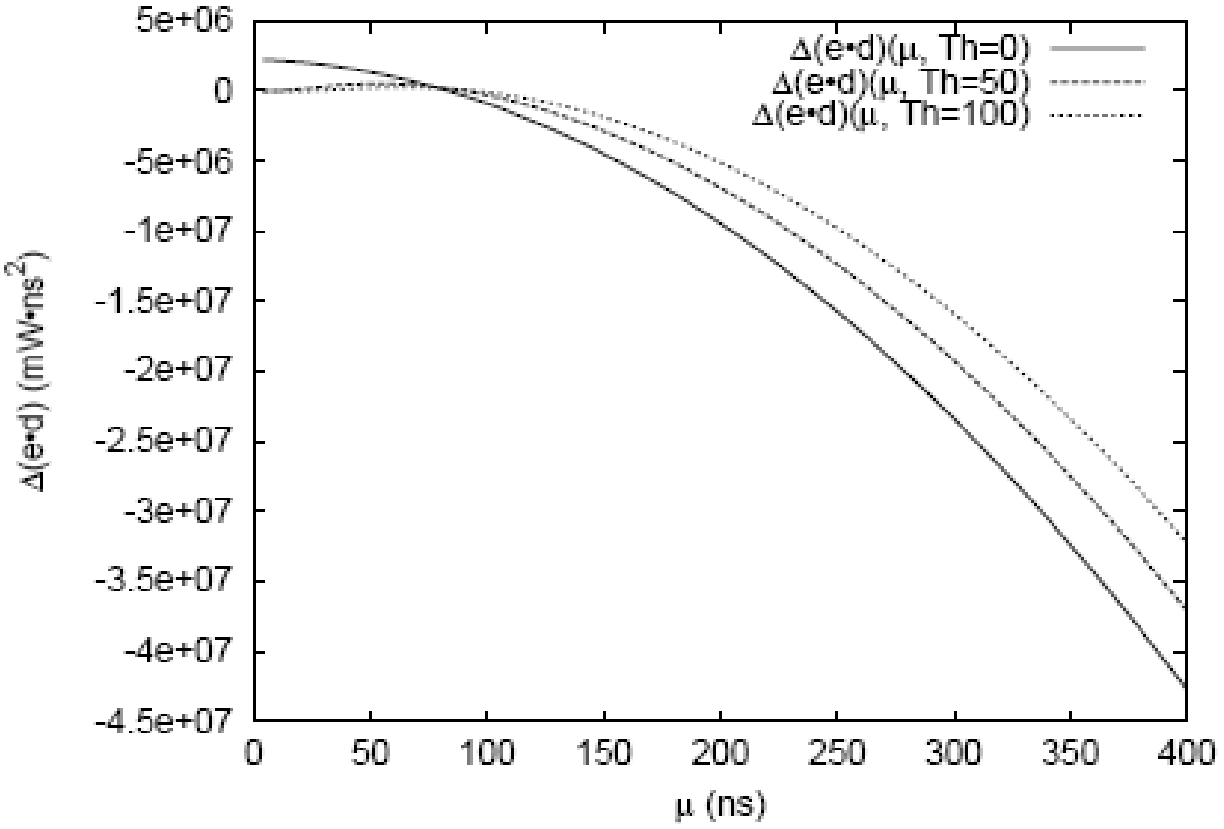

Figure 3: $\Delta(e \bullet d) = f(\mu, Th), Th = 0, 50, 100ns$

Validation

- Compare Energy-delay product from simulation and model
- Compress95 – within 5% of Simulation results
- Winword – 50% in some cases

Model Validation – Simulation Graphs

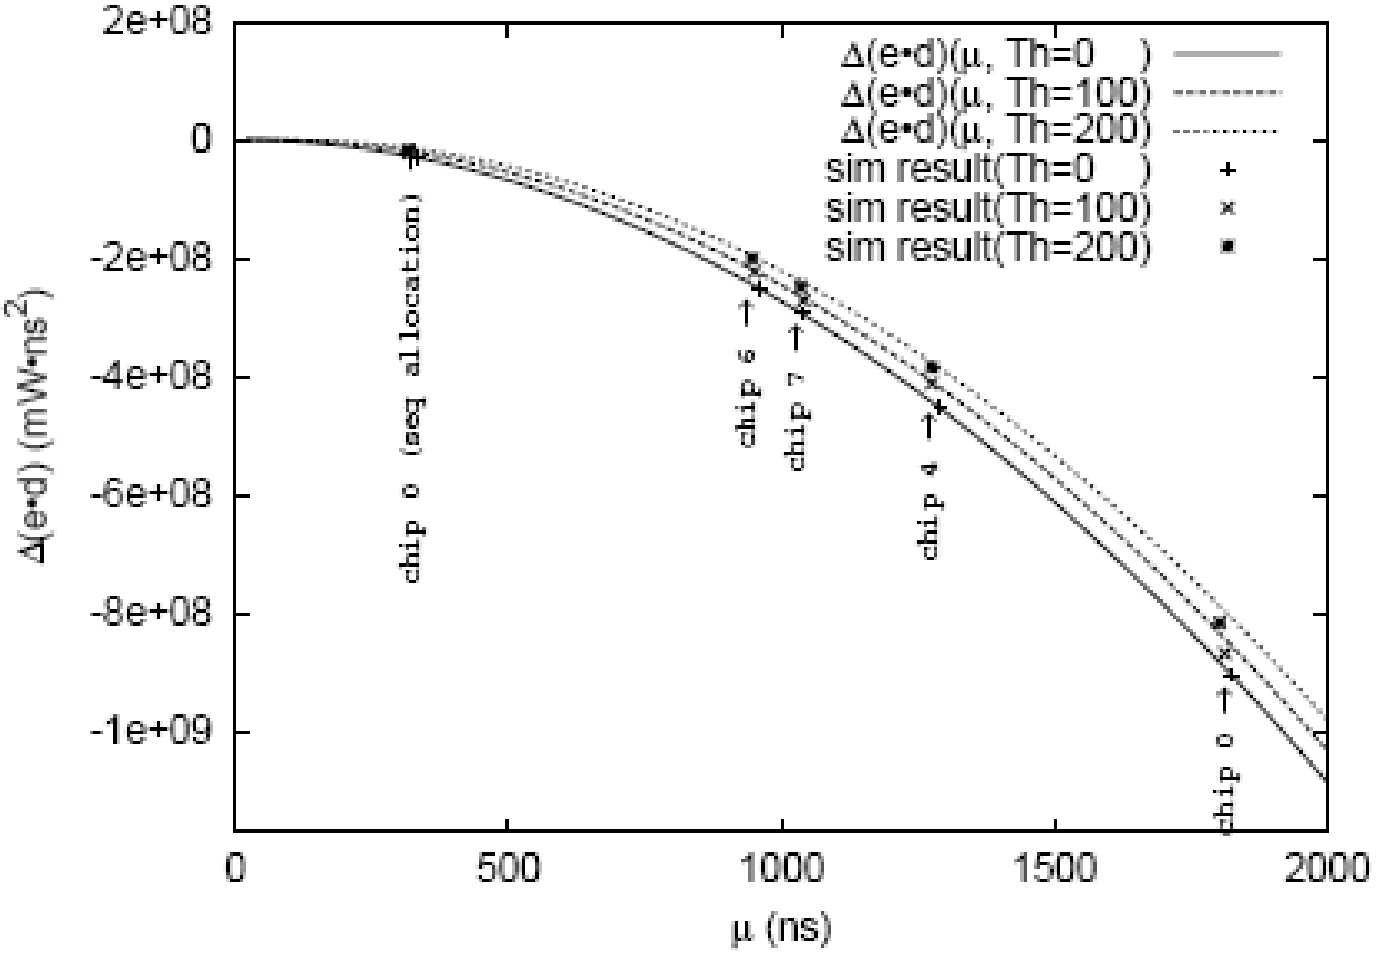

Conclusions

- Qualitative results are same for both simulation and model
- DRAM chips should transition to low power state when they are idle