

Energy Efficient Data Encoding in DRAM channels exploiting Data Value Similarity

Hoseok Seol, Wongyu Shin, Jaemin Jang,
Jungwhan Choi, Jinwoong Suh, Lee-Sup Kim

Department of Electrical Engineering






Outline

- 1. Introduction**
2. BD-Encoding
3. Evaluation Results
4. Conclusion

Modern DRAM Interface

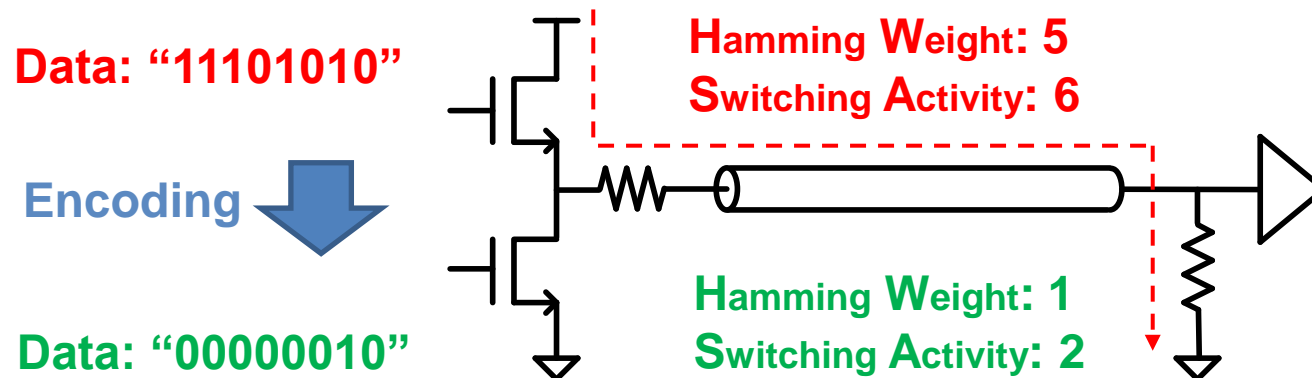
- **DRAM off-chip data bus** consumes significant energy.
- Data Bus Energy: Switching + Termination (dominant)
- Modern DRAMs introduce **asymmetric termination**.
 - ⇒ **Pseudo Open Drain (POD)**: DDR4, GDDR4/5
 - ⇒ **Low Voltage Swing Terminated Logic (LVSTL)**: LPDDR4



Hamming Weight & Interface Energy

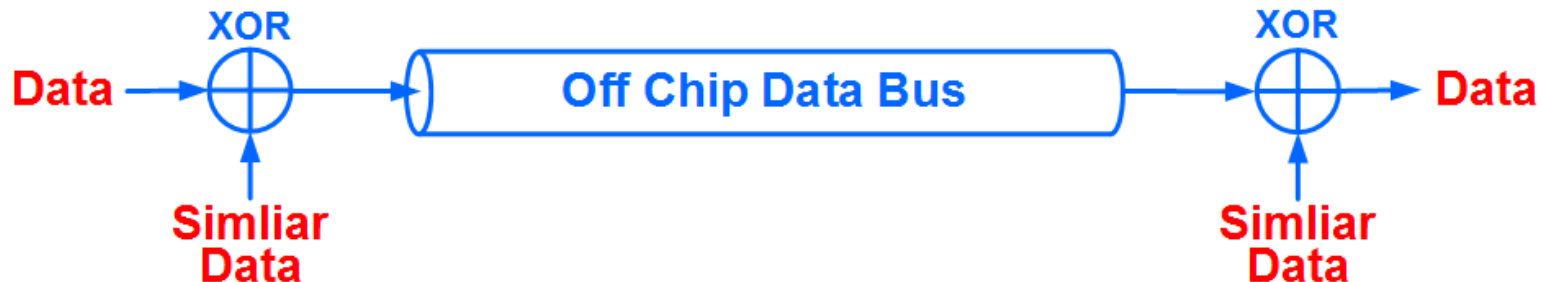
- **Hamming Weight:** number of 1's in a string of bits.
- **Decreasing Hamming Weight** reduces both the **termination** and **switching energy**.
- We propose novel **data encoding** to reduce data bus energy.

Ex) LVSTL interface



Bitwise Difference (BD) Encoding

- **Observation: Similar data words** are sent over the DRAM data bus.
- **Key Idea:** Transfer the **bit-wise difference** between a **current data word** and the **most similar data words**.

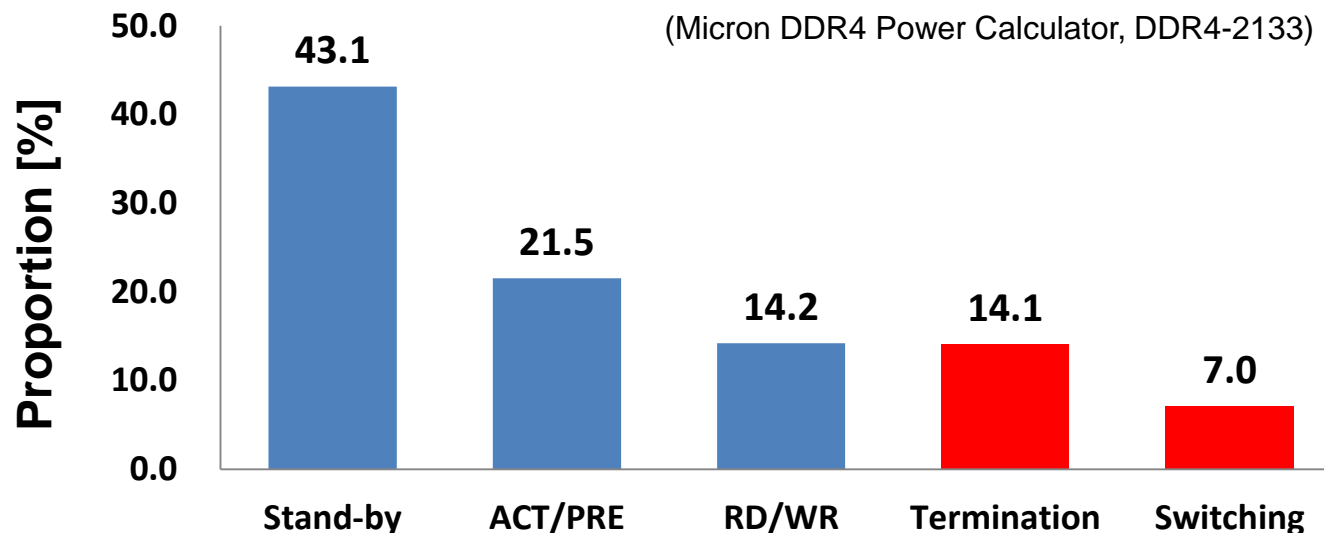


- **Energy Reduction: 58.3% of termination** and **45.3% of switching energy**.

Outline

1. Introduction
- 2. BD-Encoding**
3. Evaluation Results
4. Conclusion

Motivation



< Energy dissipation in DRAM sub-system >

➤ Energy dissipated in DDR4 data bus:


Termination (14.1%) + Switching Activity (7%)

Observation: Data Value Similarity

Transfer	libquantum	mcf
1	38 ad b3 00 18 83 24 00	18 67 df aa aa 2a 00 00
2	58 ad b3 00 18 83 24 00	01 00 00 00 00 00 00 00
3	78 ad b3 00 18 83 24 00	98 53 b8 aa aa 2a 00 00
4	98 ad b3 00 18 83 24 00	08 63 b8 aa aa 2a 00 00
5	a8 ad b3 00 18 83 24 00	00 00 00 00 00 00 00 00
6	c8 ad b3 00 18 83 24 00	00 27 bd aa aa 2a 00 00

- **Observation: Similar data words** are sent over the **DRAM data bus**.

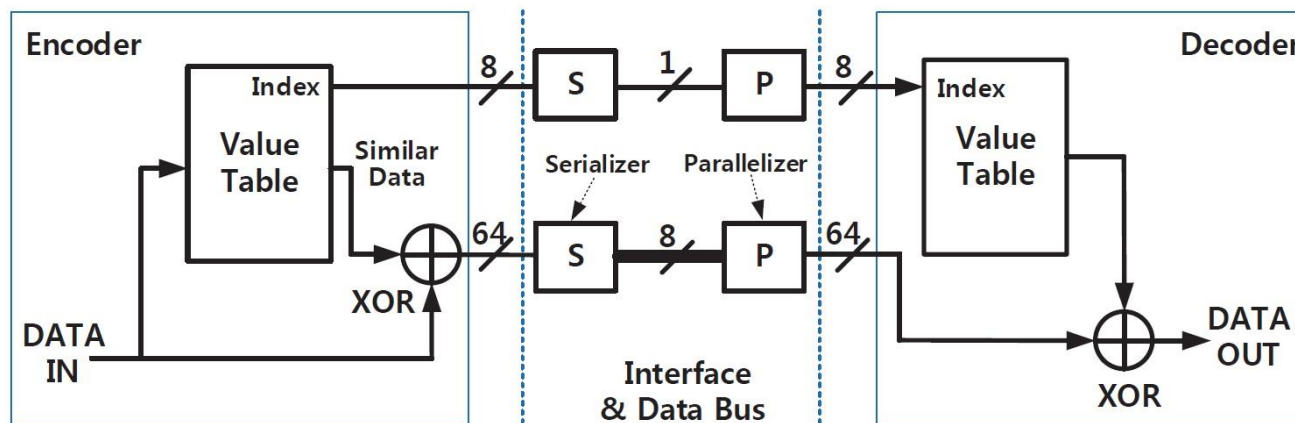
Observation: Data Value Similarity



< Probability of 90% data matching among 64 recent data words >

- **All the workloads in SPEC 2006 have Data Value Similarity.**
- **The probability of the similar data occurrence (with recent 64 data words) is 72% in SPEC 2006 workloads.**

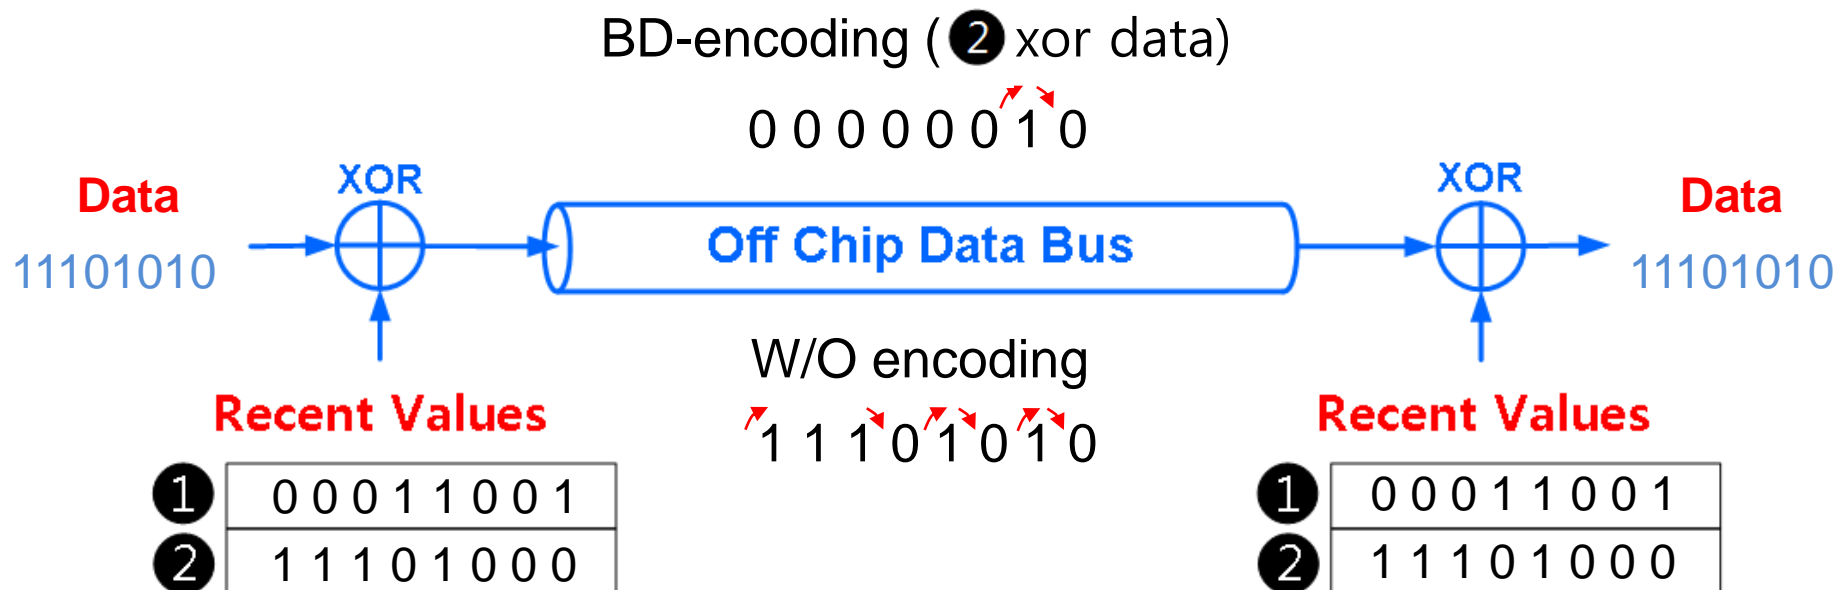
Bitwise Difference Coder



< Overall Structure of BD-coder >

- Recent data is stored in both tables in Encoder / Decoder
- When transfer data, **search the most similar data word.**
- If similar data exists, transfer **1) bitwise difference, 2) index NO.**
If not, transfer the original data.

Example of BD-encoding



	W/O encoding	BD-encoding
Hamming Weight	5	1
Switching Activity	6	2

Hardware Overheads

➤ **Coder** (data table 64 entries)

- ✓ Area: 0.044% of commodity DDR4
- ✓ Latency: 2.3ns (Transmitter), 0.7ns (Receiver)
- ✓ Energy: 7pJ (Transmitter), 2pJ (Receiver)
- ✓ Designed by 65nm logic process

➤ **Index Line**

- ✓ a single extra line per 8 data lines.
- ✓ can be shared with DBI / DM pins in DDR4.

Outline

1. Introduction
2. BD-Encoding
- 3. Evaluation Results**
4. Conclusion

Methodology

Component	Parameters
Processor	Gem5, X86, 3.3GHz
Caches	L1 I-cache : 32KB, 4way L1 D-cache : 64KB, 4way L2 cache : 2MB, 8way
DRAM	DDR4-2133, 8GB
Interface	Pseudo Open Drain (DDR4) Termination Energy Calculation: Micron DDR4 Power Calculator Switching Energy Calculation: $E = CV^2$ Channel capacitance: 15 [pF]
Workloads	SPEC CPU 2006


Comparison Points

- **Data Bus Inversion** [M.stan, TVLSI '95]
 - ⇒ Transfer inverted data if the hamming weight of inverted one is smaller.
 - ⇒ Adopted in the commodity DRAMs (GDDR4/5, DDR4, LPDDR4)

- **Power Protocol** [K.Basu, MICRO'02], **Frequent Value Encoding** [J.Yang, ISLPED'01]
 - ⇒ Transfer the table index instead of data when current data is the same as data transferred recently.

- **Variable Length Value Encoder** [D.suresh, ICCD'05]
 - ⇒ Transfer the table index instead of data when current data is partly matched with data transferred recently.

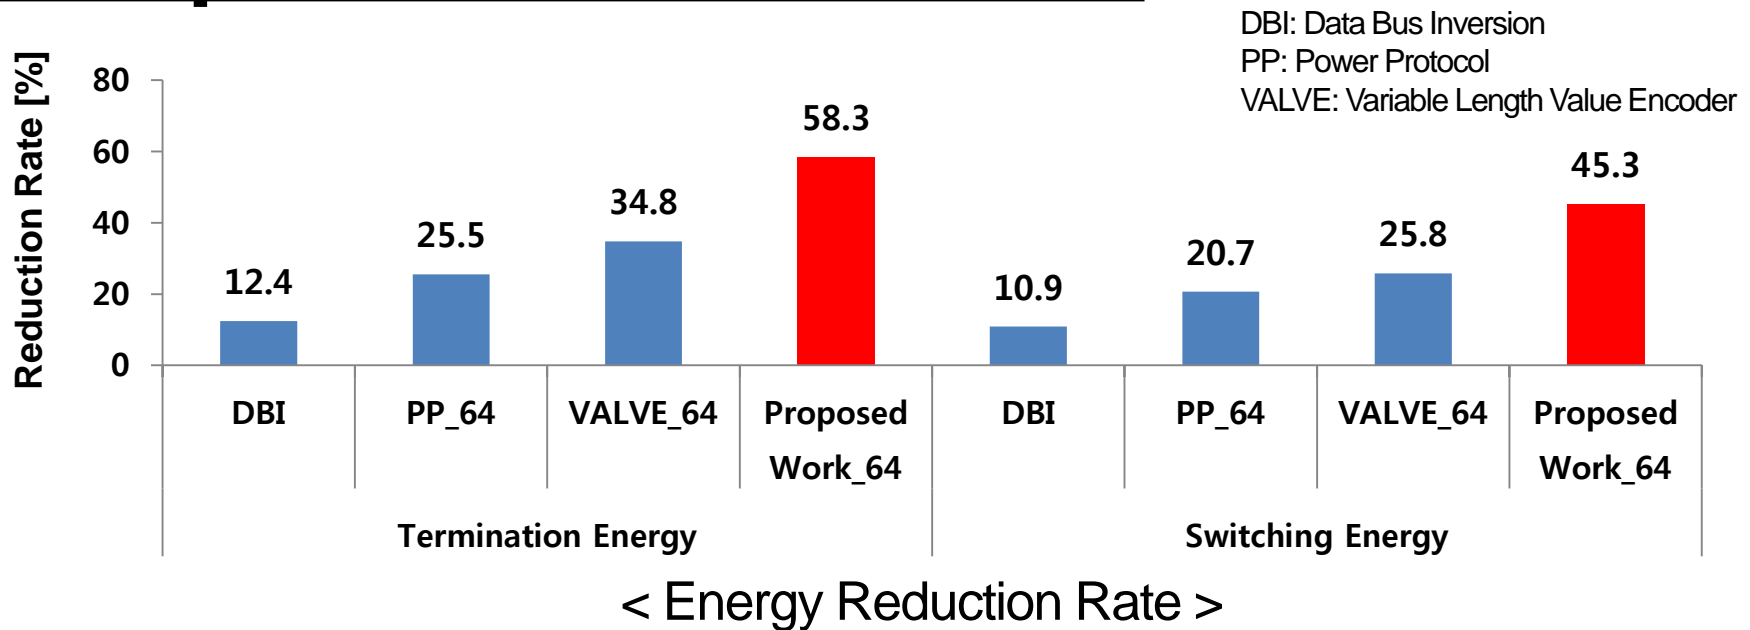
Hamming Weight Reduction



< Hamming Weight Reduction Rate of BD-Encoding >

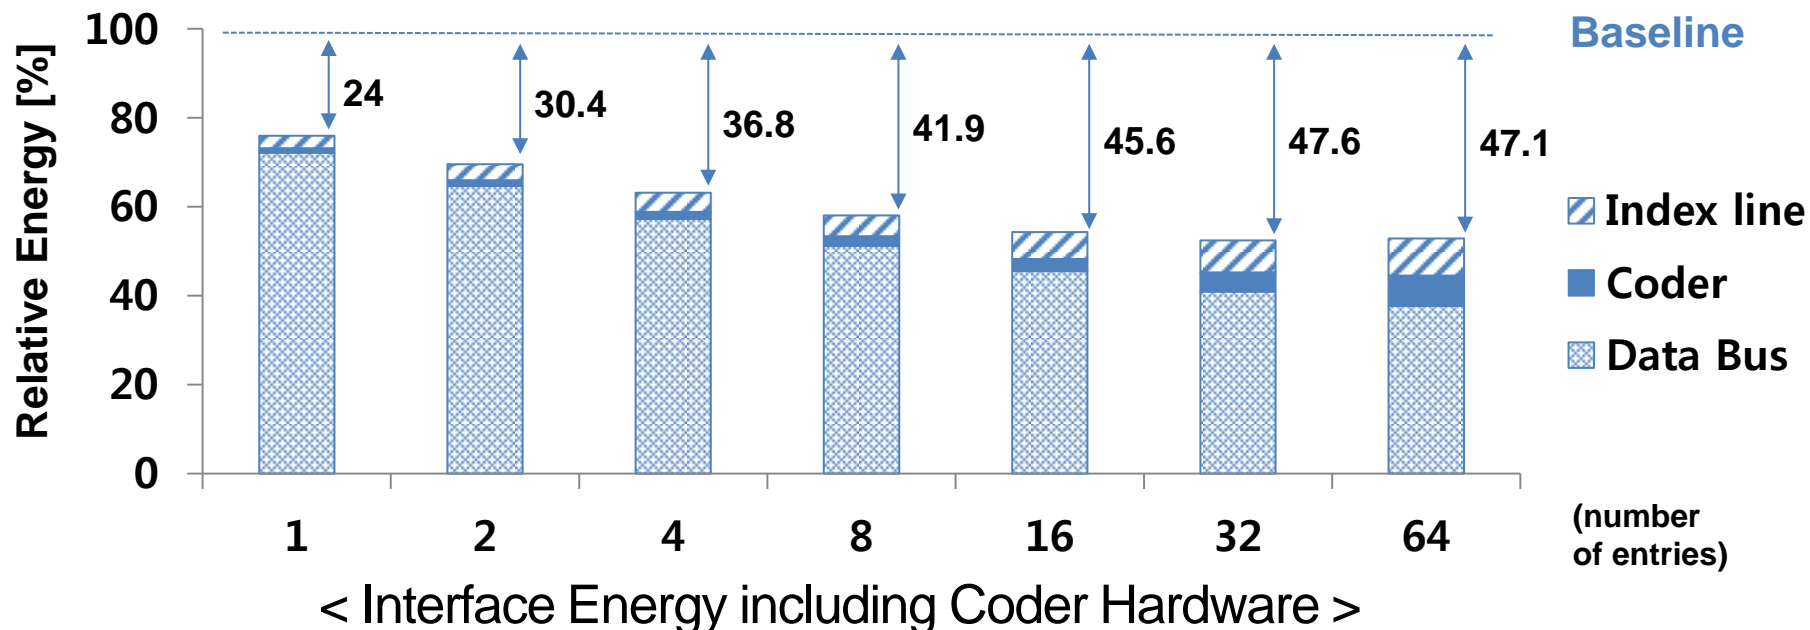
- **BD-Encoding decreases** the hamming weights in **all workloads** (the least effect in bzip: 29%)
- The results increase **as the number of table entries increases** (28-58% for 1-64 table entries)

Comparison to Prior Works



- **BD-encoding** reduces 58.3% of the termination and 45.3% of the switching energy.
- The probability for **similar data occurrence** is much higher than that for **the same data** ⇒ BD-encoding shows better results than Power Protocol and VALVE.

Interface Energy Reduction



- BD-encoding reduces overall interface energy **including coder hardware energy** (24-47.6% for 1-64 entries)
- **Optimal number of entries exists (32ea)** due to **overhead of index line and coder hardware.**

Outline

1. Introduction
2. BD-Encoding
3. Evaluation Results
- 4. Conclusion**

Conclusion

- **Reducing hamming weight** decreases both the **termination** and **switching** energy.
- Data Value Similarity: **Similar data words** are sent over the **DRAM data bus**.
- Bitwise Different Encoding: Transfer the **bit-wise difference** between a **current data word** and the **most similar data word recently transferred**.
- Evaluation Results: Reduce **58.3%** of **termination** and **45.3%** of **switching energy**.