ELPL Embedded Low-Power Laboratory

Embedded System Application

4190.303C

2010 Spring Semester

Memory Power Management

Naehyuck Chang
Dept. of EECS/CSE
Seoul National University
naehyuck@snu.ac.kr

Outline

- Concept of high-level power management
- SRAM power management
- SDRAM power states
- SDRAM mode control
- High-level memory power management

Concept of high-level power management

- Low-level energy optimization
  - Has been contributing over dozens of years
  - Enhancement of devices and components
  - General solution that applicable to almost all kinds of use
  - City bus service example
    - Objective: more gas mileage
    - New buses, engine swap, aluminum bodies, new transmissions, etc.
  - In the semiconductor world
    - NMOS
    - CMOS
    - MTCMOS

Gas-efficient engine

Light-weight bus

Low-loss transmission

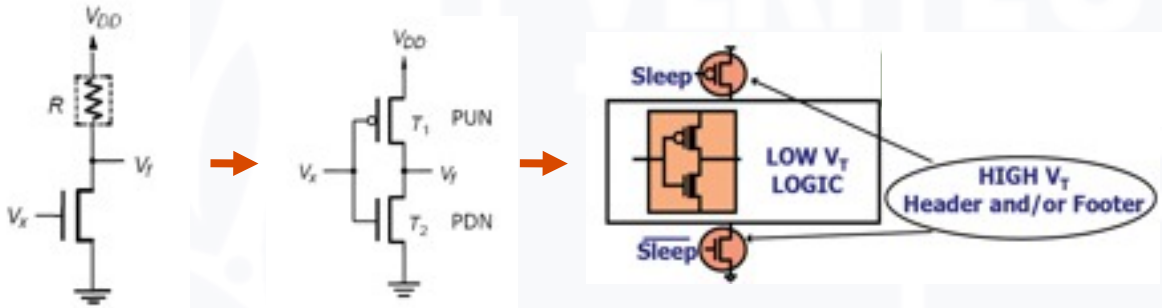

Concept of high-level power management

- System-software-level energy optimization
  - City bus service example
    - Optimal speed, engine rpm, shift position scheduling w/original hardware
      - Analysis of a target route
      - Use of component characteristics

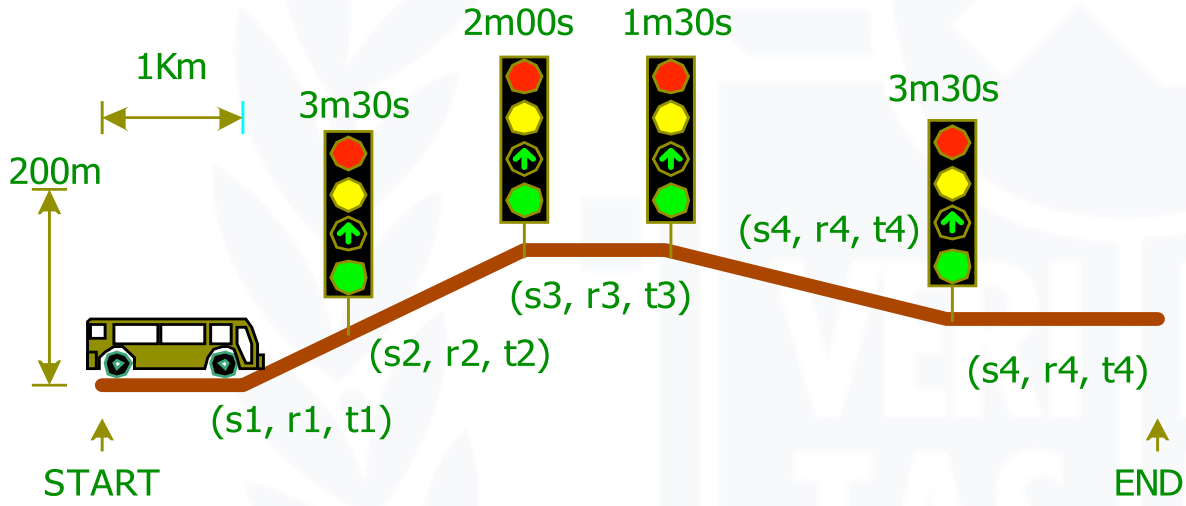

System-level approaches give us bigger chance to minimize energy consumption!

Concept of high-level power management

- Level of abstraction: engine idle gas consumption
  - Model 1: linear gas consumption per speed: g = mv
  - Model 2: counting idle gas consumption when v=0: g = mv + I
  - Model 3: counting engine restarting cost

- Applicable gas saving techniques when a vehicle is temporarily parked
  - **Technique 1**: linear gas consumption model
    - No policy when a vehicle is stopping
  - **Technique 2:** Idle gas consumption
    - Stop engine whenever a vehicle is stopped
  - **Technique 3:** Restarting cost
    - Stop engine when stopping time is more than 2 minutes for instance

Proper energy characterization is a primary concern of quality high-level power saving approach

Concept of high-level power management

- Components with multiple internal states
  - Each state has different functionality
  - Each state consumes different amount of power
  - Generally power consumption corresponds to service levels
- Conventionally abstracted as power state machines
  - State diagram with
    - Power and service annotation on states
    - Power and delay annotation on edges

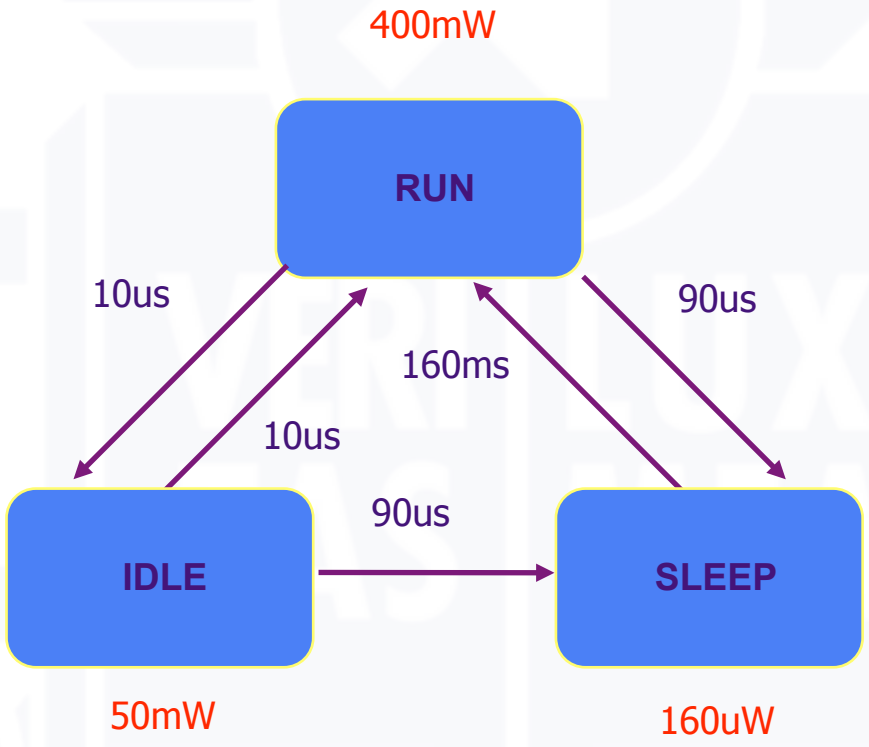

Concept of high-level power management

- Dynamic power management (DPM)
  - Reduce power according to workloads
  - Shutdown only during long idle time

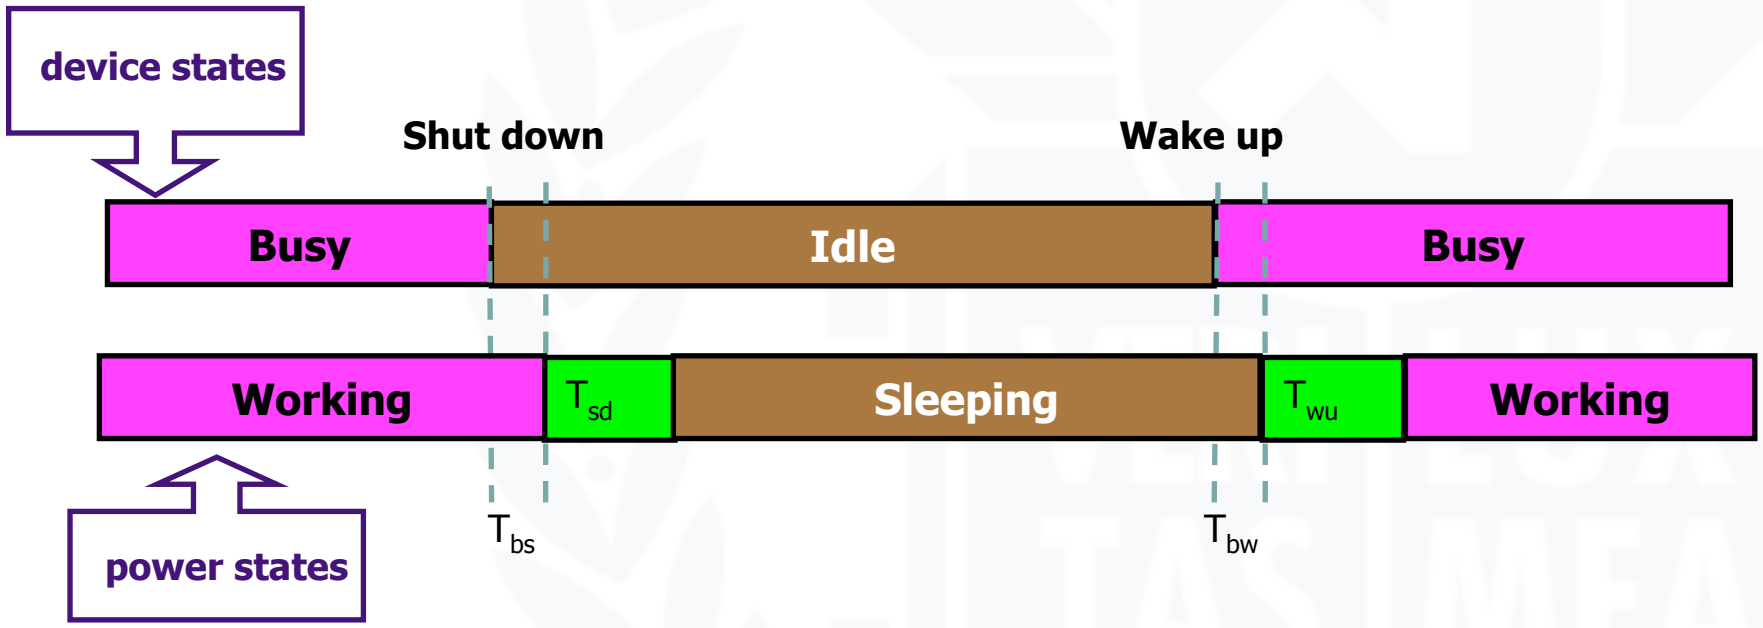

T_{sd}: shutdown delay

T_{wu}: wakeup delay

T_{bs}: time before shutdown

T_{bw}: time before wakeup

Concept of high-level power management

- Challenge
  - Predicting the future
- Break even time: Tbe
  - Shortest idle period for energy saving

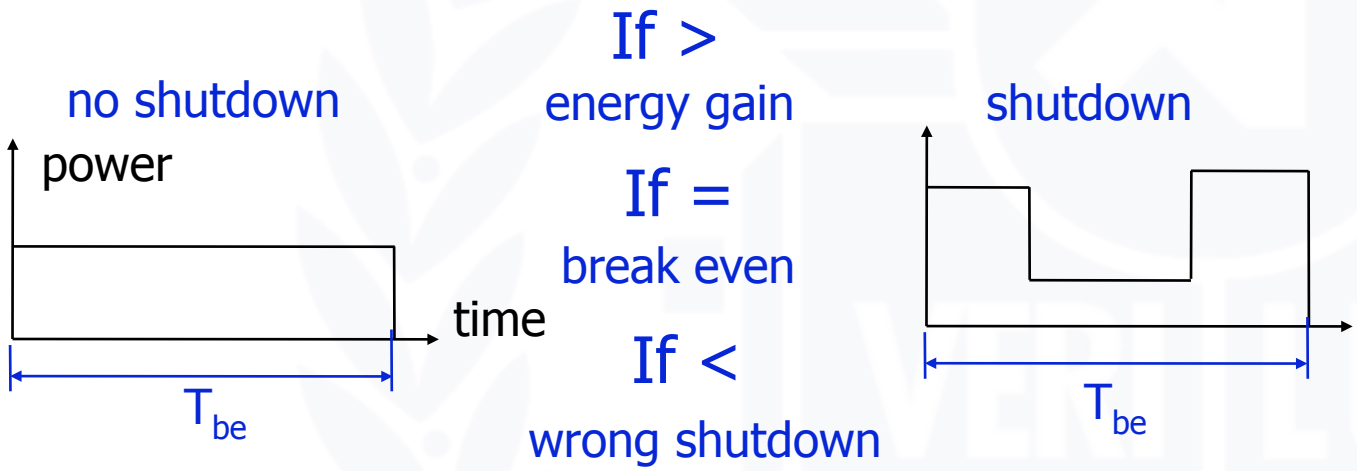

Idle period shorter than T_{be} is useless for energy saving

Concept of high-level power management

- Challenge
  - Predicting the future
- Break even time: Tbe
  - Shortest idle period for energy saving

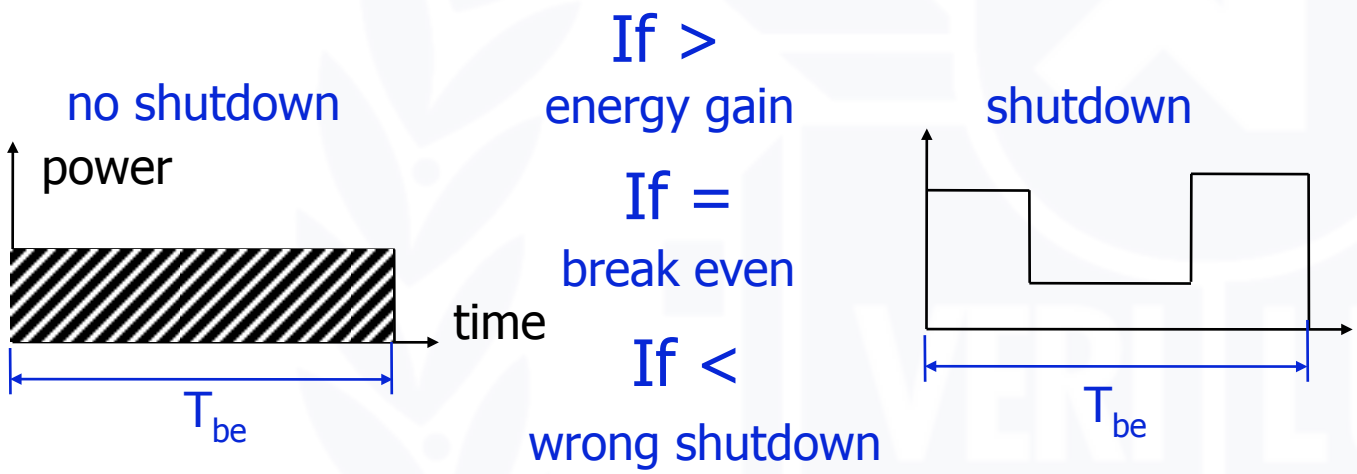

Idle period shorter than T_{be} is useless for energy saving

Concept of high-level power management

- Challenge
  - Predicting the future
- Break even time: Tbe
  - Shortest idle period for energy saving

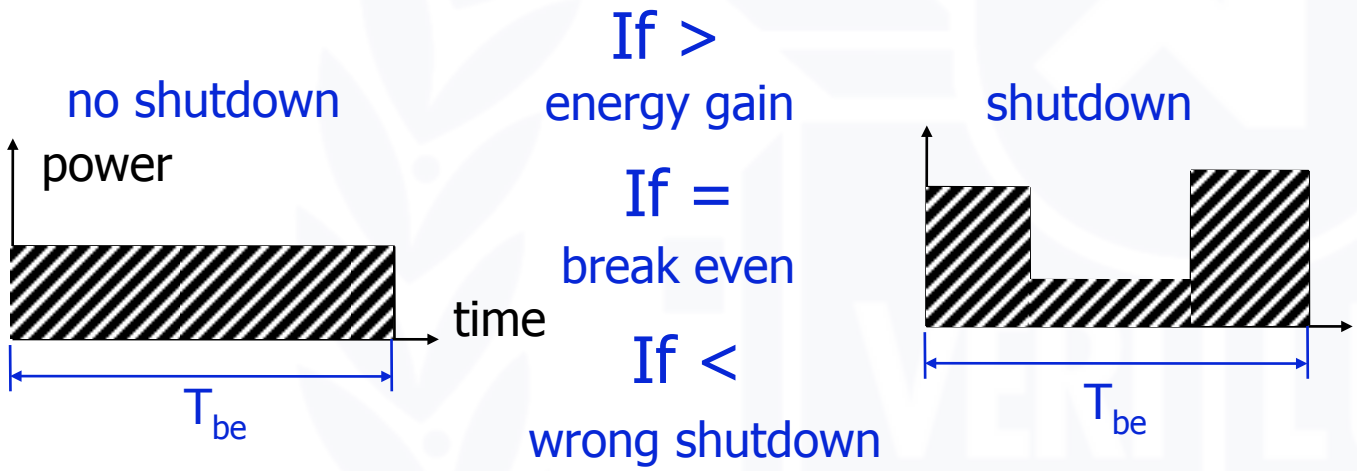

Idle period shorter than T_{be} is useless for energy saving

Concept of high-level power management

- When to use power management?
  - When $T_{BE} < T^{avg}_{idle}$
    - Average idle periods are long enough
    - Transition delay is short enough
    - Transition power is low enough
    - Sleep power is low enough
  - When designing system for a known workload
    - Criteria for component specification and design

SRAM Power Management

- Old-fashioned CMOS low-power SRAM
  - Two chip select pins
  - CS1*: active low
  - CS2: active high

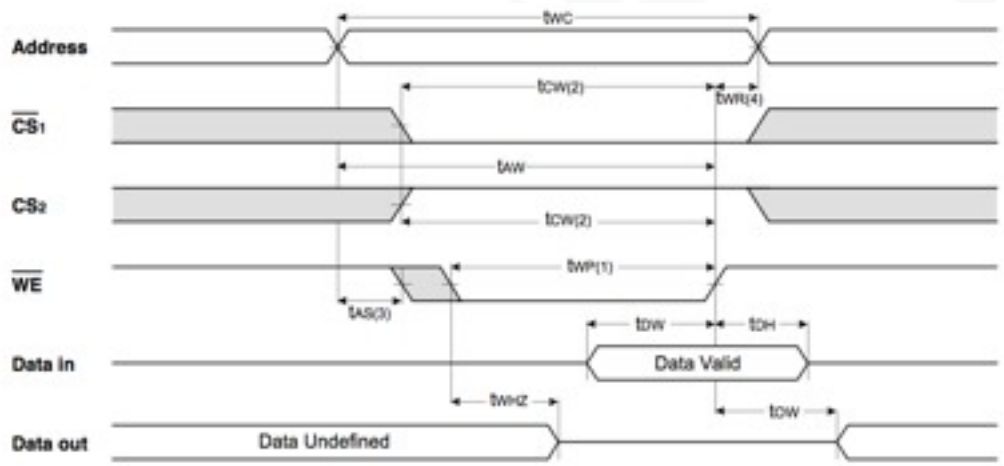

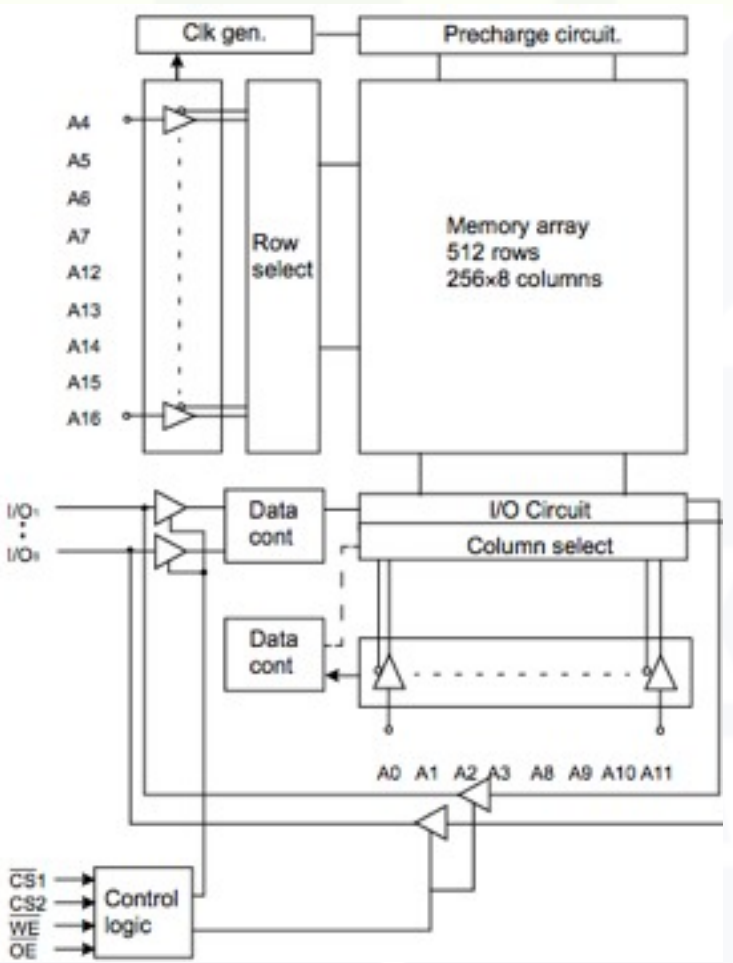

SRAM Power Management

- Low-power mode
  - Very low data retention current
  - Disable chip select and lower VDD
  - Recovery overhead exits

Item	Symbol	Test Condition		Min	Typ	Max	Unit
Vcc for data retention	V_{DR}	$\overline{CS_1}$[1]≥Vcc-0.2V		2.0	-	5.5	V
Data retention current	I_{DR}	Vcc=3.0V $\overline{CS_1}$≥Vcc-0.2V, CS_2≥Vcc-0.2V or CS_2≤0.2V Other Input+0~Vcc	KM681000BL KM681000BL-L	- -	1 0.5	50 10	µA
			KM681000BLE KM681000BLE-L	- -	- -	50 25	
			KM681000BLI KM681000BLI-L	- -	- -	50 25	
Data retention set-up time	t_{RDR}	See data retention waveform		0	-	-	ms
Recovery time	t_{RDR}			5	-	-	

SRAM Power Management

- How do send the SRAM to the data retention mode

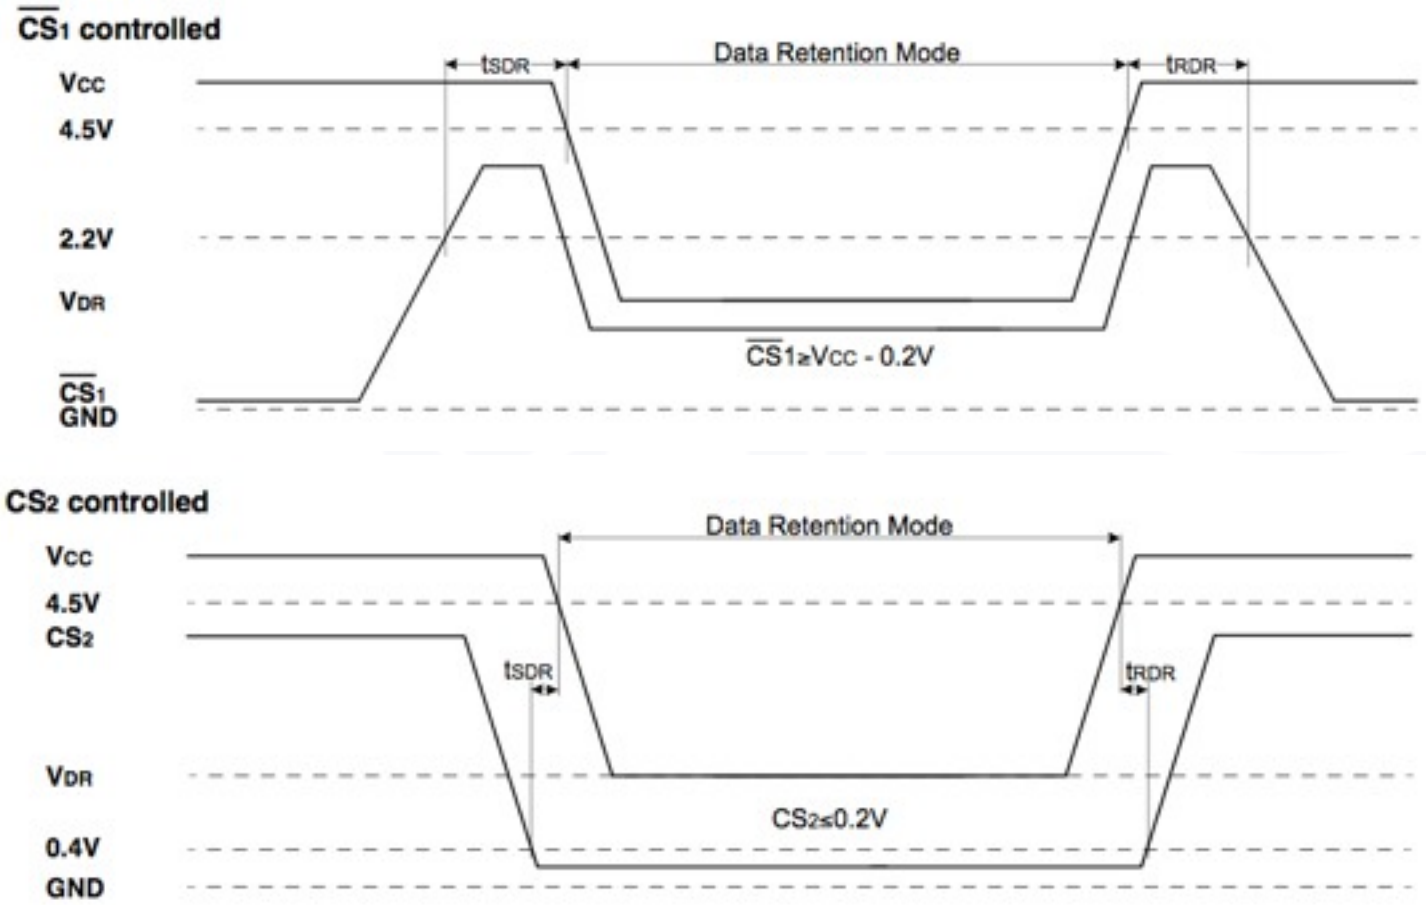

SRAM Power Management

- Battery backup
  - No true non-volatile, high-performance memory device exists
    - Non-volatile
    - Balanced read and write performance
    - Battery backed SRAM in an alternative solution
- Low-cost battery backup circuit
  - No battery charging sequence
  - Diode loss
  - SRAM operating VDD is lower than other device VDD

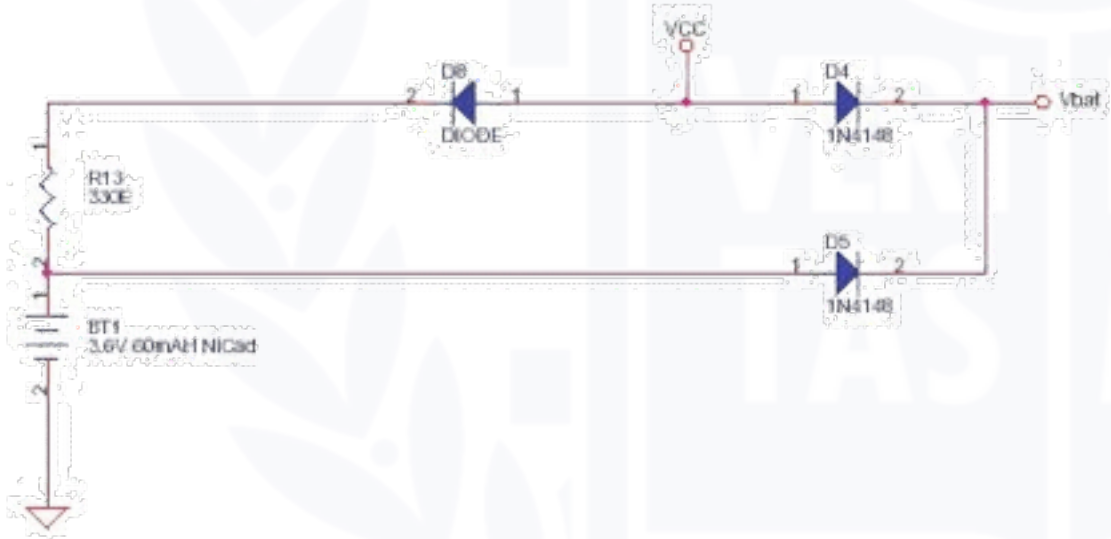

SRAM Power Management

- Non-rechargeable battery
  - Lithium battery
    - Very low self discharge
    - Very low data retention current
    - Dallas Semiconductor DS1259

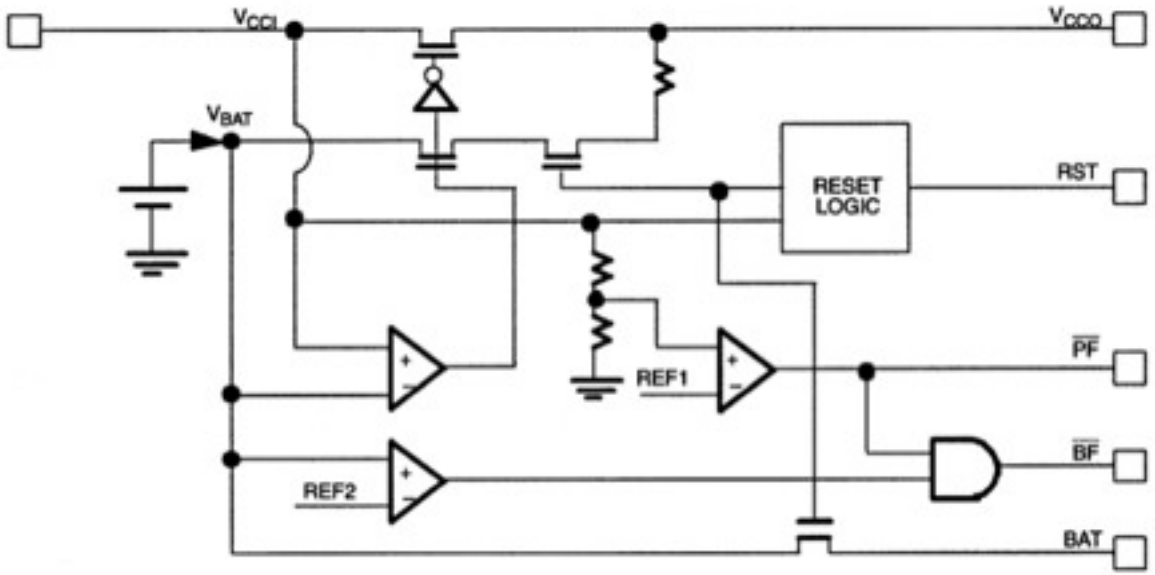

SRAM Power Management

- Rechargeable battery
  - Precision battery charge sequence
  - No diode loss

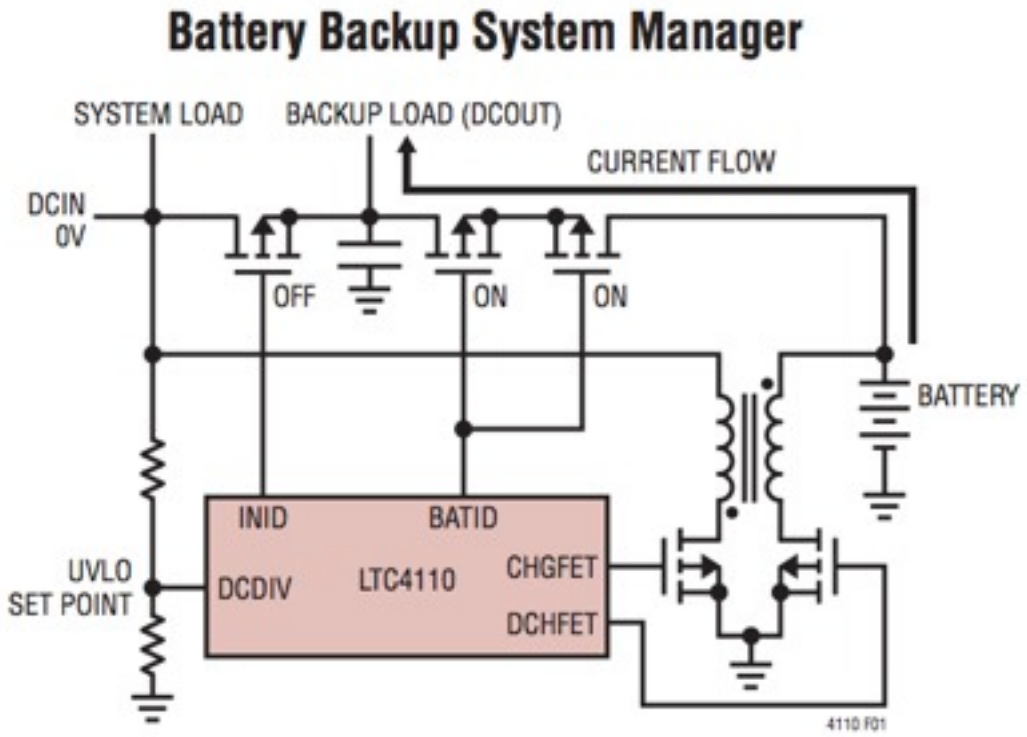

SDRAM Power States

- SDRAM power calculation

Power Calculation

Total Power = Core Power + I/O Power

$(I_{DD}4 \times V_{DD}) + (C \times f/2 \times V_{DD}Q^2 \times \text{number of I/Os} /2)$

Mobile SDRAM

$P = (90mA \times 2.5V) + (10pf \times \frac{100\ MHz}{2} \times 1.8V^2 \times 16 /2)$

$P = 238mW$

Standard SDRAM

$P = (150mA \times 3.3V) + (10pf \times \frac{100\ MHz}{2} \times 3.3V^2 \times 16 /2)$

$P = 538mW$

SDRAM Power States

- Power supply current (IDD) specification by the power states

I_{DD} Specifications and Conditions

PARAMETER/CONDITION		SYMBOL	SDR[1,3] (MAX)	128Mb Mobile SDRAM[2,4] (MAX)	UNITS
Operating Current: Active Mode; Burst = 2; READ or WRITE; ^{t}RC = ^{t}RC (MIN); CAS latency = 3		I_{DD1}	120	150	mA
Standby Current: Power-Down Mode; CKE = LOW; All banks idle		I_{DD2}	2	0.350	mA
Standby Current: Active Mode; CS# = HIGH; CKE = HIGH; All banks active after ^{t}RCD met; No accesses in progress		I_{DD3}	50	35	mA
Operating Current: Burst Mode; Continuous burst; READ or WRITE; All banks active, CAS latency = 3		I_{DD4}	150	90	mA
Auto Refresh Current: CAS latency = 3; CKE, CS# = HIGH	^{t}RFC = ^{t}RFC (MIN)	I_{DD5}	310	210	mA
Self Refresh Current: CKE 0.2V		I_{DD7}	2	0.100 to 0.35[5]	mA

NOTE: 1. Part Number MT48LC8M16A2TG-75@3.3V
2. Part Number MT48V8M16LFFC-8@2.5V
3. V_{DD}, $V_{DD}Q$ = +3.3V for x16 SDRAM
4. V_{DD} = +2.5V, $V_{DD}Q$ = 2.5V/1.8V
5. Using TCSR from 15°C to 70°C

SDRAM Power States

- Detailed power states

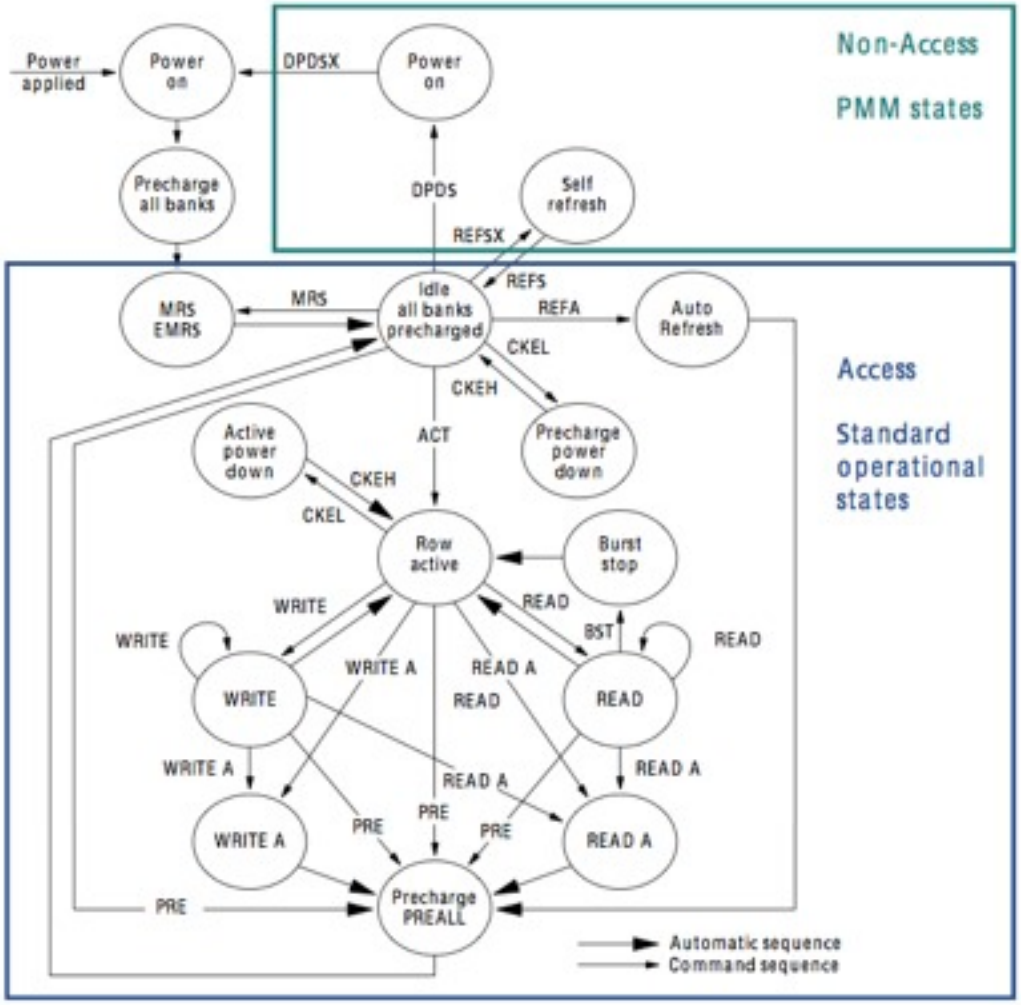

SDRAM Power States

- Low-level power management
  - The amount of current consumed is directly proportional to the self refresh rate

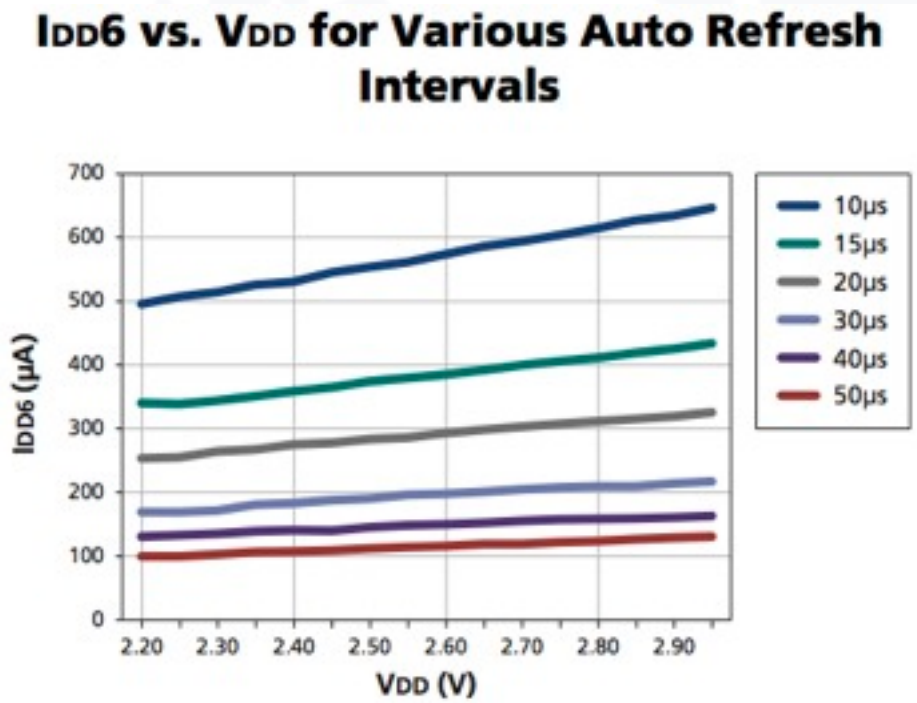

SDRAM Power States

- Temperature compensated self refresh (TCSR)
  - In self refresh operation, power can be saved if the internal self refresh intervals can be adjusted for the ambient temperature
  - Using a temperature sensor
  - Increased temperatures cause SDRAM cells to lose a charge at a faster rate

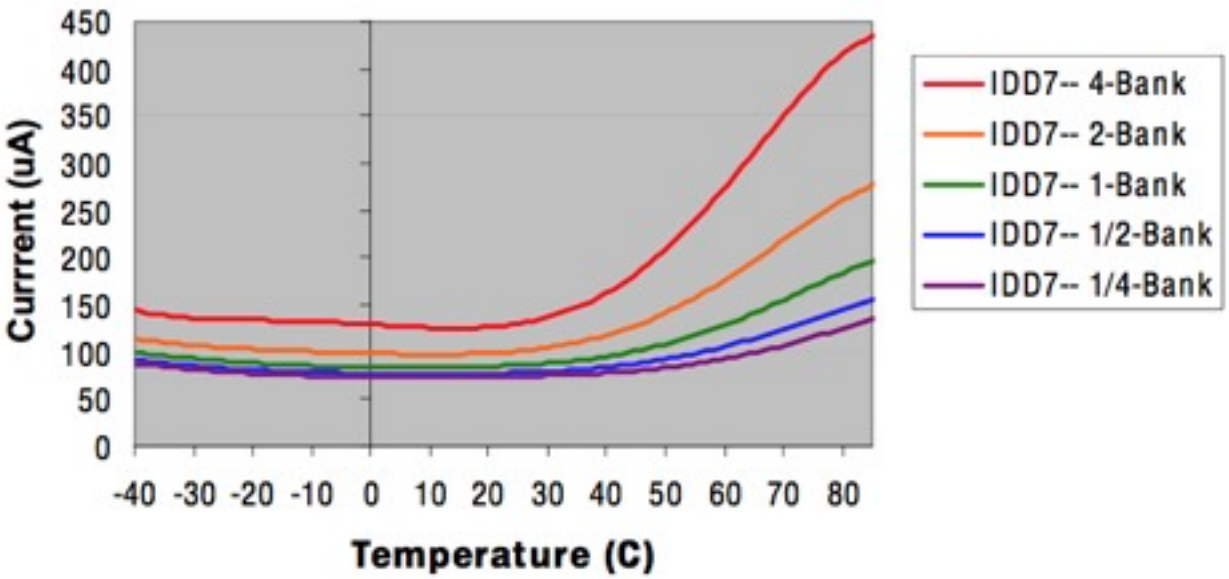

SDRAM Power States

- Partial array self refresh (PASR)
  - In self refresh operation, the refresh operation can be limited to the portion of the memory's array where data will be stored
    - If all of the array is not needed to store data
- Deep power-down (DPD)
  - In some applications, actual data retention in the DRAM is not required most of the time
    - DRAM can incorporate DPD to turn off most or all of the on-board array voltage generators
- How are they effective?

Typical Use Profile

Operation	Duty Cycle (Percentage of Clock Cycle)
Power Management Modes	
Deep Power-Down (DPD)	50%
Self Refresh (PASR)	30%
Standard SDRAM Modes	20%

SDRAM Power States

- SDRAM power down
  - In DDR2 SDRAM devices, power-down occurs when CKE is registered LOW with a DESELECT or NOP command

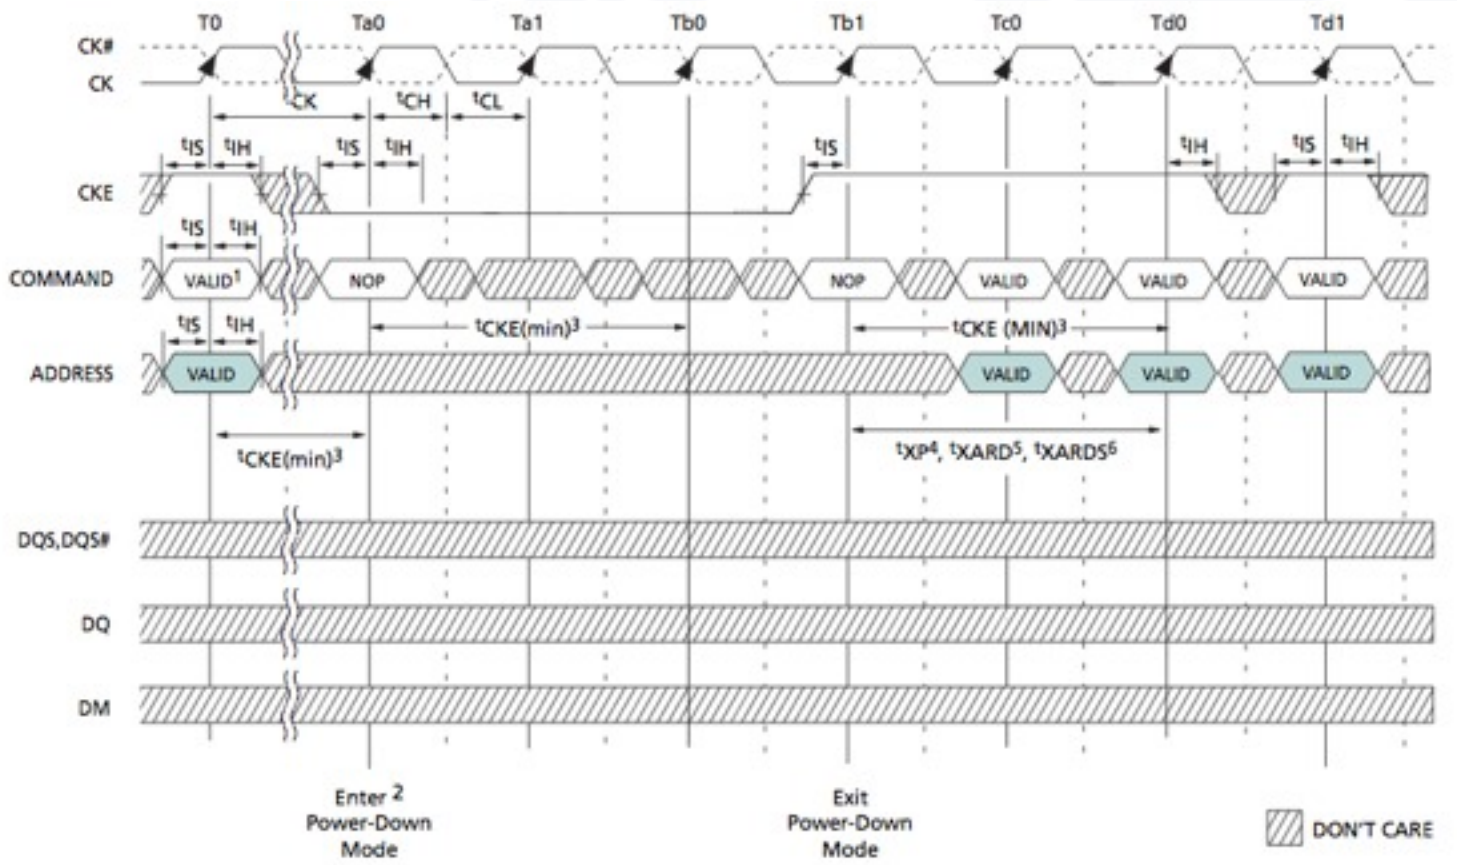

SDRAM Power States

- Precharge power down
  - If this command is a PRECHARGE (or if the device is already in the idle state)

Precharge power-down current: All banks idle; $^{t}CK = {}^{t}CK$ (I_{DD}); CKE is LOW; Other control and address bus inputs are stable; Data bus inputs are floating	I_{DD2P}	x4, x8, x16	5	5	5	mA
Precharge quiet standby current: All banks idle; $^{t}CK = {}^{t}CK$ (I_{DD}); CKE is HIGH, CS# is HIGH; Other control and address bus inputs are stable; Data bus inputs are floating	I_{DD2Q}	x4, x8	40	35	25	mA
		x16	50	35	25	

- Active power down
  - If this command is an ACTIVE (or if at least one row is already active)

Active power-down current: All banks open; $^{t}CK = {}^{t}CK$ (I_{DD}); CKE is LOW; Other control and address bus inputs are stable; Data bus inputs are floating	I_{DD3Pf}	Fast PDN exit MR12 = 0	30	25	20	mA
	I_{DD3Ps}	Slow PDN exit MR12 = 1	6	6	6	
Active standby current: All banks open; $^{t}CK = {}^{t}CK$ (I_{DD}), $^{t}RAS = {}^{t}RAS$ MAX (I_{DD}), $^{t}RP = {}^{t}RP$ (I_{DD}); CKE is HIGH, CS# is HIGH between valid commands; Other control and address bus inputs are switching; Data bus inputs are switching	I_{DD3N}	x4, x8	50	40	30	mA
		x16	55	40	30	

SDRAM Power States

- Fast exit
  - MR12 = 0
- Slow exit
  - MR12 = 1

AC Characteristics			-187E		-25E		-25		-3E		-3		-37E		-5E			
Parameter		Symbol	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Units	Notes
Power-Down: Exit active power-down to READ command	MR12 = 0	^{t}XARD	3	–	2	–	2	–	2	–	2	–	2	–	2	–	^{t}CK	18
	MR12 = 1		10 - AL	–	8 - AL	–	8 - AL	–	7 - AL	–	7 - AL	–	6 - AL	–	6 - AL	–	^{t}CK	18
Power-Down: Exit precharge power-down to any nonREAD command		^{t}XP	3	–	2	–	2	–	2	–	2	–	2	–	2	–	^{t}CK	18
Power-Down: CKE MIN HIGH/LOW time		^{t}CKE	MIN = 3 MAX = n/a														^{t}CK	18, 44

Active power-down current: All banks open; ^{t}CK = ^{t}CK (I_{DD}); CKE is LOW; Other control and address bus inputs are stable; Data bus inputs are floating	I_{DD3Pf}	Fast PDN exit MR12 = 0	30	25	20	mA
	I_{DD3Ps}	Slow PDN exit MR12 = 1	6	6	6	

SDRAM Power States

- Implementation

Extended Mode Register

BA1	BA0	A12	A11	A10	A9	A8	A7	A6	A5	A4	A3	A2	A1	A0
1	0	Reserved								TCSR[1]		PASR[2]		

A4	A3	Max. Case Temp.
0	0	70°C
0	1	45°C
1	0	15°C
1	1	85°C

A2	A1	A0	Self Refresh Coverage
0	0	0	All Banks
0	0	1	Banks 0 and 1 (BA1 = 0)
0	1	0	Bank 0 (BA1 = BA0 = 0)
0	1	1	Reserved
1	0	0	Reserved
1	0	1	Lower Half Bank 0[3,5]
1	1	0	Lower Quarter Bank 0[4,5]
1	1	1	Reserved

NOTE:
1. Temperature Compensated Self Refresh
2. Partial Array Self Refresh
3. Row Address 2 MSB = 0
4. Row Address MSB = 0
5. Available on future devices. Contact factory.

High-Level Memory Power Management

- SDRAM power down is effective power management
  - As far as power down decision making is appropriate
  - Break-even time should be carefully considered
  - Misprediction results in negative energy gain
- Other approaches for high-level memory system power management
  - Bus encoding
  - SDRAM mode control

Commercial Data Sheets

- Power values are described for a particular operating condition

Parameter/Condition	Symbol	Configuration	-3	-37E	-5E	Units
Operating one bank active-precharge current: $^{t}CK = {}^{t}CK (I_{DD})$, $^{t}RC = {}^{t}RC (I_{DD})$, $^{t}RAS = {}^{t}RAS\ MIN (I_{DD})$; CKE is HIGH, CS# is HIGH between valid commands; Address bus inputs are switching; Data bus inputs are switching	I_{DD0}	x4, x8	90	80	75	mA
		x16	90	80	75	
Operating one bank active-read-precharge current: I_{OUT} = 0mA; BL = 4, CL = CL (I_{DD}), AL = 0; $^{t}CK = {}^{t}CK (I_{DD})$, $^{t}RC = {}^{t}RC (I_{DD})$, $^{t}RAS = {}^{t}RAS\ MIN (I_{DD})$, $^{t}RCD = {}^{t}RCD (I_{DD})$; CKE is HIGH, CS# is HIGH between valid commands; Address bus inputs are switching; Data pattern is same as I_{DD4W}	I_{DD1}	x4, x8	100	90	85	mA
		x16	100	90	85	
Precharge power-down current: All banks idle; $^{t}CK = {}^{t}CK$ ... CKE is LOW; Other control and address bus inputs are switching	I_{DD2P}	x4, x8, x16	5	5	5	mA
Precharge ... $^{t}CK = {}^{t}CK$ ... control and address bus inputs are switching	...	...	...	...	...	...

I_{DD} Parameters	-3	-37E	-5E	Units
CL (I_{DD})	5	4	3	^{t}CK
$^{t}RCD (I_{DD})$	15	15	15	ns
$^{t}RC (I_{DD})$	60	60	55	ns
$^{t}RRD (I_{DD})$ - x4/x8 (1KB)	7.5	7.5	7.5	ns
$^{t}RRD (I_{DD})$ - x16 (2KB)	10	10	10	ns
$^{t}CK (I_{DD})$	3	3.75	5	ns

Micron Power Calculator

- Three steps
  - Power subcomponents are calculated based on data sheet specifications
  - Power is derated based on the command scheduling in the system
  - Power is derated to the system's actual operating VDD and clock frequency
- Power derating
  - Power calculations have assumed a system operating at worst-case VDD (Psch)
  - Clock frequency in the system is the same as the frequency defined in the data sheet (Psys)
  - However, most systems operate at different clock frequencies or operating voltages than the ones defined in the data sheet
- VDD derating

$$Psys(XXX) = Psch(XXX) \times \left(\frac{\text{use } V_{DD}}{\text{Max spec } V_{DD}}\right)^2$$

- Clock frequency derating

$$Psys(XXX) = Psch(XXX) \times \frac{use\_freq}{spec\_freq}$$

Micron Power Calculator

- Some parameters are frequency dependent while others are not

$$Psys(PRE\_PDN) = Psch(PRE\_PDN) \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(ACT\_PDN) = Psch(ACT\_PDN) \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(PRE\_STBY) = Psch(PRE\_STBY)\left[\frac{use\_freq}{spec\_freq}\right] \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(ACT\_STBY) = Psch(ACT\_STBY)\left[\frac{use\_freq}{spec\_freq}\right] \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(ACT) = Psch(ACT) \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(WR) = Psch(WR) \times \left[\frac{use\_freq}{spec\_freq}\right] \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(RD) = Psch(WRRD) \times \left[\frac{use\_freq}{spec\_freq}\right] \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

$$Psys(REF) = Psch(REF) \times \left[\frac{use\ V_{DD}}{Max\ spec\ V_{DD}}\right]^2$$

Energy State Machine

- Energy state machine
  - Energy state machine is a finite state machine
  - States denote static power consumption
  - Transitions denote dynamic power consumption

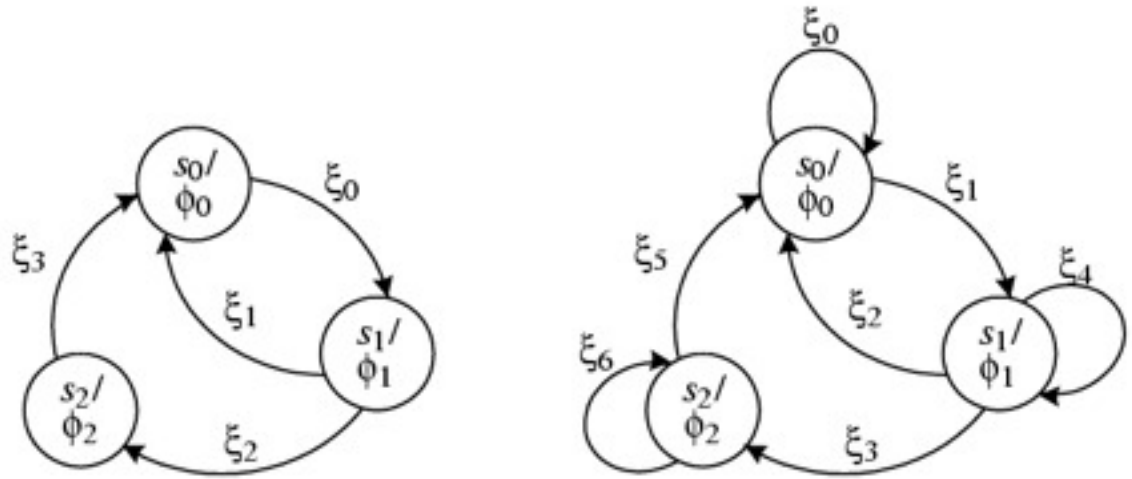

(a) asynchronous energy state machine

(b) synchronous energy state machine

Energy State Machine

- Energy state machine (Memory device)
  - SRAM
  - DRAM
  - SDRAM
  - NAND Flash
  - etc.

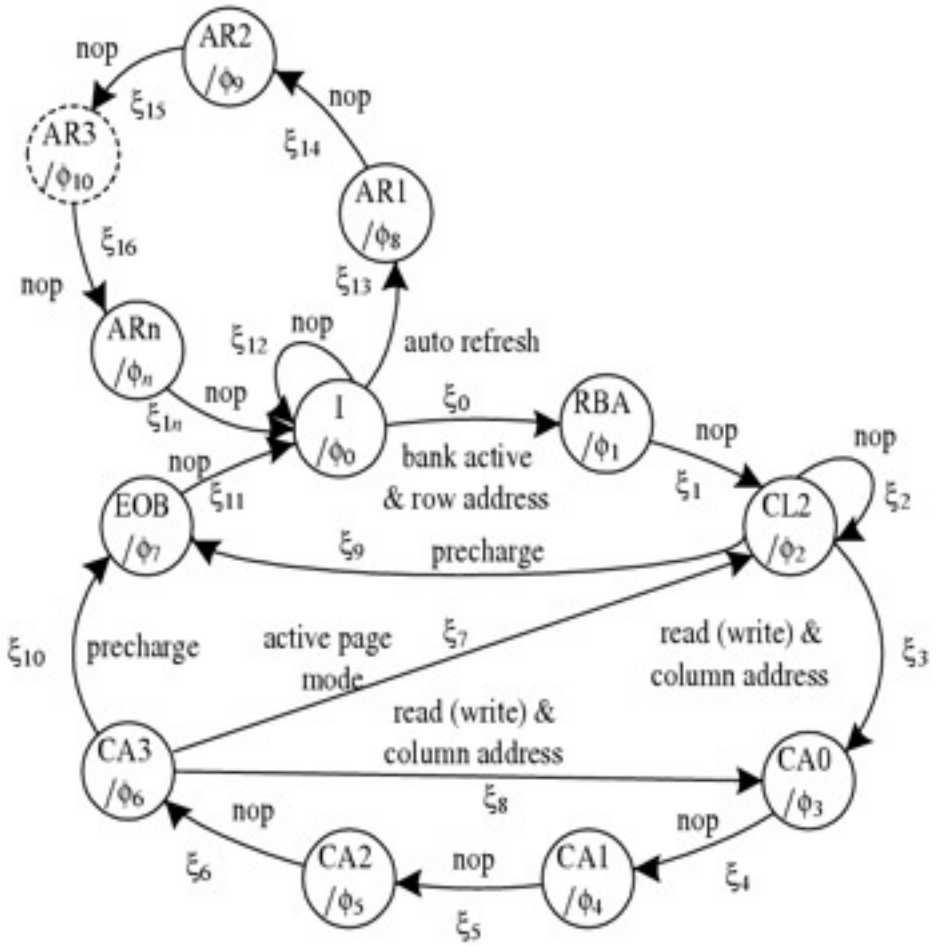

Energy State Machine

- Annotation of energy state machine (Event-accurate energy measurement)

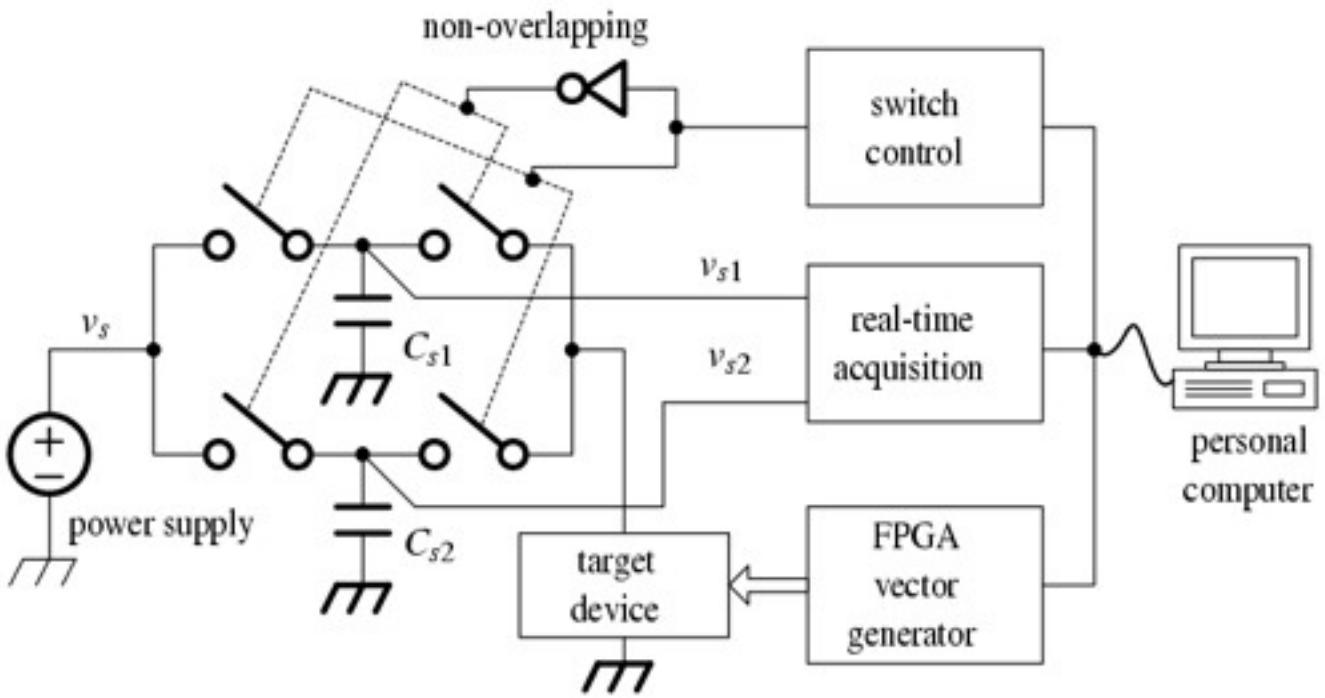

Energy State Machine

- Annotation of energy state machine (Leakage energy consumption triggered by an asynchronous strobe signal)
  - Leakage energy consumption is denoted by

$$\phi_j = \frac{C_s}{2}\frac{v_s^2(t+\Delta t) - v_s^2(t)}{\Delta t}$$

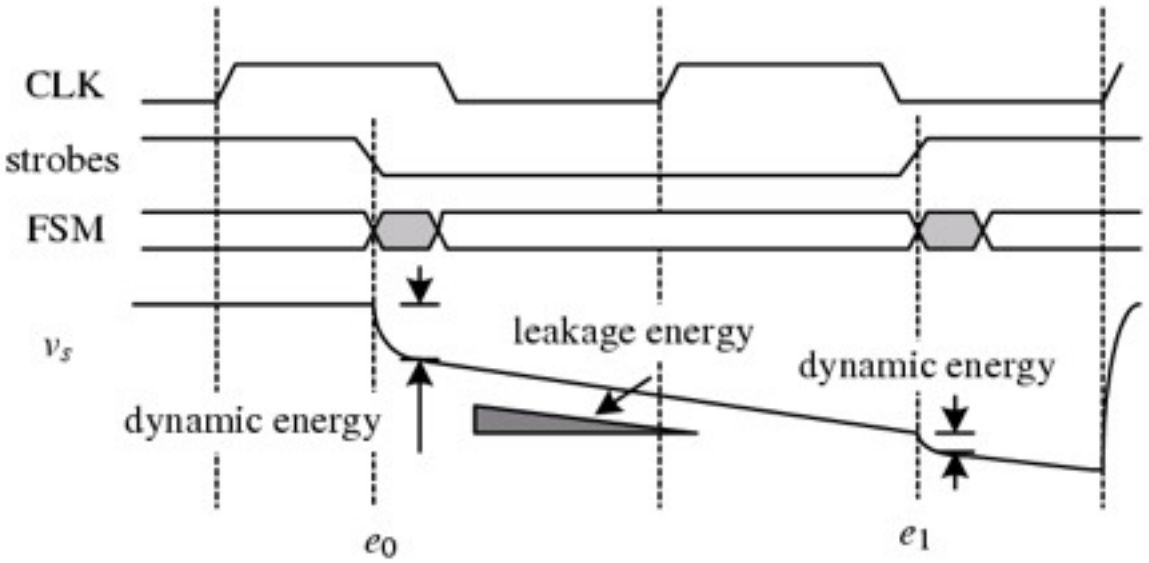

Memory Bus Encoding

- Bus encoding concept

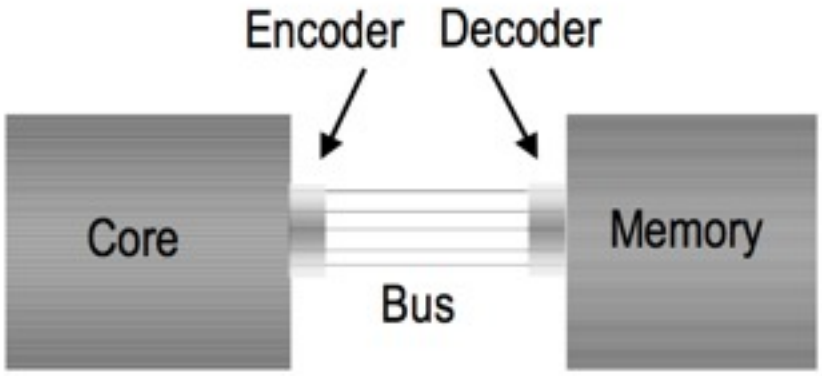

- Bus power cost
  - Power cost
    - HDD: Hamming distance dependent power
    - WDS: Weight dependent static power
  - HDD
    - Due to the switching capacitance of the load and PCB track
  - WDS
    - Due to the static power of the bus driver and termination resistors

Memory Bus Encoding

- Clock frequencies of processor and memory system doesn't affect the number of transactions
- Execution time decreases as clock frequency increases

Parameter		Address bus		Data bus	
		HDD	WDS	HDD	WDS
Processor clock (f_P) MHz	100	1.27/4.46	11.79/41.36	4.13/14.48	1.05/3.67
	133	1.27/5.71	9.20/41.32	4.13/18.54	1.05/4.70
	200	1.27/7.97	6.58/41.24	4.13/25.86	1.05/6.55
	266	1.27/9.90	5.29/41.17	4.13/32.13	1.05/8.14
	400	1.27/13.06	4.00/41.04	4.13/42.37	1.05/10.74
	800	1.27/19.03	2.73/40.78	4.13/61.73	1.05/15.64
f_M @66MHz, 8KB/4word/2-way-set-associative cache					
Memory (f_M) clock MHz	33	1.27/7.75	6.72/40.94	4.13/25.16	2.09/12.75
	66	1.27/9.90	5.29/41.17	4.13/32.13	1.05/8.14
	83	1.27/10.50	5.00/41.23	4.13/34.05	0.84/6.90
	100	1.27/10.68	4.92/41.31	4.13/34.63	0.70/5.85
	133	1.27/11.22	4.69/41.35	4.13/36.41	0.52/4.61
f_P @266MHz, 8KB/4word/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

Memory Bus Encoding

- Static power consumption

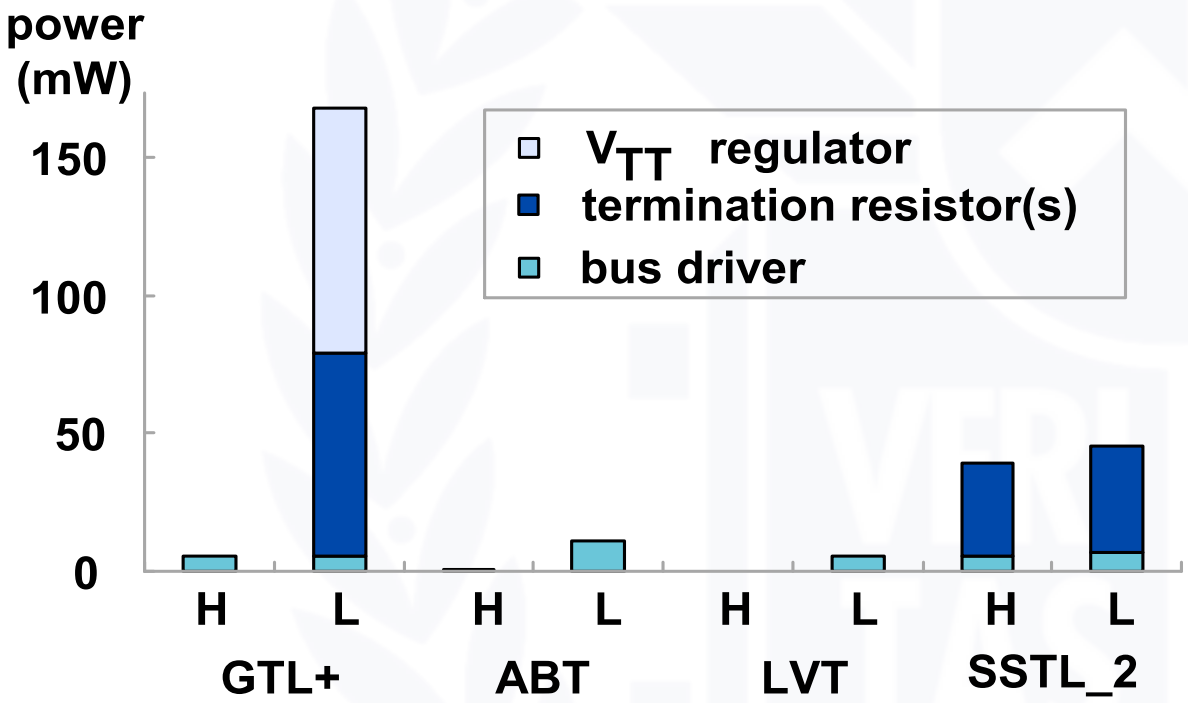

Memory Bus Encoding

- Power consumption coefficients

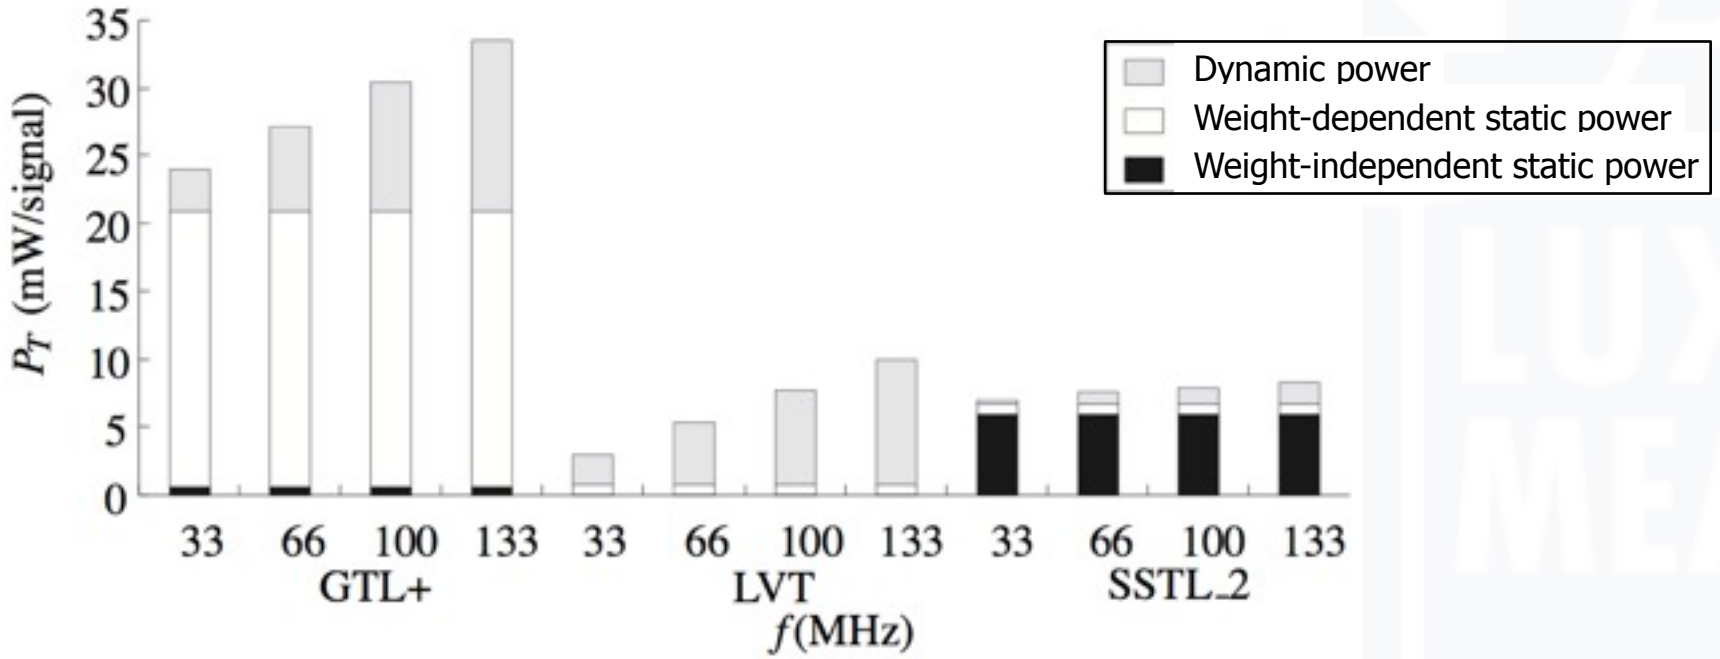

Memory Bus Encoding

- Bus invert coding
  - Original bus signal + INV signal
  - If the next bus power cost is high, invert the bus data and enable INV signal

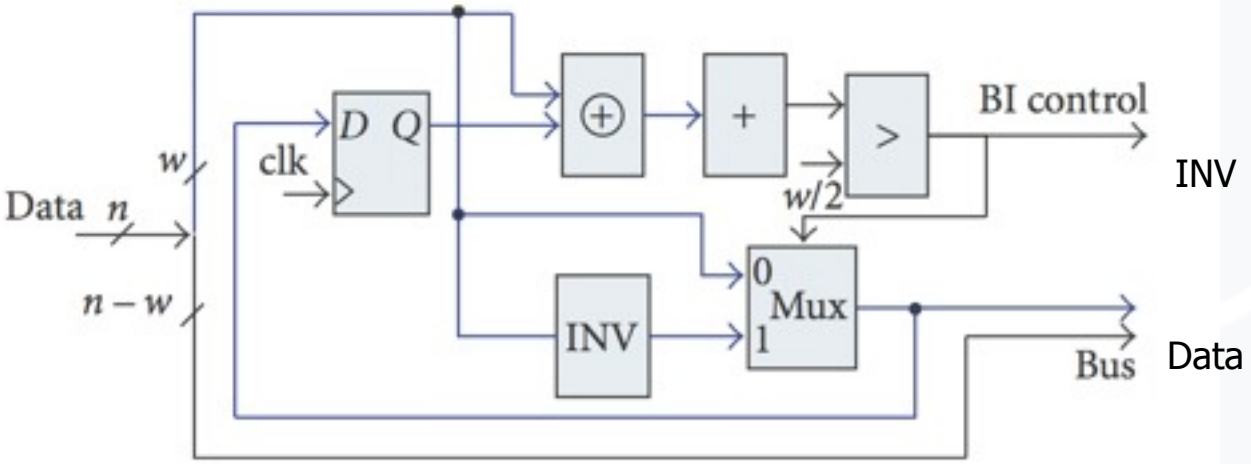

  - Example
    - 00000000, 11111111 → 00000000 (0), 00000000 (1)
  - Inversion decision is the key
    - HDD based, WDS based, or both

Memory Bus Encoding

- HDD and WDS based inversion decision contract with each other

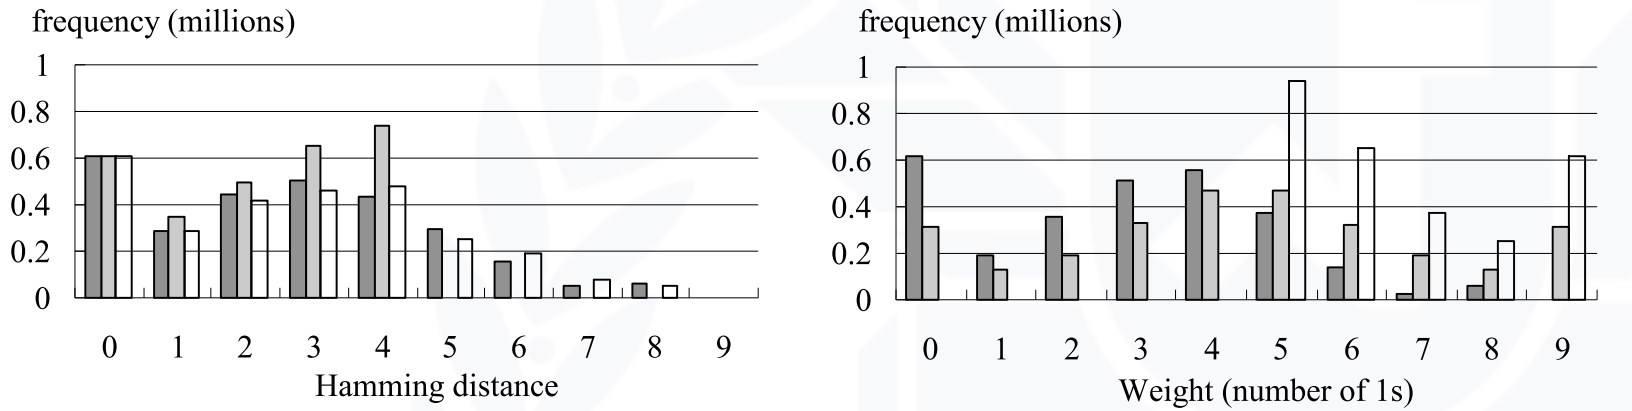

(c) JPEG decompressor (DJPEG)

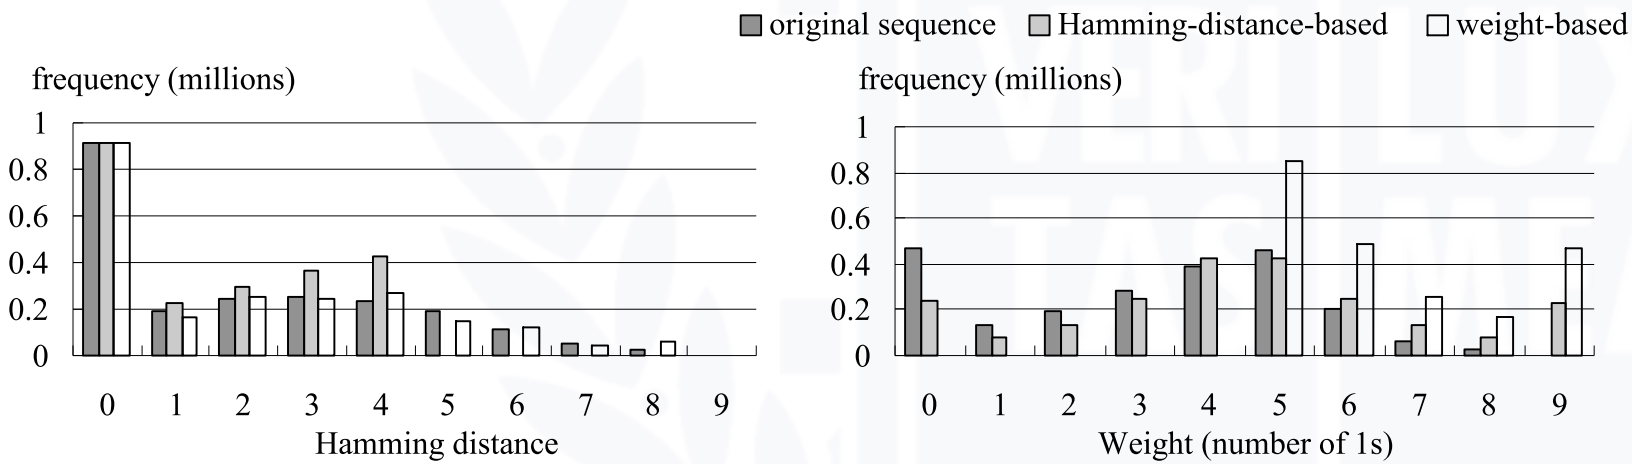

(d) MPEG4 decoder

Memory Bus Encoding

- In the view of energy consumption of the data bus
  - Weight-based bus-invert coding scheme is superior for MP3 decoder
  - Hamming-distance-based bus-invert coding scheme is superior for JPEG compressor, JPEG decompressor and MPEG4 decoder

Application	Decision scheme	HDD	WDS	Total	Reduction ratio (%)
MP3	no encoding	0.18	0.40	0.58	0.0
	H-based	0.13	0.20	0.33	43.9
	W-based	0.16	0.02	0.18	69.9
CJPEG	no encoding	4.13	1.05	5.18	0.0
	H-based	3.37	1.01	4.38	15.4
	W-based	4.28	0.57	4.85	6.4
DJPEG	no encoding	4.15	1.17	5.32	0.0
	H-based	3.37	1.02	4.39	17.5
	W-based	4.25	0.53	4.78	10.0
MPEG4	no encoding	2.50	0.85	3.35	0.0
	H-based	1.95	0.80	2.75	17.8
	W-based	2.58	0.44	3.02	9.9

24M instructions, (unit: mJ)

Energy Reduction Practices

- Larger cache block size increases the HDD energy of the data bus, but decreases the HDD energy of the address bus

Parameter		Address bus		Data bus	
		HDD	WDS	HDD	WDS
Cache size KB	1	5.78/22.44	10.09/39.18	16.44/63.83	4.48/17.40
	2	4.63/20.55	8.82/39.20	13.70/60.84	3.41/15.13
	4	3.91/19.01	8.05/39.15	12.14/59.05	2.73/13.29
	8	1.27/9.90	5.29/41.17	4.13/32.13	1.05/8.14
	16	0.62/5.59	4.65/42.03	2.17/19.61	0.53/4.79
f_M @66MHz, f_P @266MHz, 4word/2-way-set-associative cache					
Cache set	1	1.30/10.04	5.31/41.15	4.11/31.86	1.14/8.86
	2	1.27/9.90	5.29/41.17	4.13/32.13	1.05/8.14
	4	1.36/10.36	5.40/41.13	4.45/33.90	1.09/8.27
	8	1.47/10.87	5.60/41.46	4.81/35.63	1.15/8.52
f_M @66MHz, f_P @266MHz, 8KB/4word cache					
Block size	4	1.27/9.90	5.29/41.17	4.13/32.13	1.05/8.14
	8	0.80/6.68	5.13/42.58	5.15/42.79	1.48/12.28
f_M @66MHz, f_P @266MHz, 8KB/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

DRAM Mode Control

- Auto-precharge

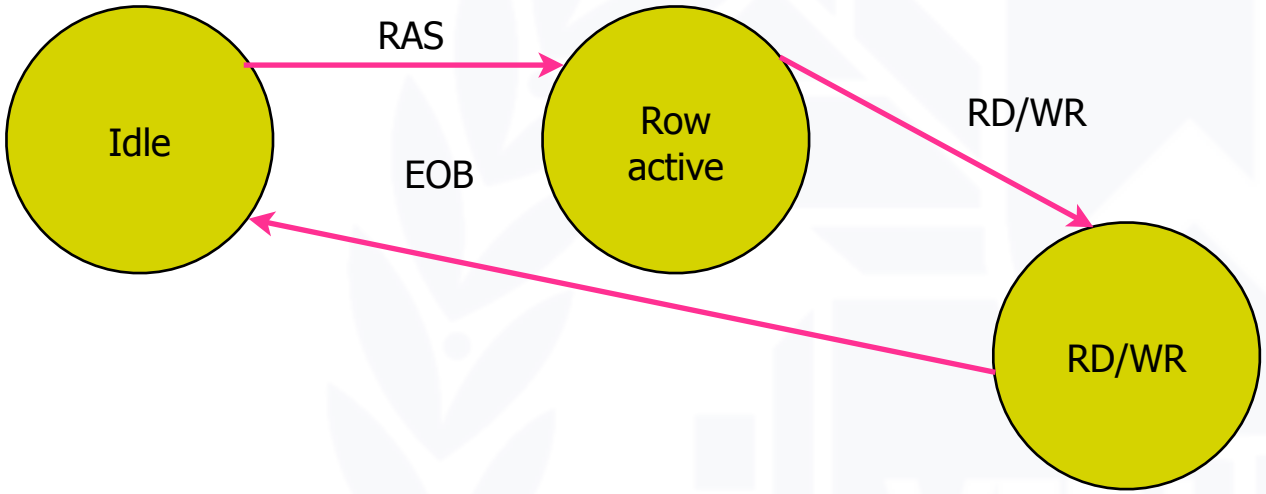

- After a burst-mode access, the controller closes the row, and the SDRAM remains in the idle mode

DRAM Mode Control

- Active page

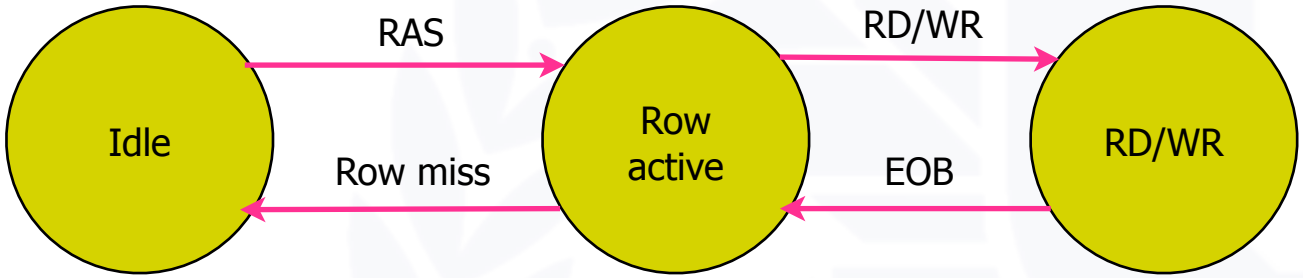

- The SDRAM may also remain in the row-active state after a burst mode access
- If the row address of the next access is equal to the current one (row hit), the controller does not need to re-issue the row address, and thus data can be directly forwarded to the sense amplifier
- If the next access refers the different row address (row miss) while the SDRAM remains in the row-active mode expecting the next access will hit the same row, the controller need to close the row and re-open a new row

DRAM Mode Control

- Auto-precharge policy for a mid-performance system
  - CL: common-mode leakage power
  - CD: common-mode dynamic power

Parameter		CL	CD	WDD	Total
Processor clock (f_P) MHz	100	19.0/66.6	65.8/230.8	2.3/8.1	87.1/305.5
	133	15.3/68.9	63.5/285.2	2.3/10.4	81.2/364.4
	200	11.7/73.0	61.2/383.1	2.3/14.5	75.1/470.6
	266	9.8/76.5	60.0/467.1	2.3/18.0	72.2/561.6
	400	8.0/82.3	58.9/604.3	2.3/23.7	69.2/710.3
	800	6.2/93.2	57.8/863.6	2.3/34.5	66.3/991.3
f_M @66MHz, 8KB/4word/2-way-set-associative cache					
Memory (f_M) clock MHz	33	14.2/86.8	61.1/372.3	2.3/14.1	77.7/473.2
	66	9.8/76.5	60.0/467.1	2.3/18.0	72.2/561.6
	83	8.9/73.8	59.8/493.3	2.3/19.0	71.1/586.1
	100	8.8/73.7	59.7/501.1	2.3/19.4	70.8/594.2
	133	8.0/70.5	59.6/525.4	2.3/20.4	69.9/616.2
f_P @266MHz, 8KB/4word/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

DRAM Mode Control

- Auto-precharge policy for a mid-performance system
  - CL: common-mode leakage power
  - CD: common-mode dynamic power

Parameter		CL	CD	WDD	Total
Cache size KB	1	24.7/96.1	275.8/1070.9	10.8/41.8	311.3/1208.7
	2	20.9/93.0	222.9/990.0	9.0/40.1	252.9/1123.1
	4	18.6/90.6	191.4/930.7	8.0/39.0	218.1/1060.4
	8	9.8/76.5	60.0/467.1	2.3/18.0	72.2/561.6
	16	7.6/69.1	30.0/270.9	1.1/9.8	38.7/349.8
f_M @66MHz, f_P @266MHz, 4word/2-way-set-associative cache					
Cache set	1	9.9/77.0	60.6/469.7	2.2/17.2	72.8/563.9
	2	9.8/76.5	60.0/467.1	2.3/18.0	72.2/561.6
	4	10.1/77.2	64.9/494.3	2.5/19.3	77.6/590.8
	8	10.5/78.1	70.8/524.5	2.8/20.8	84.2/623.4
f_M @66MHz, f_P @266MHz, 8KB/4word cache					
Block size	4	9.8/76.5	60.0/467.1	2.3/18.0	72.2/561.6
	8	9.5/78.9	66.2/550.0	2.2/17.8	77.9/646.8
f_M @66MHz, f_P @266MHz, 8KB/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

DRAM Mode Control

- Auto-precharge policy for a high-performance system
  - CL: common-mode leakage power
  - CD: common-mode dynamic power

Parameter		CL	CD	WDD	Total
Processor clock (f_P) MHz	133	27.0/133.6	25.7/126.9	1.1/5.2	55.1/272.5
	200	19.6/140.6	26.4/182.8	1.0/7.5	47.1/337.6
	266	15.7/145.5	26.1/236.2	1.0/9.7	42.8/398.1
	400	11.7/153.6	25.8/331.7	1.0/13.7	38.6/505.6
	800	7.3/163.3	25.5/560.9	1.0/23.2	33.9/753.7
	1000	6.4/165.8	25.4/651.0	1.0/26.9	32.9/850.0
f_M @100MHz, 16KB/4word/2-way-set-associative cache					
Memory clock (f_M) MHz	33	15.3/155.6	26.0/258.4	1.0/10.6	42.3/430.9
	66	12.5/155.1	25.8/315.1	1.0/13.0	39.3/489.6
	83	11.8/154.2	25.8/330.2	1.0/13.6	38.7/504.6
	100	11.7/153.6	25.8/331.7	1.0/13.7	38.6/505.6
	133	11.2/153.1	25.8/346.9	1.0/14.3	38.0/521.0
f_P @400MHz, 16KB/4word/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

DRAM Mode Control

- Auto-precharge policy for a high-performance system
  - CL: common-mode leakage power
  - CD: common-mode dynamic power

Parameter		CL	CD	WDD	Total
Cache size KB	1	31.0/172.2	240.9/1332.2	9.7/54.0	281.6/1564.9
	2	26.9/172.1	194.0/1233.3	8.1/52.0	229.1/1464.0
	4	24.4/170.3	166.9/1158.3	7.3/50.7	198.6/1385.8
	8	14.9/164.9	53.5/585.9	2.2/24.2	70.6/781.5
	16	11.7/153.6	25.8/331.7	1.0/13.7	38.6/505.6
f_M @100MHz, f_P @400MHz, 4word/2-way-set-associative cache					
Cache set	1	12.7/163.9	28.3/358.6	1.1/14.5	42.1/543.5
	2	11.7/153.6	25.8/331.7	1.0/13.7	38.6/505.6
	4	11.3/149.5	25.2/325.0	1.0/13.4	37.5/494.5
	8	11.1/147.2	24.0/311.8	1.0/13.0	36.0/478.6
f_M @100MHz, f_P @400MHz, 16KB/4word cache					
Block size	4	11.7/153.6	25.8/331.7	1.0/13.7	38.6/505.6
	8	10.3/144.6	24.7/341.1	0.8/11.6	35.8/504.0
f_M @100MHz, f_P @400MHz, 16KB/2-way-set-associative cache					

JPEG compressor, 24M instructions, (unit: mJ/mW)

DRAM Mode Control

- Conventional CRB (Column, Row, Bank) address alignment

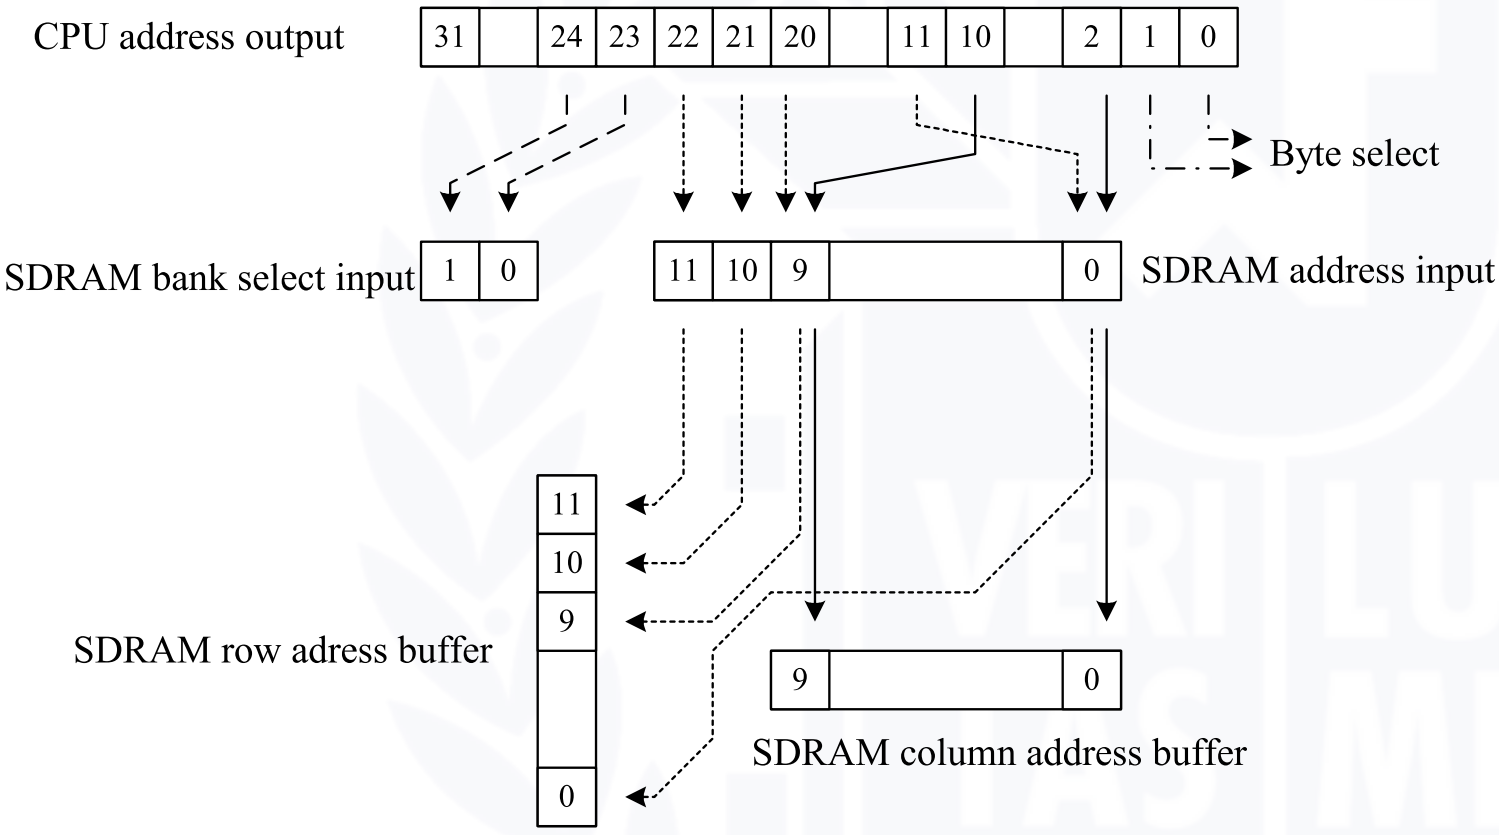

DRAM Mode Control

- CBR (Column, Bank, Row) address alignment for higher row-hit ratio

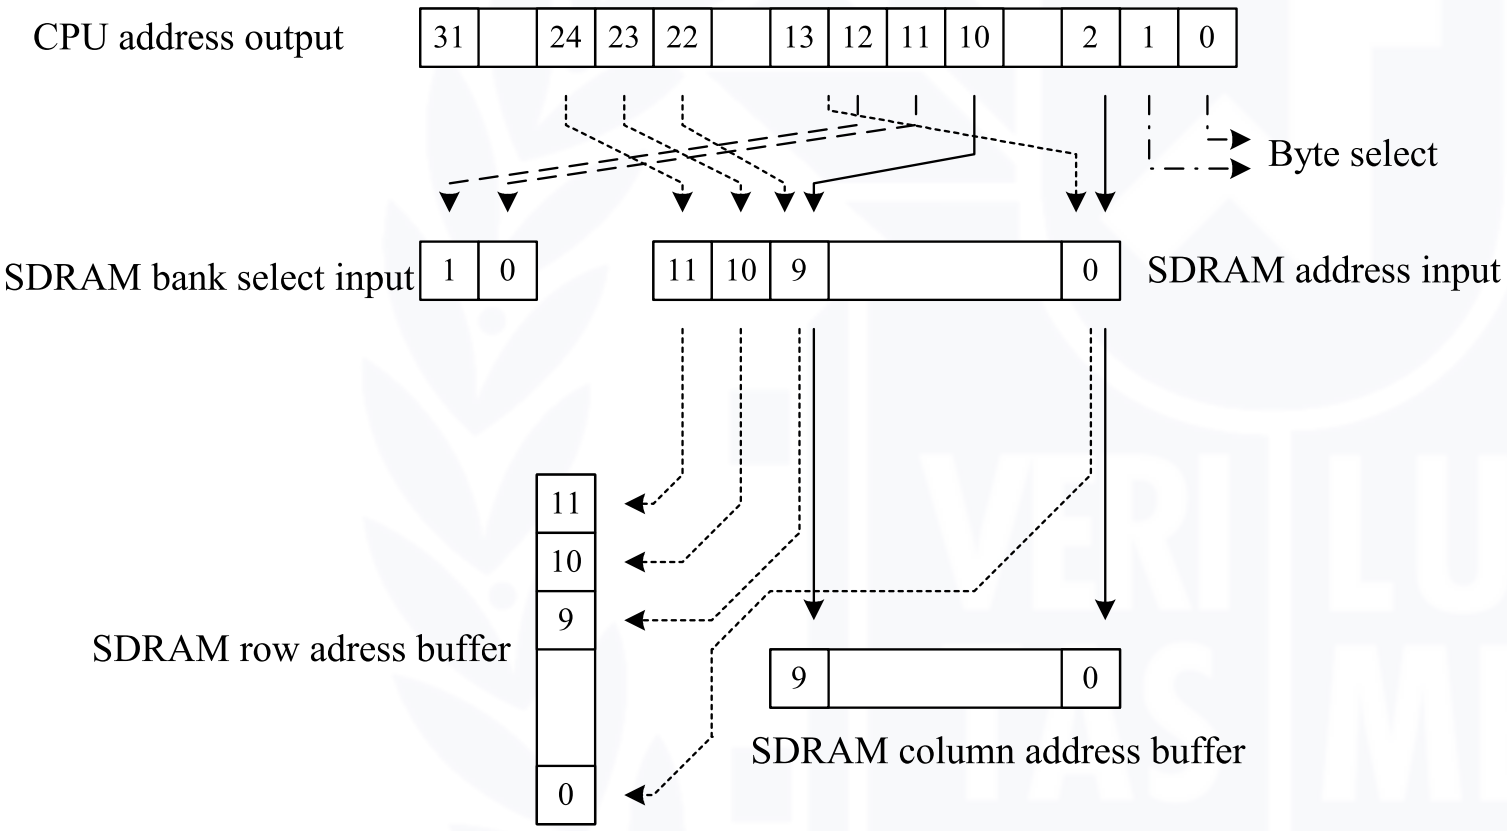

DRAM Mode Control

- High performance configuration **row-hit ratio** versus **execution time**

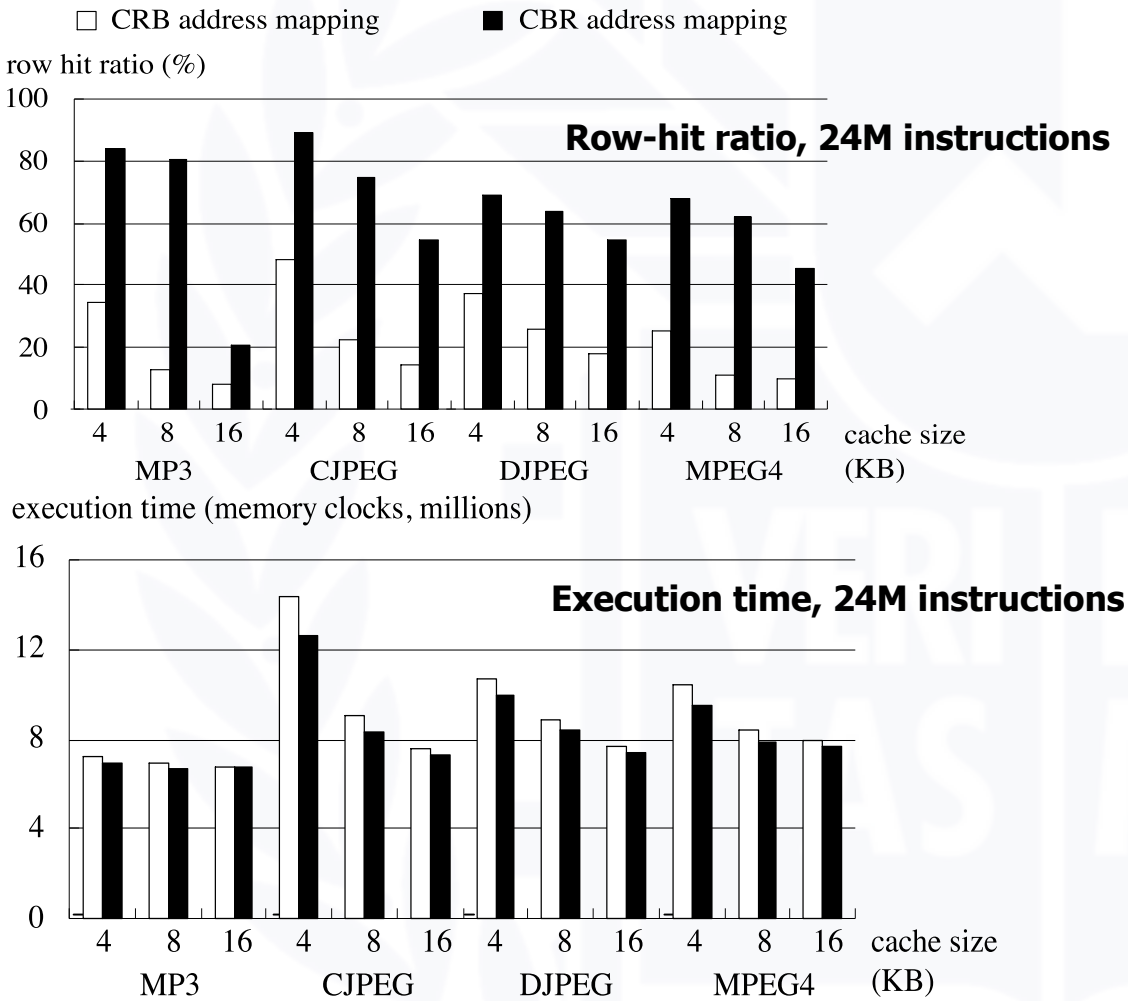

DRAM Mode Control

- High performance configuration **row-hit** ratio versus **energy consumption**

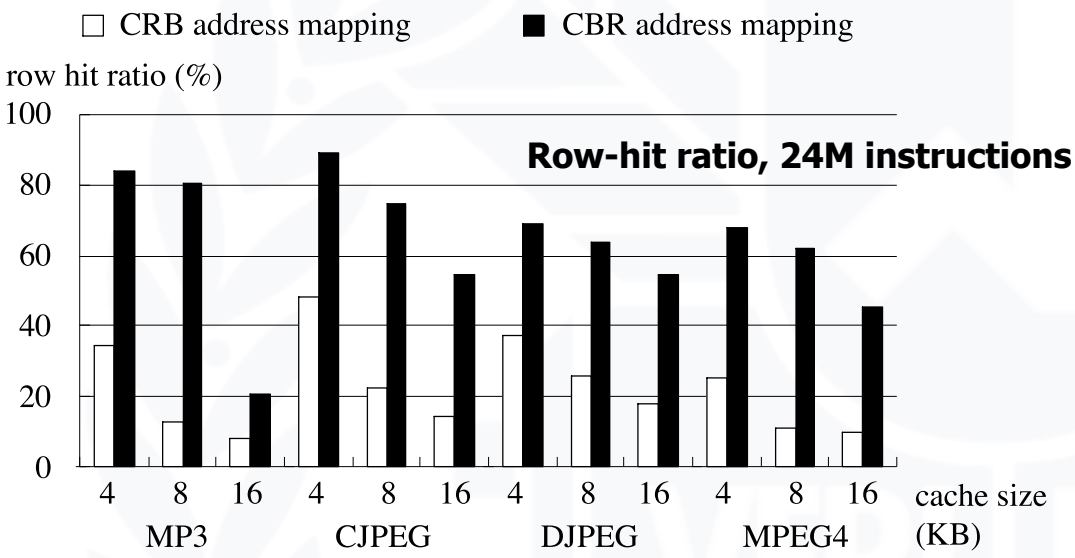

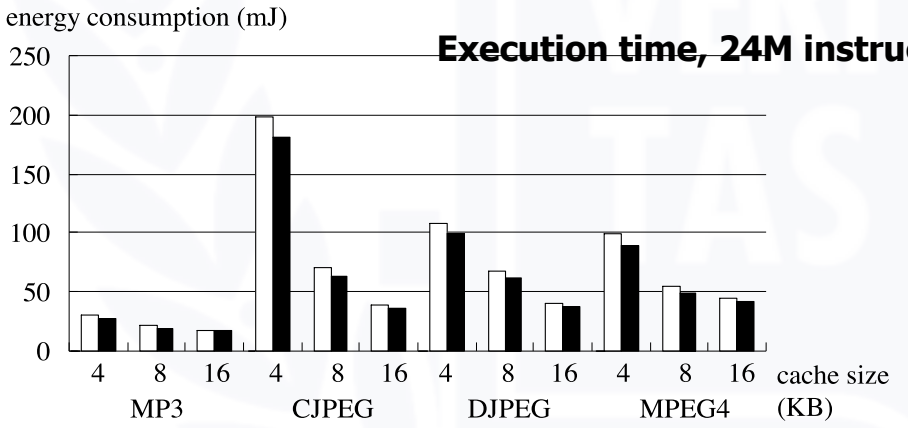

DRAM Mode Control

- SDRAM mode control
  - With elaborated bus encoding scheme, we reduce only around 1% out of total energy of SDRAM devices
    - HDD energy is not observed in SDRAM devices
    - Actual portion of WDD energy is very small
  - SDRAM mode control schemes are introduced
    - Forcing SDRAM devices to active (high energy state) or idle mode (low energy state)
    - Shutting down SDRAM devices
  - The first mode control scheme requires correct estimation of **row hit behavio**r
  - The second mode control scheme requires proper **break-even** time for shutting down the devices

DRAM Mode Control

- High-performance configuration idle clocks between successive memory operations
  - Cache hit ratio determines the vertical scale of the spectrum
  - The processor clock frequency determines the horizontal scale
  - As cache-hit ratio increases, new spectrums appear at the right side of the graph

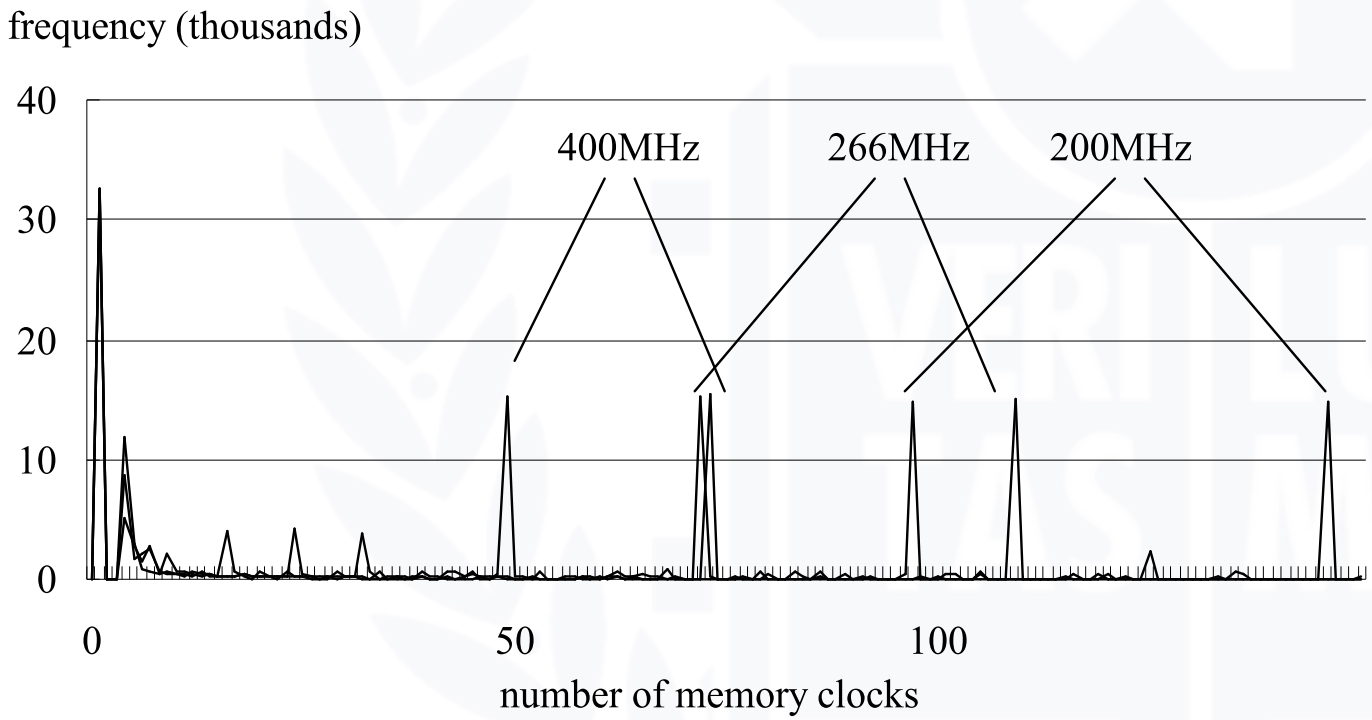

DRAM Mode Control

- High-performance configuration idle clocks between successive memory operations or successive row hits with CRB and CBR

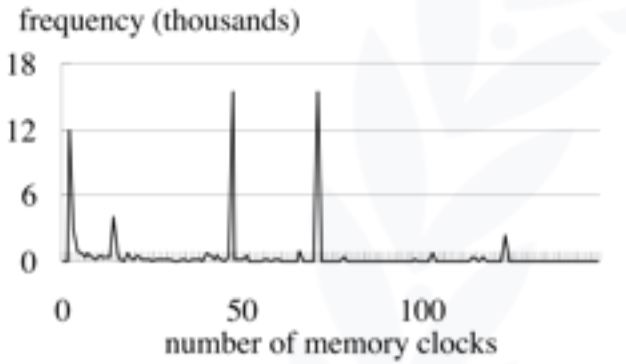

(a) idle clocks between successive memory operations (CRB)

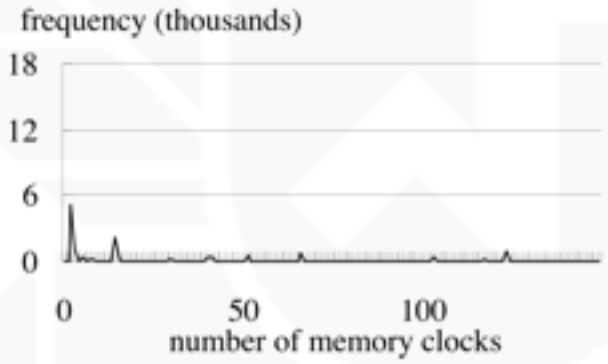

(b) idle clocks between successive row hits (CRB)

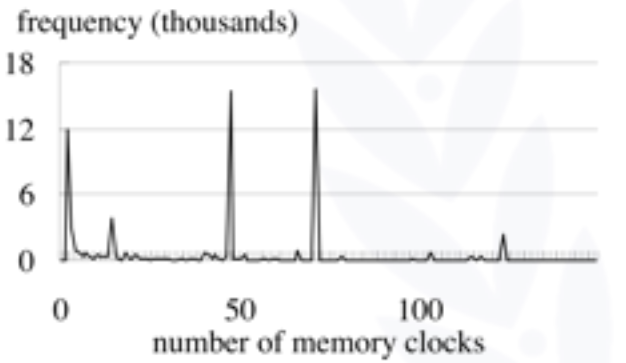

(c) idle clocks between successive memory operations (CBR)

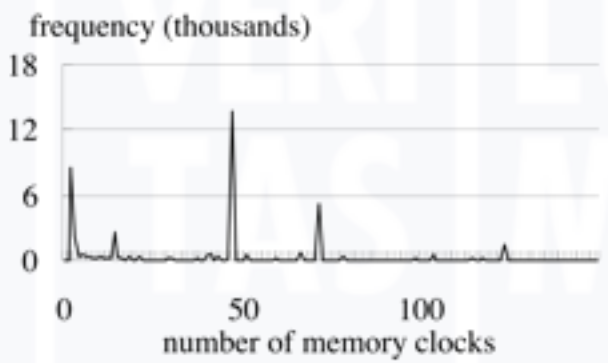

(d) idle clocks between successive row hits (CBR)

DRAM Mode Control

- Commercial SDRAM controllers commonly set up the time-out value for the delayed precharge by 256 clock steps or do not have a capability to set up the value
  - But the dominant row hit spectrum locates at 3 clocks which is a desirable time-out value
  - When we adopt CBR alignment, the magnitude of the dominant row hit spectrum at 3 clocks becomes larger than that of CBR alignment, and a new row hit spectrum appears at 48 clocks
  - But it is not a good idea to set the time-out value to 48 clocks, because it results in more energy consumption due to large CL energy in the active mode
  - The optimal time-out values
    - MP3 decoder: 3 clocks
    - JPEG compressor, JPEG decompressor: 4 clocks, 4 clocks
    - MPEG4 decoder: 10 clocks

DRAM Mode Control

- Shutting down the SDRAM is conducted by the idle time distribution
  - It is important to count in both dynamic cost and static cost in calculating mode control energy overhead
  - For I → IPD → I (when using idle-mode power down), energy cost is $(0.0038 + 1.6 \cdot 10^{-4} n)\tau$, where n is the dwell time in state IPD in clock cycles

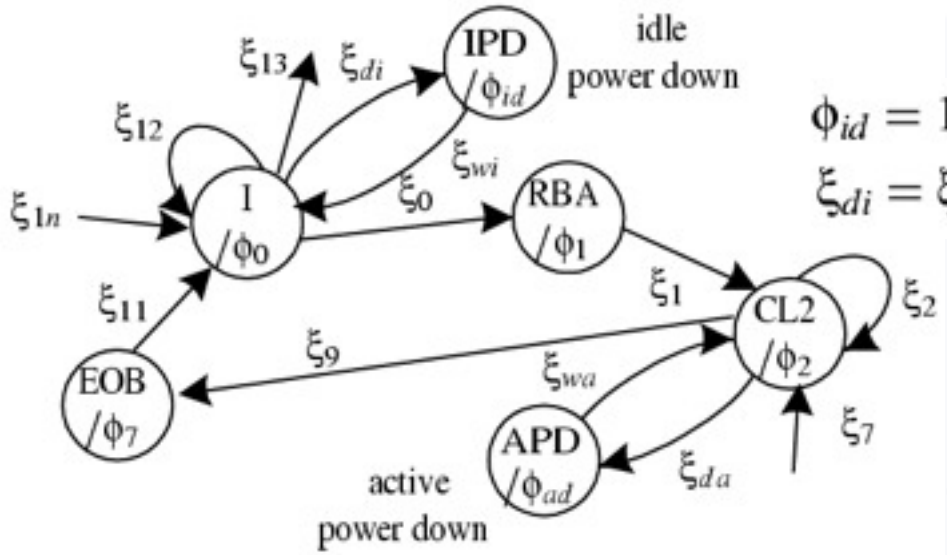

$$\phi_{id} = 1.6 \cdot 10^{-4}\tau,\ \phi_{ad} = 7.5 \cdot 10^{-4}\tau$$

$$\xi_{di} = \xi_{da} = \xi_{wi} = 0,\ \xi_{wa} = 0.54$$

(unit: nJ/bit)

DRAM Mode Control

- Energy reduction of SDRAM devices
  - Mid-performance configuration with auto-precharge policy
    - In most cases, auto-precharge policy achieves slightly better performance enhancement and lower energy consumption than active-page policy with conventional CRB address alignment
    - But, active-page policy with CBR address alignment achieves 2.9% performance enhancement on the average
    - When we use given delayed precharge policy with CBR alignment and idle-mode power down, we achieve 1.3% performance enhancement and 10.9% energy reduction on the average

App.	Reduction technique	CL	CD	WDD	Total	Reduction ratio (%)	Execution time (ms)
CJPEG	active page (CRB)	21.0	56.5	2.2	79.7	-11.8	130.5
	active page (CBR)	20.3	50.0	1.9	72.3	-1.4	123.5
	auto precharge	**9.8**	**59.2**	**2.3**	**71.3**	**0.0**	**128.5**
	auto precharge, power down	4.9	59.2	2.3	66.4	6.9	131.0
	delayed precharge (CBR)	12.6	54.6	2.0	69.2	3.0	125.0
	delayed precharge (CBR), power down	8.4	54.6	2.0	65.0	8.9	126.2

Energy consumption (mJ), 24M instructions

DRAM Mode Control

- High-performance configuration with active-page policy
  - Active-page policy with given CBR alignment not only consumes less energy but also enhances the performance than active-page policy with conventional CRB alignment
  - When we use given delayed precharge policy with CBR alignment and idle-mode power down, we achieve 2.4% performance enhancement and 25.1% energy reduction on the average
  - In addition, active-page policy with conventional CRB alignment does not increase the performance remarkably comparing with the other reduction techniques

App.	Reduction technique	CL	CD	WDD	Total	Reduction ratio (%)	Execution time (ms)
CJPEG	**active page (CRB)**	**11.7**	**25.8**	**1.0**	**38.6**	**0.0**	**76.3**
	active page (CBR)	11.5	23.4	0.9	35.8	7.3	73.6
	auto precharge	5.3	26.7	1.1	33.1	14.2	75.1
	auto precharge, power down	2.1	26.7	1.1	29.8	22.7	75.9
	delayed precharge (CBR)	6.0	26.0	1.0	33.1	14.3	74.1
	delayed precharge (CBR), power down	3.0	26.0	1.0	30.0	22.2	74.7

Energy consumption (mJ), 24M instructions

DRAM Mode Control

- Low-power SDRAM main memory system
  - The energy reduction schemes for memory buses and devices are orthogonal (independent)
  - Mid-performance configuration

Application	Original	Final	Reduction (%)
MP3	26.8	16.0	40.2
CJPEG	83.0	71.0	14.5
DJPEG	80.1	68.8	14.1
MPEG4	62.7	53.3	15.0

Energy consumption (mJ), 24M instructions

  - High-performance configuration

Application	Original	Final	Reduction (%)
MP3	21.3	11.6	45.5
CJPEG	44.9	32.8	26.8
DJPEG	47.2	34.6	26.6
MPEG4	50.4	39.0	22.7

Energy consumption (mJ), 24M instructions