

Power Management for ASICs

April 24, 2015



Prasad Subramaniam
eSilicon Corporation

Power Trends & Challenges


- System power dissipation is becoming more critical
 - On-going integration increasing overall system power
 - “Green” systems and “Green” companies imperative
 - Power management no longer limited to mobile applications
- Underlying technology adding to the challenge
 - Device leakage power increasing by process node (mitigated somewhat by FinFET and FDSOI)
 - Voltage no longer scales with process node
 - Package thermal transfer not improving
 - Tools focus on execution

Intelligent Power Management Required!

Power Management System

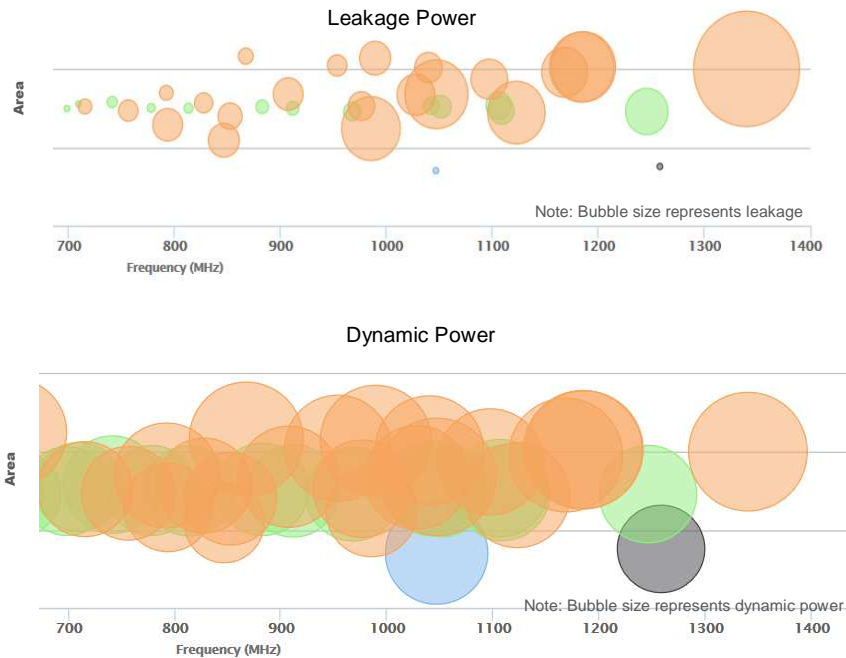


- Complete power management includes:
 - Using the right process, libraries and IP
 - Leveraging a power-aware design methodology
 - Minimizing overhead while ensuring power integrity



Power Management System

Process Technology Impact on Power



- Wide variation in leakage power seen in same technology depending on target frequency
- Faster doesn't always imply higher leakage
- Dynamic power (per MHz) is similar and primarily depends on power supply voltage

Power and Performance Variations for a JPEG Encoder Core

Library Selection is Complex



- Multiple library variables impact design PPA
 - Number of tracks
 - Channel length
 - VT
 - VDD
 - Availability of power-optimized cells like multi-bit flip-flops


Standard Cell Libraries for Low Power Design



	Features	Benefits
Low Power Baseline	Characterized over wider VDD range	Multi-VDD operation
	Multi-Vt components	Leakage minimization
	Multi-bit flipflops	Dynamic power reduction
	Decoupling capacitors	Power integrity
Multi-VDD	Level shifters	Level translation between multi-VDD islands
	Isolation cells	Control the logic inputs for interface between OFF and ON domains
	Retention flops	Saves states in shut-down mode
	Power gating switches	Enables island power-down
Long Channel Devices	Long channel standard cells	Lowers leakage on non-critical paths
Gate Length Bias	Automated biasing	Lowers leakage on non-critical paths

Impact of VDD on Power

- Using higher VDD with a high V_T and/or long channel library helps reduce leakage for the same performance



Impact of Multi-Vt Cells For Power

Cortex A7 core system in 28nm

- ULVT cells for 1GHz operation at SS, 0.81V, -40C
- Multi-Vt cells for leakage power recovery

	Leakage Power (mW)	Active Power (mW)	Total Power (mW)
ULVT	430	404	834
Mixed VT	191	374	565
Reduction	55%	7%	32%

SRAM Power Management in Memory Compilers



- Low dynamic power
 - Multi-Vt peripheral logic
 - Dual rail operation
 - Ultra low voltage operation using custom logic rule bit cell
 - Memory segmentation
- Low leakage power
 - LL bit cell
 - High Vt peripheral logic
 - Multiple sleep modes
 - Light sleep – 40% leakage power reduction with memory data retention and fast wake-up, uses array biasing; one cycle recovery
 - Deep sleep – 70% leakage power reduction with memory data retention and integrated power switches, uses periphery shutdown, ten cycle recovery
 - Shutdown – 90% leakage power reduction with integrated power switches, fifty cycle recovery time, no data retention

Single Port Optimized Low Voltage/Low Power SRAM

40LP Logic Rule Approach



Memory architecture	Word	Bit	Mux	Area (um ²) pre-shrink	Read Power	Write Power	Leakage Power	Leakage Power – shutdown mode
					(uW/MHz)	(uW/MHz)	(mW)	(mW)
					FF / 1.21V / 125C			
Reference 40LP SP HD design using 6T Bit Cell	2048	64	4	50,039	19.38	23.04	1.07	
	8192	64	16	189,764	43.78	43.03	2.52	



Standard HD architecture (top) vs LP/LV architecture (bottom)

Memory architecture	Word	Bit	Mux	Area (um ²) pre-shrink	Read Power	Write Power	Leakage Power	Leakage Power – shutdown mode
					(uW/MHz)	(uW/MHz)	(mW)	(mW)
					FF / 1.21V / 125C			
eSilicon 40LP Custom Low Power / Low Voltage SRAM using Logic Rule Bit Cell	2048	64	2	113,000	14.75	10.99	0.435	0.017
	8192	64	4	419,000	41.06	29.04	1.595	0.033
					FF / 0.85V / 125C			
	2048	64	2	113,000	6.76	5.28	0.185	0.007
	8192	64	4	419,000	19.62	14.11	0.684	0.017

ASIC Example – Network Processor

Reducing Power at the Same Performance



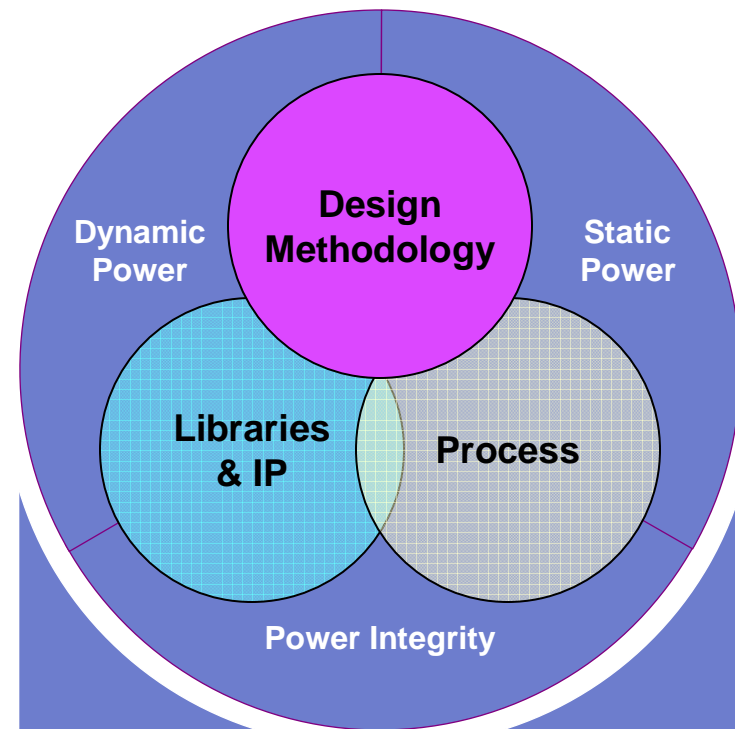

- Technology: 28nm
- 394Mb memory subsystem
- Customization
 - Standard Vt memory array operated at nominal VDD
 - Migrated memory peripheral logic to high Vt
 - Re-characterized for overdrive operating voltage

- Result = Same Performance
- Static Power Savings = 20%

Architecture	Array Leakage (mW)	Periphery Leakage (mW)	Total Leakage (mW)	Array Leakage (mW) Overdrive	Periphery Leakage (mW) Overdrive	Total Leakage (mW) Overdrive
DP SRAM	231	3726	3957	304	2844	3075
2P RF	2653	12117	14770	2769	9250	11903
SP SRAM	262	1966	2227	313	1500	1762

Nominal VDD with SVt = 21 Watts
Overdrive VDD with HVt = 16.7 Watts

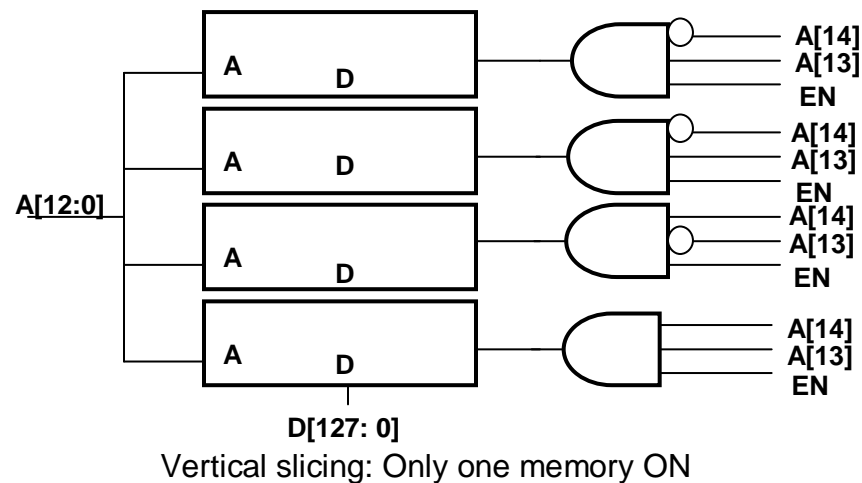

Power-Aware Design Methodology



Power Management System

Architectural Optimization

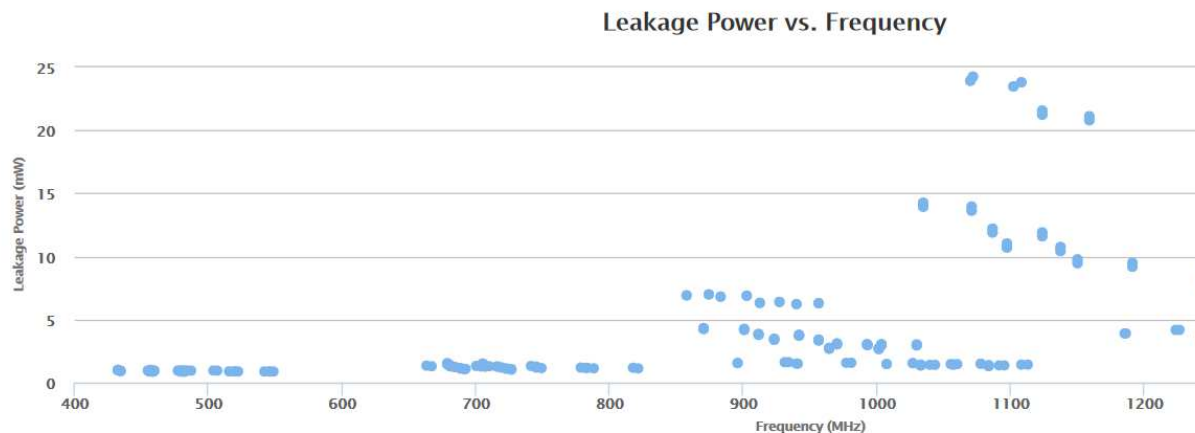
- **Example:** Memory organization



Total system power management begins at the architectural level

Memory Selection and Optimization

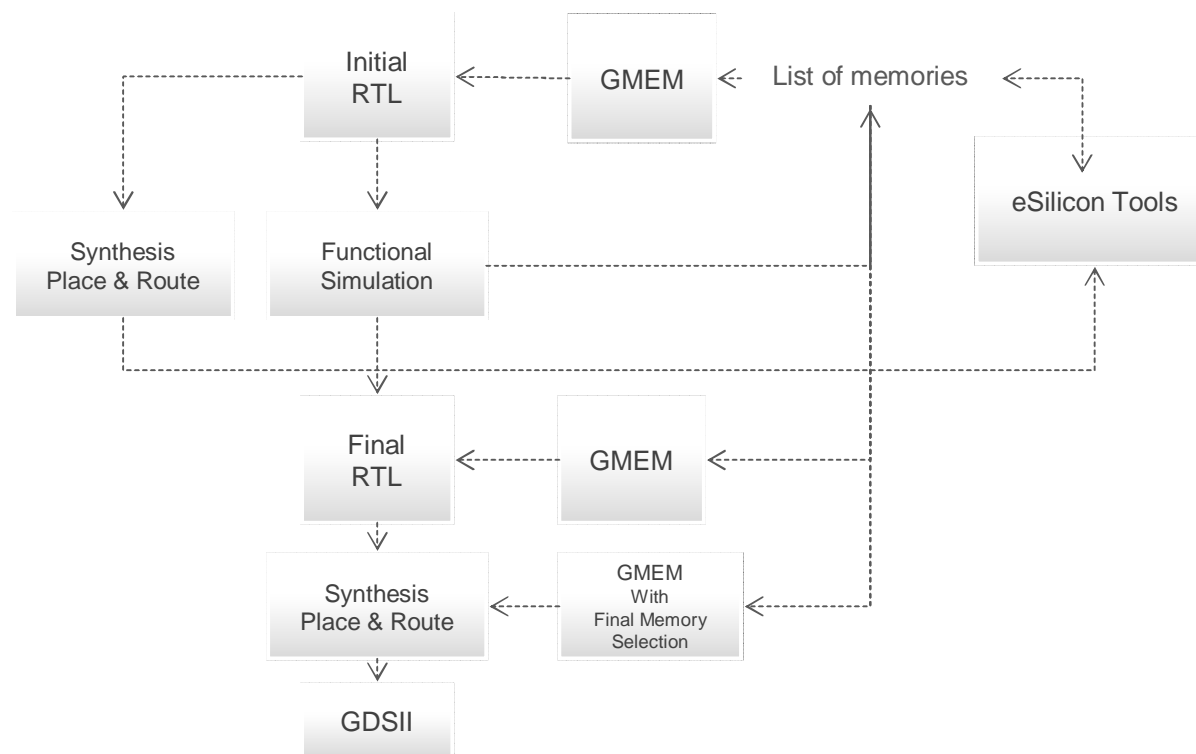
- Typical chip statistic: 80% of the memory area and power are contributed by 20% of the memories
- Memory selection is critical for area and power optimization
- EDA tools focus on logic optimization but leave memory optimization to user
- Traditional memory selection is manual and often ad-hoc
- Memories are changed primarily for functional reasons



1kx32 Single port memory choices

Memory Selection & Optimization Using Generic Memory Models



- eSilicon generic memory model (GMEM) provides automated memory optimization
 - User RTL is based on parametrized generic memory models
 - eSilicon tools select memory based on synthesis results and memory compiler constraints



Signoff for Power

- Standard signoff (28nm)
 - Voltage regulator tolerance: 5%
 - IR drop: 5%
 - Timing: SS, VDD – 10%, -40C
 - Power: FF, VDD + 5%, 125C
- Aggressive signoff (28nm)
 - Voltage regulator tolerance: 3%
 - IR drop: 3%
 - Timing: SS, VDD – 6%, 0C
 - Power: FFG, VDD + 3%, 105C
 - Lower and re-center VDD


VDD = 0.9V
 IR drop: 45mV
 Voltage regulator tolerance: 45mV
 Timing: SS, 0.81V, - 40C
 Power: FF, 0.945V, 125C



VDD = 0.864V
 IR drop: 27mV
 Voltage regulator tolerance: 27mV
 Timing: SS, 0.81V, 0C
 Power: FFG, 0.891V, 105C

Benefit

- Easier to close timing
 - Process is faster at 0C compared to -40C by about 8%
 - Can use more transistors with higher V_t
 - Can use additional margin
- Lower Power
 - FFG is more realistic as local variations in wafer average out
 - Leakage is significantly reduced by lowering V_{DDmax} and temperature
 - Leakage reduction: 63%
 - Active power reduction 13%

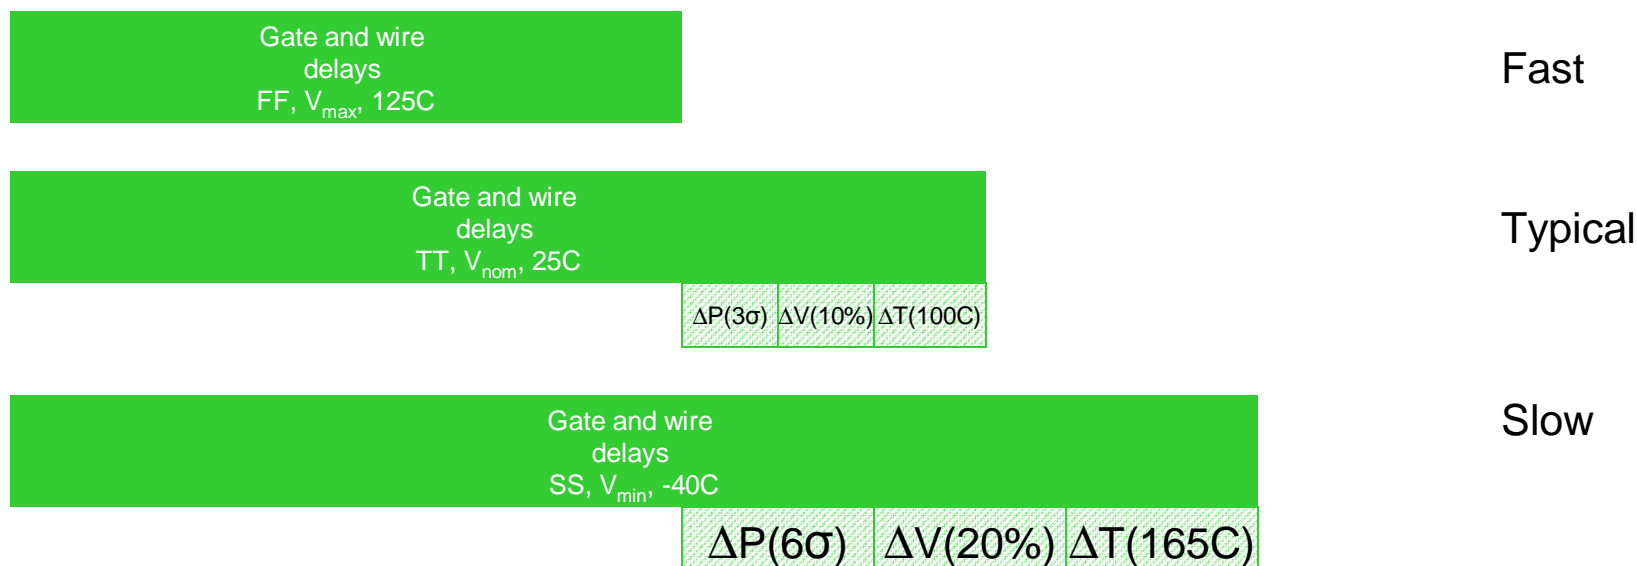


Voltage Scaling and Binning



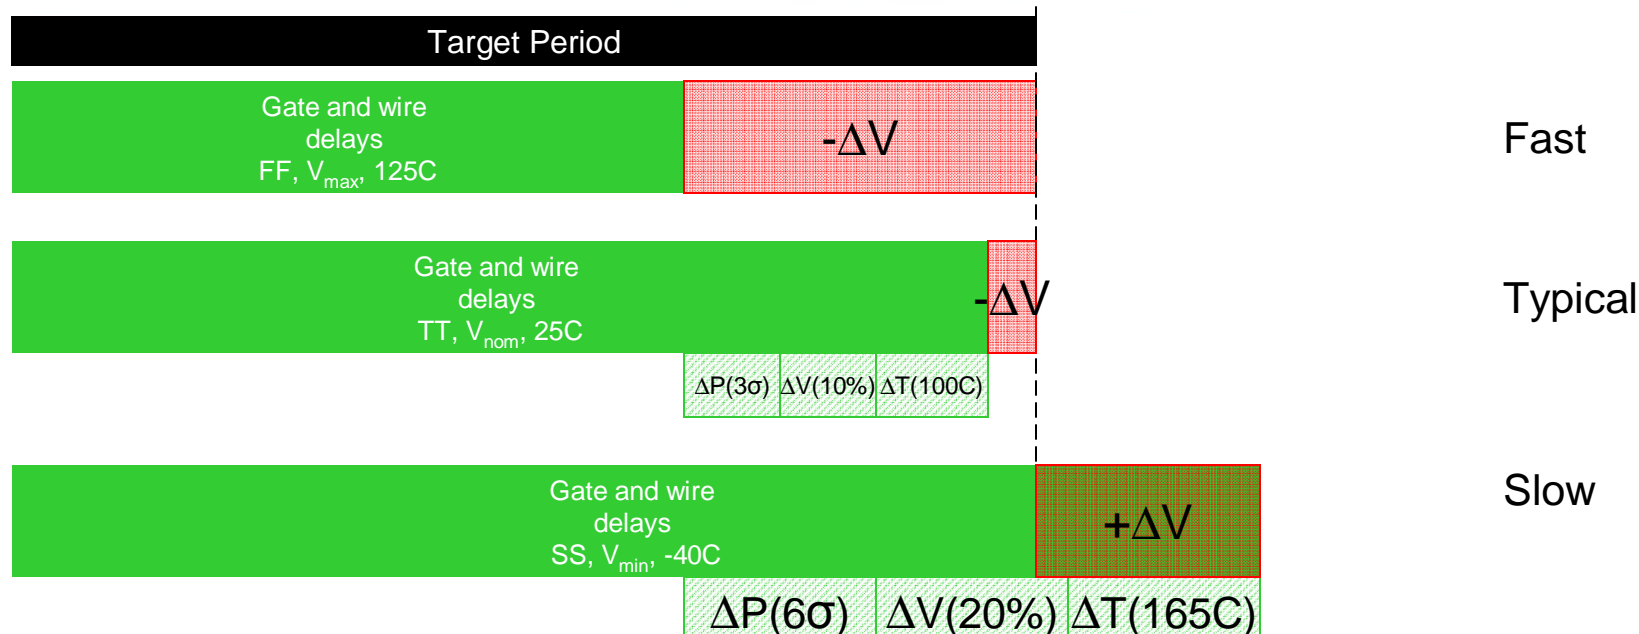
- Voltage scaling is the most effective method of reducing power for FF parts
 - Scaling can be continuous or discrete
 - Discrete scaling is equivalent to binning
- Binning allows parts to be separated based on their process corner
 - Different voltages for each bin ensure that performance is met while power is optimized

Voltage Scaling for Performance and Power



- The FF, V_{max} , 125C (Fast) part gives the maximum performance
- The TT, V_{nom} , 25C (Typical) part loses performance due to all three components P, V and T
- The worst performance is from the SS, V_{min} , -40C (Slow) part

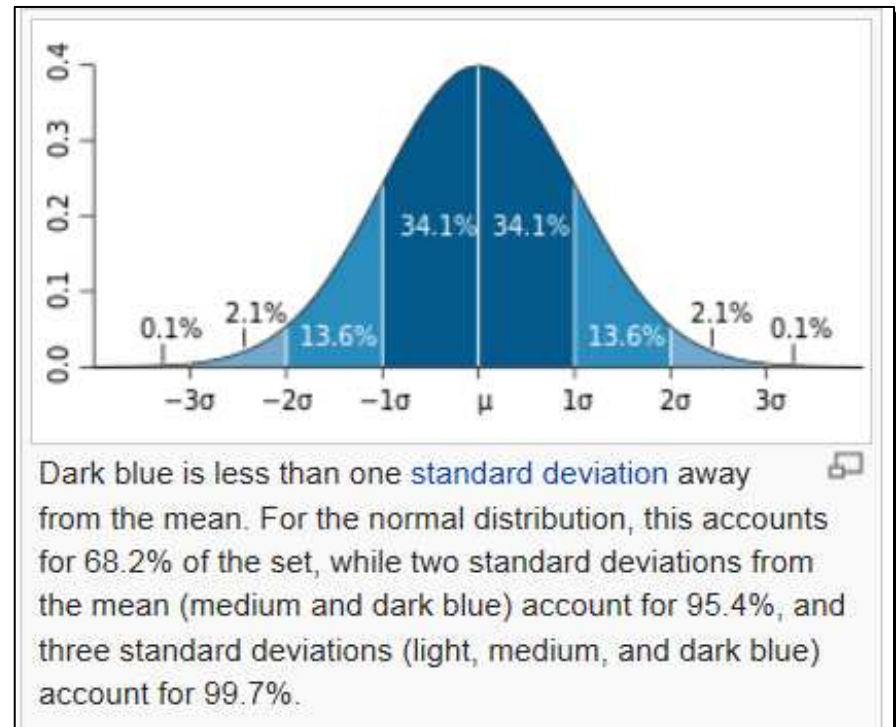
Voltage Scaling for Performance and Power



- Performance can be recovered by increasing voltage on a part
- If target frequency is lower than that of the Fast part, power can be recovered by lowering the voltage of the Fast part
- Optimum performance and power can be achieved by centering the part around the target frequency

Process and Yield Management

- Foundry process is well-controlled
- 3σ signoff is extremely conservative
- Process and yield management can yield power and performance improvements
 - Shift process by 1σ
 - Discard parts outside 2σ (4.6% yield loss)



Power Management Case Study

- Customer requires maximum total power dissipation to not exceed 40W
- Tapeout ready database power at FFG, 1.05V, 105C is **66.58W**
- Assumptions:
 - Typical voltage at 1V
 - Voltage regulator tolerance is $\pm 50\text{mV}$
 - IR Drop is 50mV
 - Operating frequency of 600MHz
- How do we meet customer's power target?

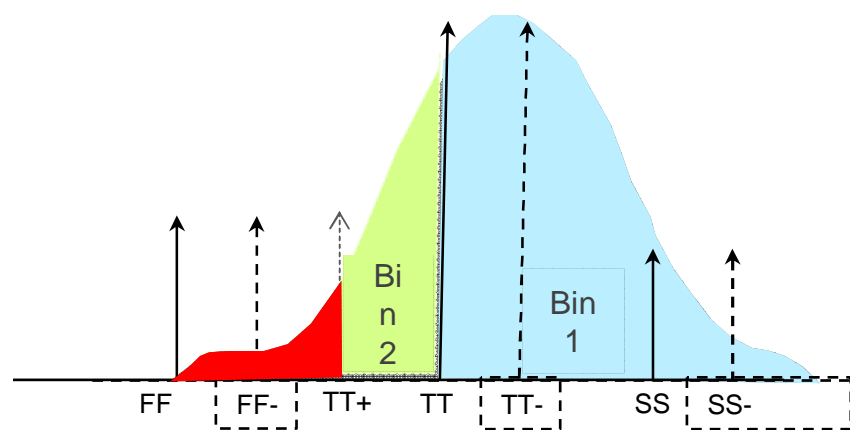
	Total (W)	Dynamic (W)	Leakage (W)
Logic	35.02	25.74	9.28
Memory	20.78	7.76	13.02
CAM	3.24	1.17	2.06
SerDes (AVDD)	7.48	6.19	1.29
IO ring + rest	0.08	0.07	0.01
Total	66.58	40.93	25.65

Power Management Approach

- Any strategy for power reduction to achieve 40W requires
 - Use of two power supplies
 - Power supply for SerDes (1V typical)
 - Power supply for Core
 - Lower frequency operation
- Core supply voltage tolerance should be as low as possible
 - Customer is unwilling to reduce voltage tolerance below 50mV
- Minimum core voltage is 0.81V (memory VDDmin)
- Use IR drop based on actual data from power analysis
 - 30mV simulated at 125C with >100W chip power (FF, 1.05V, 125C)
 - 15mV assumed for 105C with < 50W chip power (linear scaling)

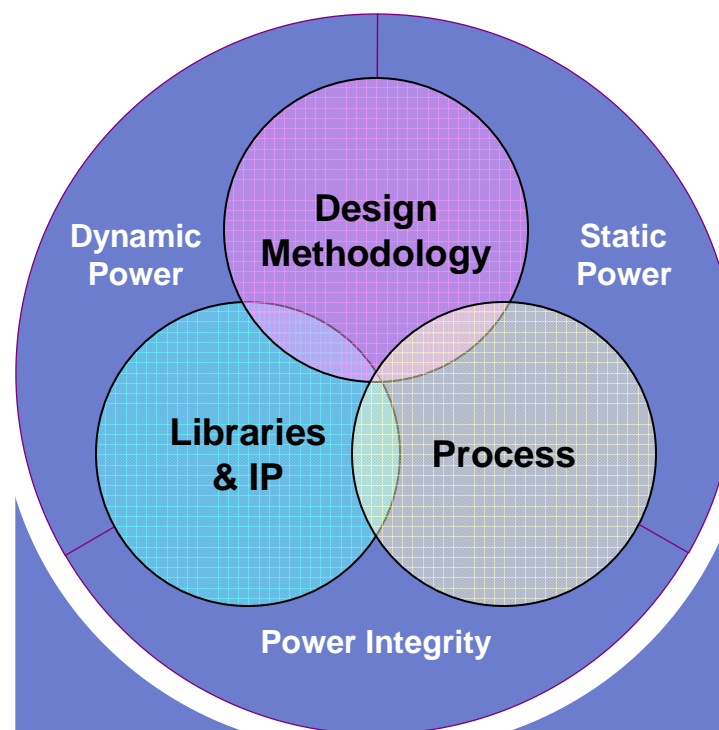
Power Management Solution

- Operating Frequency is 450MHz
- Separate power supplies for SerDes and Core
- SerDes VDD = 1V \pm 50mV
- Two bins, bins separated at TT
- Process skewed by one sigma
- Bin 1 (TT to SS-)
- Core VDD = 0.93 \pm 50mV
- Power at TT, 0.98V, 105C **35.96W**
- Worst timing at SS-, 0.865V, -40C
- Bin 2 (TT+ to TT) No yield loss
- Core VDD = 0.875 \pm 50mV
- Power at FFG-, 0.925V, 105C **40.2W**
- Worst timing at TT, 0.81V, -40C
- Bin 2 (TT+ to TT) 2.1%yield loss
- Core VDD = 0.875 \pm 50mV
- Power at TT+, 0.925V, 105C **36.35W**
- Worst timing at TT, 0.81V, -40C




Power Management System Summary

- A complete system approach from concept to production
 - Lower dynamic and static power
 - Power integrity ensured
- Leveraging the best process, IP, libraries, power aware tools and power management methodology
- Power management solutions complement traditional EDA solutions



Power Management System

Intelligent Power Management Delivered



eSilicon®

Enabling Your Silicon Success™