



SHAPING THE NEXT GENERATION OF ELECTRONICS

**JUNE 23-27, 2024**

MOSCONE WEST CENTER  
SAN FRANCISCO, CA, USA

# GPU Power Tracking and Optimization Using Emulation

Sandesh Saokar, Intel CGAI

Veera Dasari, Intel CGAI

Navid Farazmand, Intel CGAI



# Motivation

- Ability to track dynamic power using real-world use-cases/workloads at GPU IP level rather than bottoms up projections based on unit level tests
- Enable exploration of biases in real-world data-patterns and design scenarios
- Avoid late surprises in IP power by left-shifting the pre-silicon power tracking process
- Establish a robust data collection and visualization/analysis mechanism to monitor the trends
- Need a way to differentiate power savings/degradation from back-end physical design recipe changes, software changes versus RTL changes throughout the course of project execution

# Main Idea

## Real-world Scenarios

Use real-world workloads on emulation (EMU) to estimate and optimize dynamic power

- This can be accomplished using EMU workloads and available industry leading fast-synthesis/RTL based power estimation solutions

## Implementation agnostic RTL analysis

RTL based estimation allows physical design tool-methodology independent approach

- The relative power trend will be predominantly driven by RTL changes. No impact of physical design recipes/goodness
- Establish reasonable correlation between RTL-based versus sign-off netlist-based power estimates
- We can compare relative power delta with delta from sign-off tools to quantify impact of Physical design optimizations

## Fast Turn-around-time (TAT)

Enable GPU IP level analysis in reasonable timeframe

- RTL power relative trend tracking TAT at GPU IP level is < 48hrs for enabled emulation workloads (multi-million cycles)
- TAT improvement mainly coming from removing physical design dependency and automated regression flow
- All regression data stored to a database, using PowerBI and custom web apps for data visualization and analytics

# RTL-based Average Power Estimation Flow

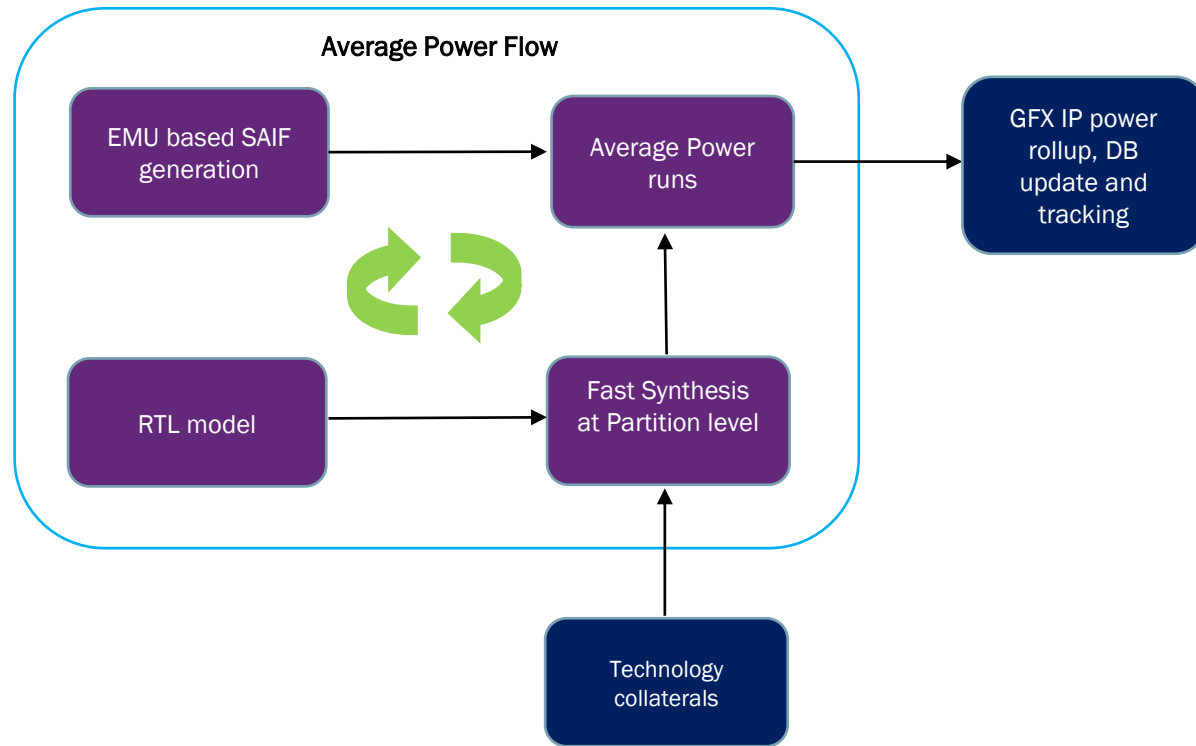


Fig 1a: Figure depicting RTL-based power estimation methodology

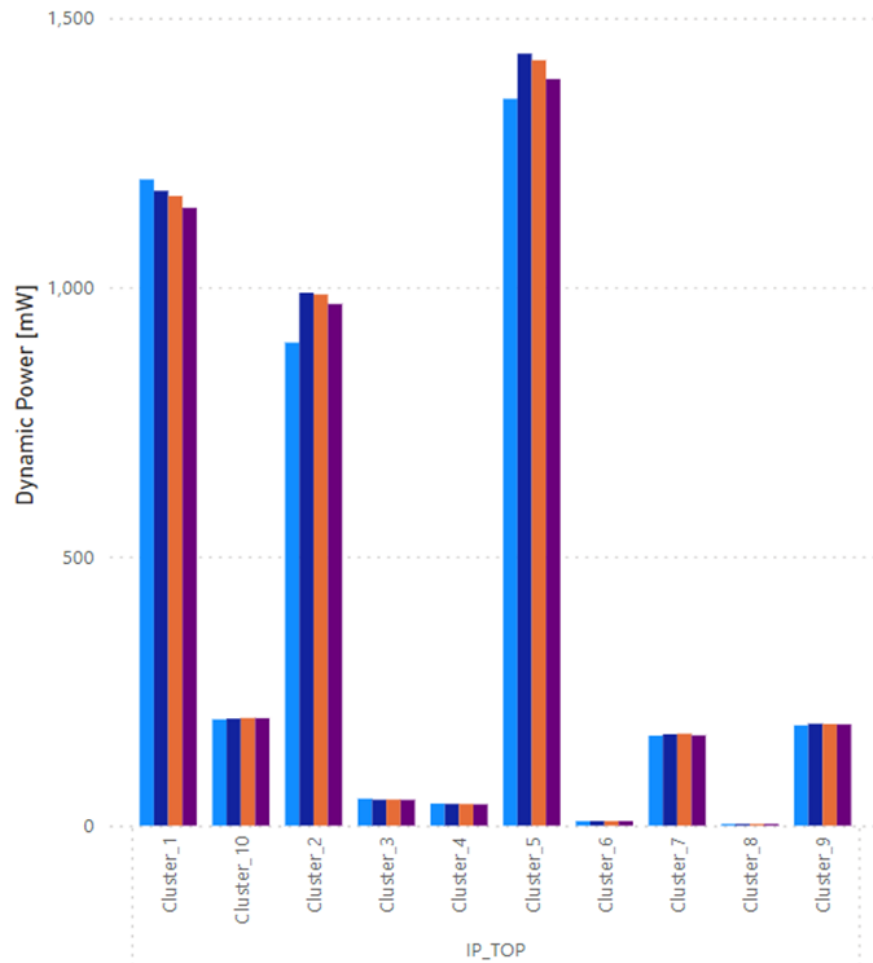
Sum of Dynamic Power [mW]		Column Labels		
Row Labels	2023WW20	2023WW22	Abs Delta	% Delta
IP_TOP	5383.85	5560.37	176.52	3.28%
Functional	4902.05	5071.22	169.17	3.45%
Cluster_1	1200.22	1179.08	-21.14	-1.76%
Cluster_10	197.42	198.76	1.34	0.68%
Cluster_11	101.3	108.13	6.83	6.74%
Cluster_12	67.1	62.86	-4.24	-6.32%
Cluster_13	54.78	55.21	0.43	0.78%
Cluster_14	96.1	95.8	-0.3	-0.31%
Cluster_15	84.19	81.31	-2.88	-3.42%
Cluster_16	400.96	410.33	9.37	2.34%
Cluster_2	897.14	989.81	92.67	10.33%
Inst_1	500.01	590.86	90.85	18.17%
Inst_2	346.49	347.91	1.42	0.41%
top	50.64	51.04	0.4	0.79%
Cluster_3	50.07	48.09	-1.98	-3.95%
Cluster_4	40.97	40.33	-0.64	-1.56%
Cluster_5	1350.01	1433.87	83.86	6.21%
Cluster_6	8.3	8.48	0.18	2.17%
Cluster_7	167.1	169.76	2.66	1.59%
Cluster_8	0.1	0.1	0	0.00%
Cluster_9	186.29	189.3	3.01	1.62%
Infra	481.8	489.15	7.35	1.53%
DFX	23.6	22.9	-0.7	-2.97%
Globals	120.8	119.23	-1.57	-1.30%
Misc	36.7	38.79	2.09	5.69%
Repeater	300.7	308.23	7.53	2.50%

Fig 1b. Typical results from RTL-based Average Power Estimation Flow

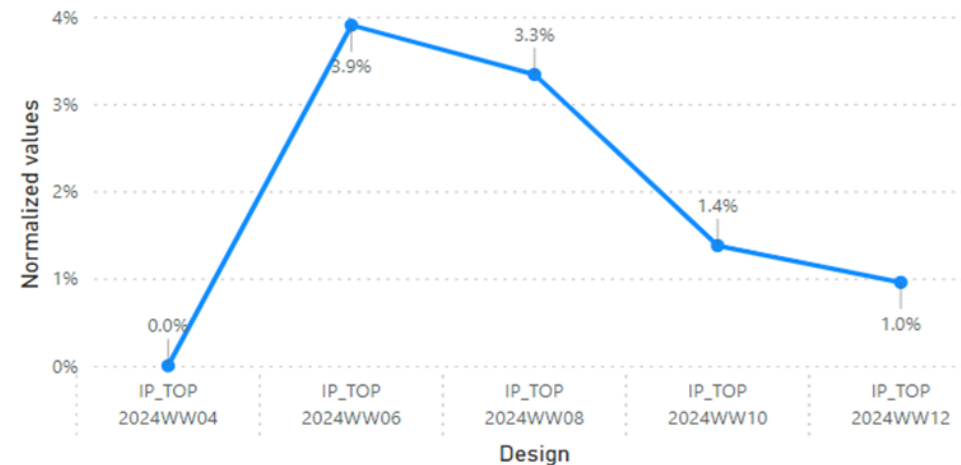
# Regression data visualization and analysis

Dynamic Power [mW] trend by hierarchy and Workweek

Workweek ● 2024WW04 ● 2024WW06 ● 2024WW08 ● 2024WW10



Trend - Normalized Dynamic Power Delta



Category-wise breakdown of Dynamic Power

● Memory power ● Sequential Power ● Clock Power ● Combo Power

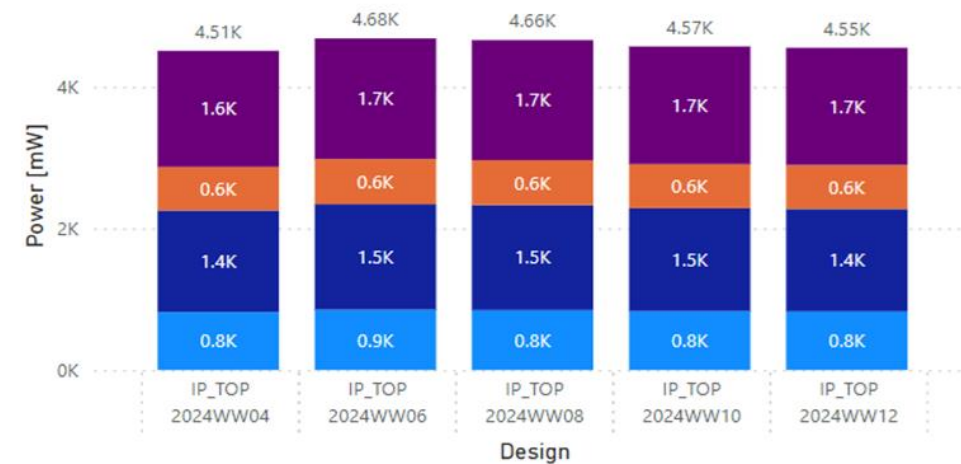
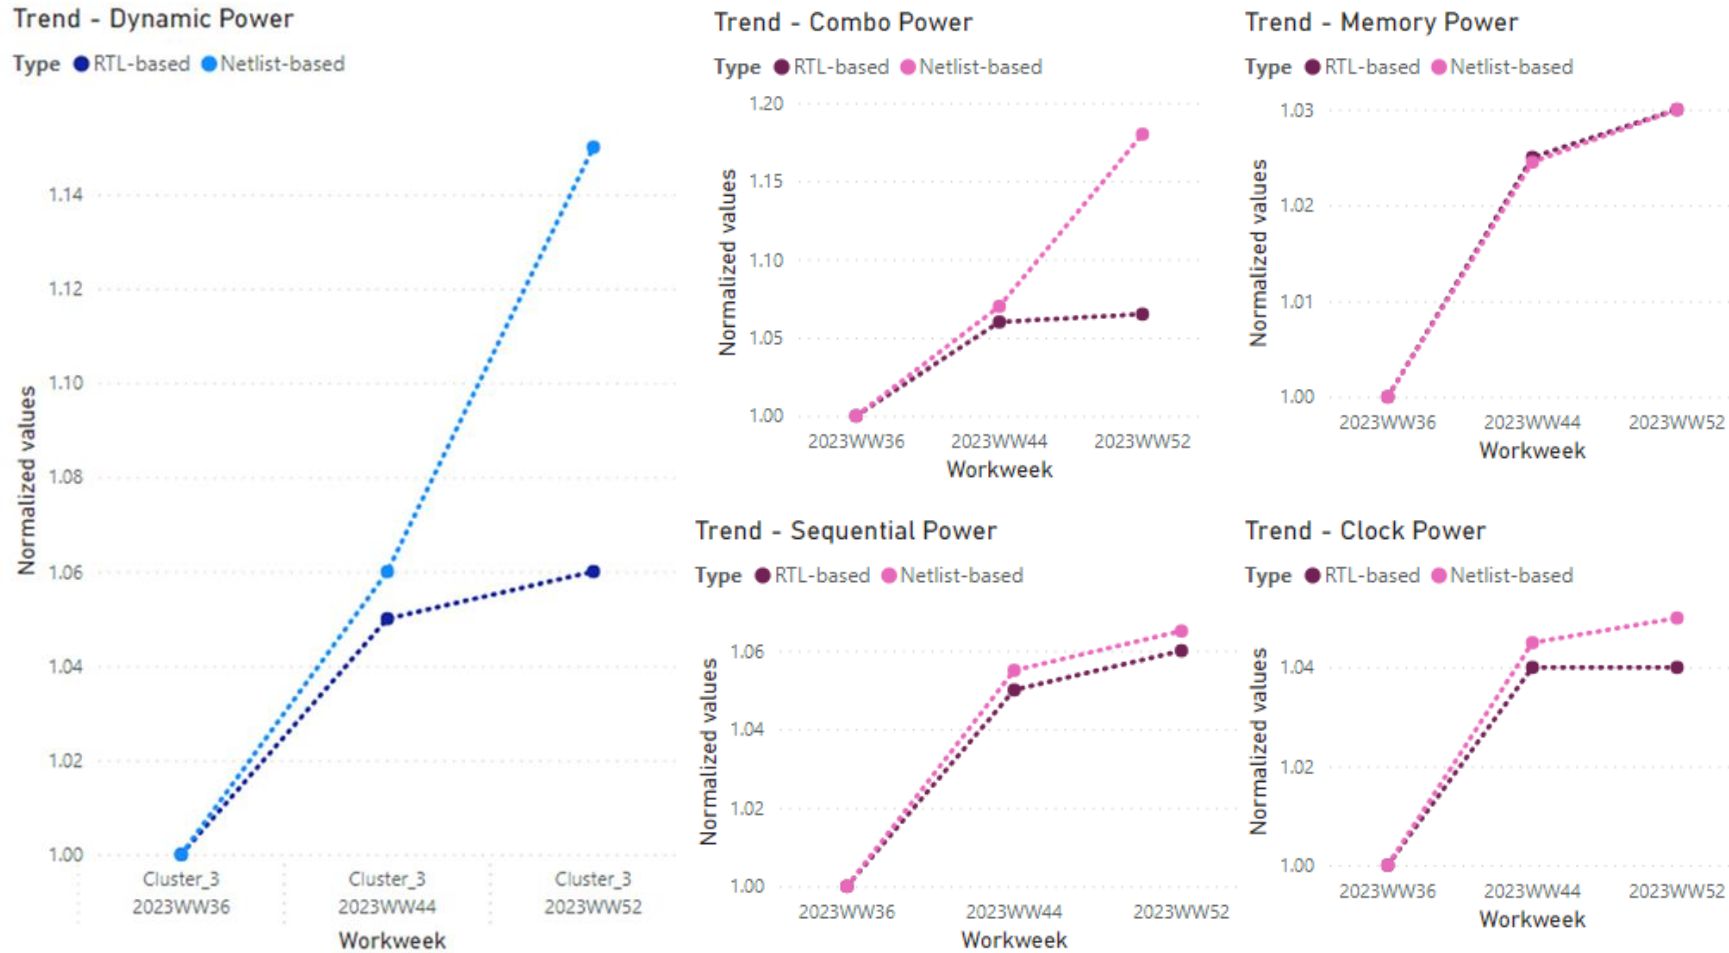


Fig 2: Charts depicting birds eye view of power trend. Uses synthetic power data for demonstration purposes only.

# Tracking RTL, Physical Design, SW Power impact



- Established correlation helps understand divergences in RTL versus NL based trends
- Divergences can often point to the physical design convergence impact on power
- Symptom is manifested as significant divergence in combo power trends due to physical design push for meeting stringent timing requirements.
- Floorplan changes or routing challenges comes at the cost of sizing up cells, detour paths etc. which directly impact combo-power

Fig 4a: Charts depicting impact of physical design convergence on power

# Tracking RTL, Physical Design, SW Power impact

## Trend - Normalized Dynamic Power

Software Driver Version ● ver\_1.1 ● ver\_1.2



- Especially helpful on products where software (driver/compiler) changes can happen during early RTL design phase
- For example: Product with minor incremental design updates based on an older architecture

Fig 4b: Same RTL but workloads are run with a different software driver version to access the impact

# RTL based power optimization flow

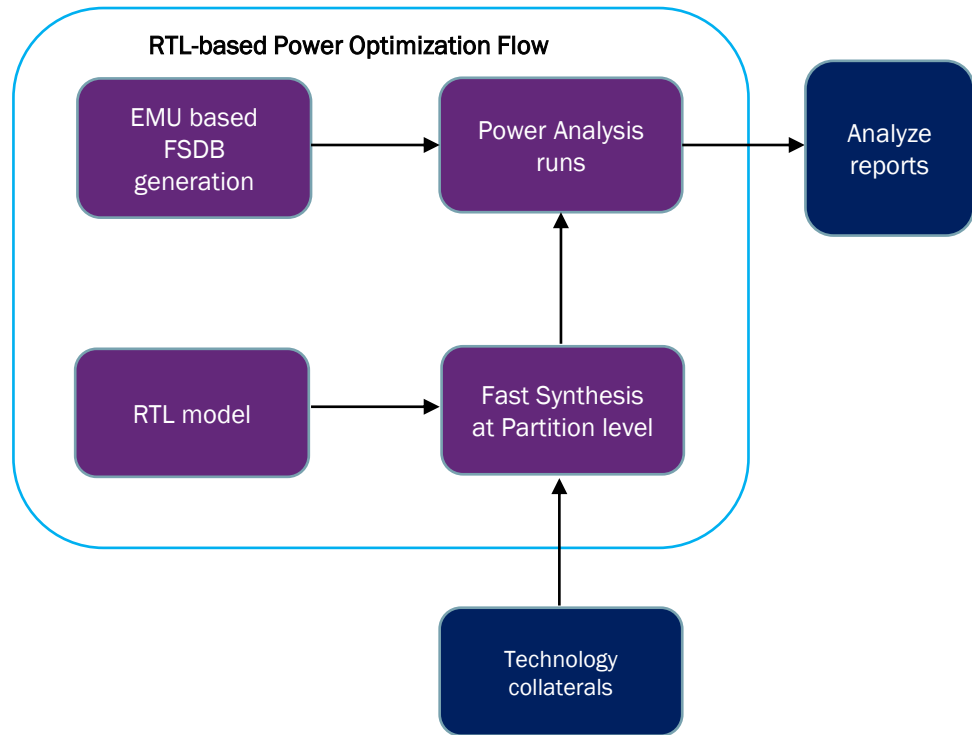


Fig 5a: Figure depicting RTL-based power optimization methodology

Results from RTL-based Power Optimization Flow

Cell Name	Instance	Total Clock Cycles	No-Access Cycles	Read Cycles	Write Cycles	Redundant Read	Redundant Writes
memdphc66x384	ip_top.inst1.mem0	1760	0	1352	408	38	4
memdphc66x384	ip_top.inst1.mem1	2032	0	1432	600	40	3
memdphc66x384	ip_top.inst1.mem2	1654	0	1352	302	17	2
memdphc66x384	ip_top.inst1.mem3	2712	0	1906	806	42	9
memdphc66x384	ip_top.inst1.mem4	1442	0	1360	82	2	2
memdphc66x384	ip_top.inst1.mem5	1541	0	1352	189	16	5
memdphc66x384	ip_top.inst1.mem6	1701	0	1312	389	41	6
memdphc66x384	ip_top.inst1.mem7	2098	0	1312	786	36	4
memdphc66x384	ip_top.inst1.mem8	2889	0	1998	891	92	8

Fig 5b: Snippet from redundant memory access reports

Gating Technique	Width	Static Power (nW)	Dynamic Power (nW)	WCPP (nW)	Initial DACGE	Flop Name
Instantiated	32	116	7710	7710	0	fub_inst.sarb_bindresrc_bufsize[31:0]
Instantiated	25	64.2	6220	6010	0	fub_inst.sarb_bindsamp_baddr[56:32]
Instantiated	25	96.8	5890	5690	0	fub_inst.sarb_bindresrc_baddr[56:32]
Instantiated	25	180	5700	5450	0	fub_inst.sarb_dynstate_baddr[56:32]
Instantiated	25	168	5060	4820	0	fub_inst.sarb_surfstate_baddr[56:32]
Instantiated	20	161	4390	4390	0	fub_inst.sarb_dynstate_bufsize[31:12]
Instantiated	20	60.6	5320	3300	0	fub_inst.sarb_dynstate_baddr[31:12]
Instantiated	20	166	4600	3090	0	fub_inst.sarb_bindsamp_baddr[31:12]
Instantiated	20	77.6	4390	2950	0	fub_inst.sarb_surfstate_baddr[31:12]
Instantiated	20	118	4330	2920	0	fub_inst.sarb_bindresrc_baddr[31:12]

Fig 5c: Snippet from Data Aware Clock Gating Efficiency (DACGE) report

# Redundant memory accesses

- Subsequent reads or writes to the same address in the memory is redundant and leads to power wastage
- Memories can contribute to significant dynamic power consumption

Dynamic Power Breakdown

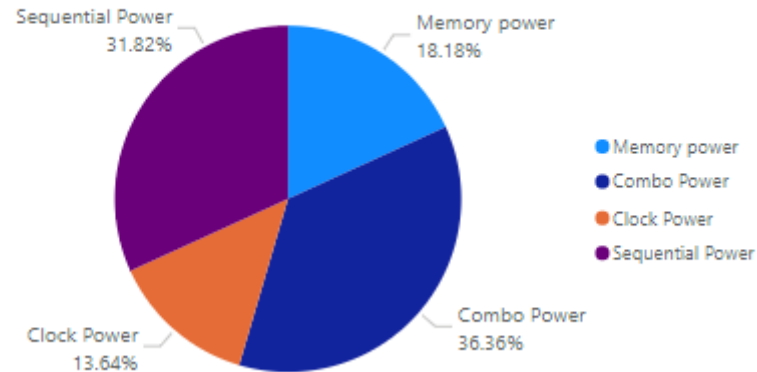


Fig 6a: Example memory dynamic power contribution to total dynamic power

Cell Name	Instance	Total Clock Cycles	No-Access Cycles	Read Cycles	Write Cycles	Redundant Read	Redundant Writes
memdphc66x384	ip_top.inst1.mem0	1760	0	1352	408	38	4
memdphc66x384	ip_top.inst1.mem1	2032	0	1432	600	40	3
memdphc66x384	ip_top.inst1.mem2	1654	0	1352	302	17	2
memdphc66x384	ip_top.inst1.mem3	2712	0	1906	806	42	9
memdphc66x384	ip_top.inst1.mem4	1442	0	1360	82	2	2
memdphc66x384	ip_top.inst1.mem5	1541	0	1352	189	16	5
memdphc66x384	ip_top.inst1.mem6	1701	0	1312	389	41	6
memdphc66x384	ip_top.inst1.mem7	2098	0	1312	786	36	4
memdphc66x384	ip_top.inst1.mem8	2889	0	1998	891	92	8

Fig 6b: Snippet from redundant memory access reports

# Clock Gating Efficiency Regression and Tracking

Track both idle and active DACGE (Data Aware Clock Gating Efficiency) reports

- Ideally clock pin of flop should toggle only when there is valid D pin transition for DACGE to be 100%
- For active use-cases/workloads, target to achieve DACGE > 98%
- During IDLE, target to achieve DACGE > 99.5%

Trend - Hierarchical DACGE

Workweek ● 2024WW04 ● 2024WW06 ● 2024WW10 ● 2024WW12

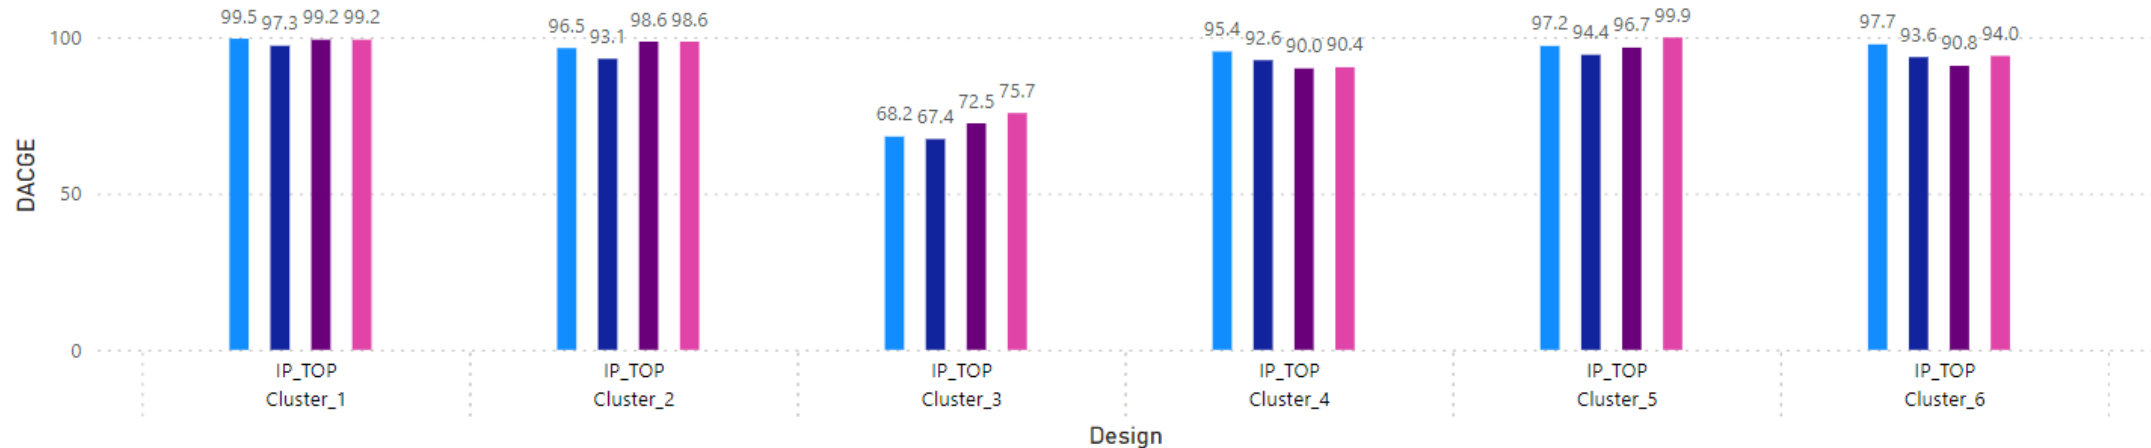


Fig 7a: Cluster wise DACGE trend tracking

# Summary

- Successfully identified and fixed regressions up to **3%** in dynamic power throughout the course of project between successive versions of the GFX IP
- Achieved significant power reduction up to **1.5%** at the GFX IP level as measured on the workloads
- These savings are on top of any planned RTL/architectural power savings and typically the targets for such savings are limited to <5% during project execution. So, this is significant savings from a single methodology.
- Avoided late power related surprises by left-shifting the power tracking process
- Discovered areas where physical design was compromising power for timing

# Thank You