

An ML-based System-Level Framework for Back-Annotating Low-Level Effects

Katayoon Basharkhah, Zain Navabi


School of Electrical and Computer Engineering, College of Engineering, University of Tehran, Tehran, Iran



1- Paper and pencil profiling is still on the stage

Going back to the history of IC performance characteristics, one comes up with passive diagrams in the datasheet. This requires a paper and pencil analysis of low-level properties like power and thermal dissipation.

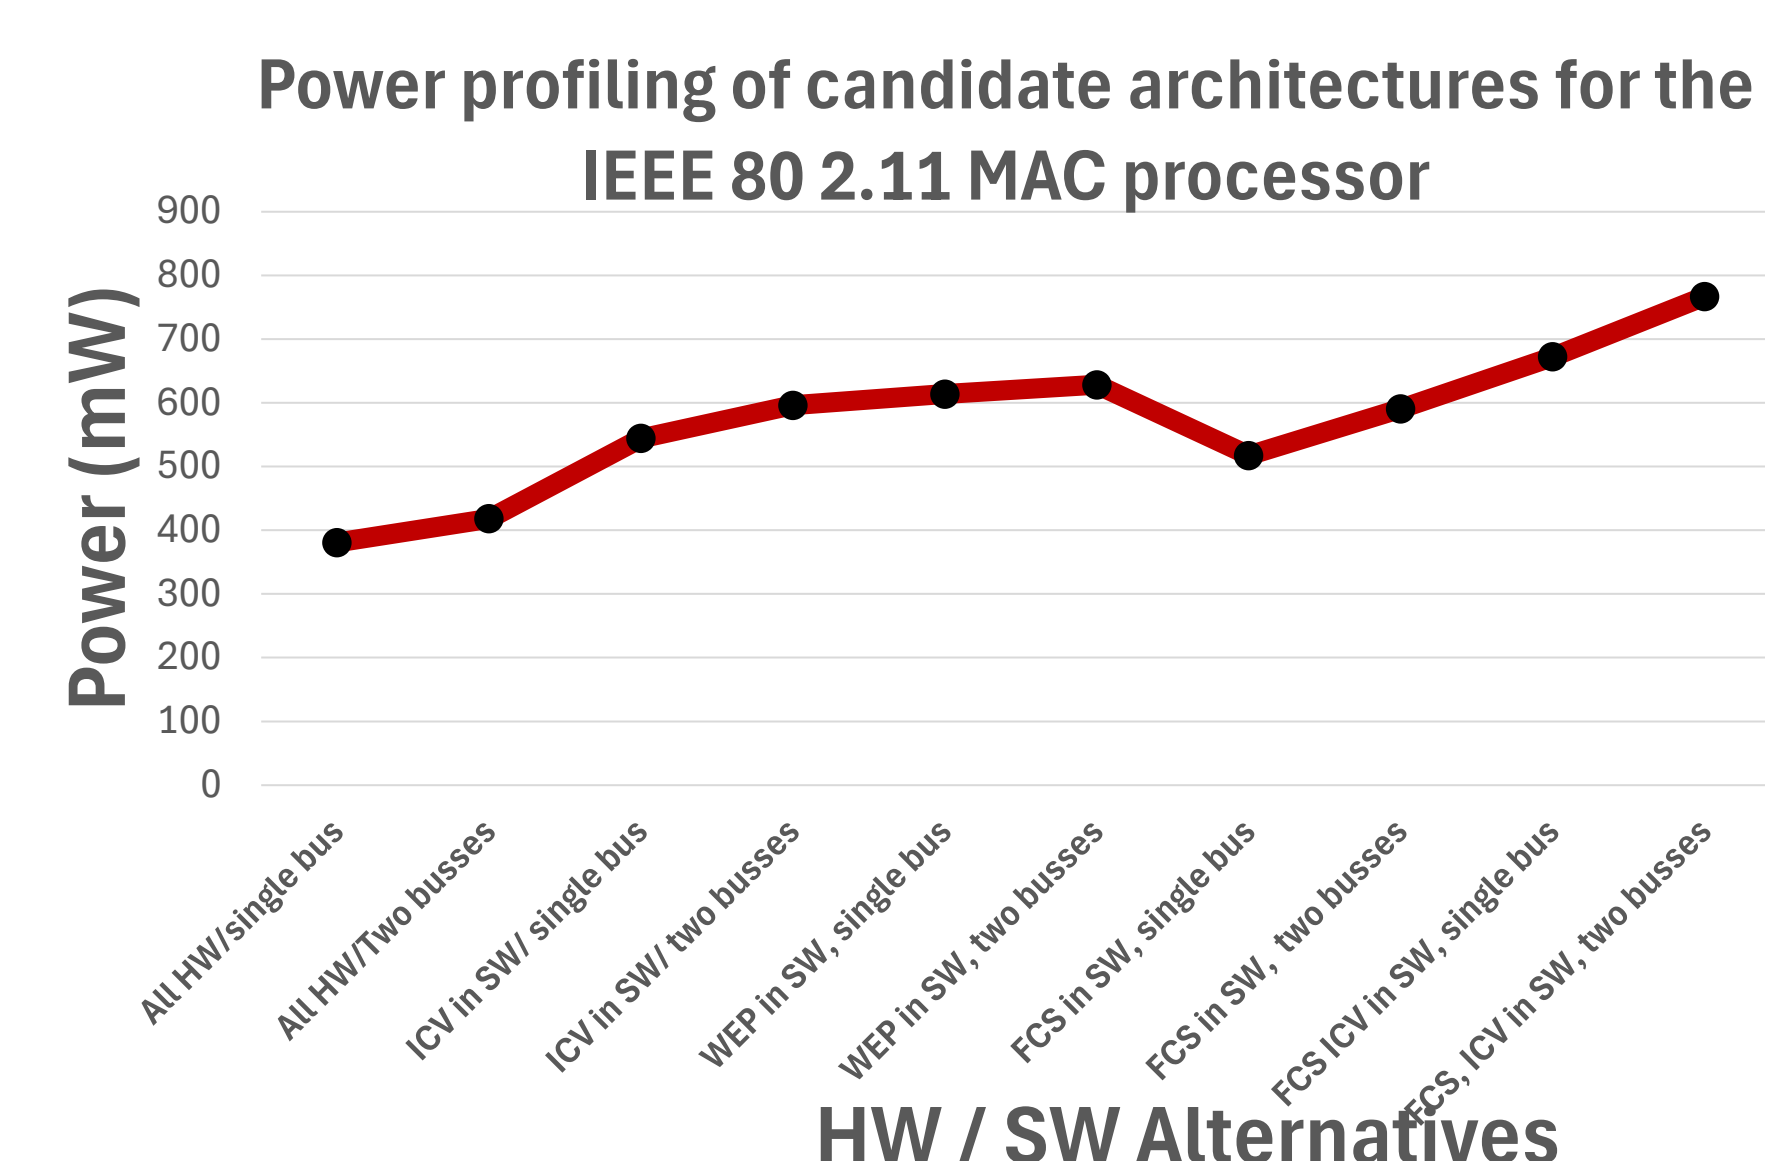
- Despite moving to very complex IP cores and processors, this coarse estimation of power still survives with similar information in a different format like a power spreadsheet.
- The power criteria are far more away from what real workload and benchmarks impose in a full-system stack



Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Configuration	Processor IA Core Frequency	Graphics core Frequency	Thermal Design Power (TDP) [w]	Scenario Design Power (SDP) [w]	Notes
H-Processor Line BGA	6-Core GT2 45W	Base	2.0 GHz to 2.9 GHz	1.1GHz to 1.2 GHz	45		
		Configurable TDP-Down	1.6 GHz to 2.4GHz	0.35 GHz	~34.5		
		LPM	0.8 GHz	0.35 GHz	~34.5		
	4-Core GT2 45W	Base	2.3 GHz to 3.0 GHz	1 GHz to 1.1 GHz	45		
		Configurable TDP-Down	1.8 GHz to 2.6 GHz	0.35 GHz	~34.5		
		LPM	0.8 GHz	0.35 GHz	~34.5		
U-Processor Line BGA	4-Core GT3 28W with OPC	Base	2.3 GHz to 2.7GHz	1.05 GHz to 1.2 GHz	28	N/A	1,9,10, 11,12, 15
		Configurable TDP-Down	1.7 GHz to 1.9GHz	0.3 GHz	~22.5		
		LPM	0.4 GHz	0.3 GHz	~22.5		
U-Processor Line BGA	2-Core GT3 28W with OPC	Base	2.7GHz to 3GHz	1.05GHz	28	N/A	1,9,10, 11,12, 15
		Configurable TDP-Down	1.7 GHz to 2.2GHz	0.3 GHz	~22.5		
		LPM	0.4 GHz	0.3 GHz	~22.5		

2- DSE needs more than minimal datasheet

- Core integrators need an accurate and fast power analysis of integrated IP cores
 - HW/SW partitioning
 - Workloads
 - Frequency
- Existing minimal datasheets are no more informative
- Going down to a full-system gate-level is impossible



3- We propose an active/live datasheet

- IP developer gives a live and active datasheet to the core integrator
 - High-level power model
 - A coded surrogate model
- Removes the low-level simulation burden for the core integrator

Intel's Desktop TDPs No Longer Useful to Predict CPU Power Consumption

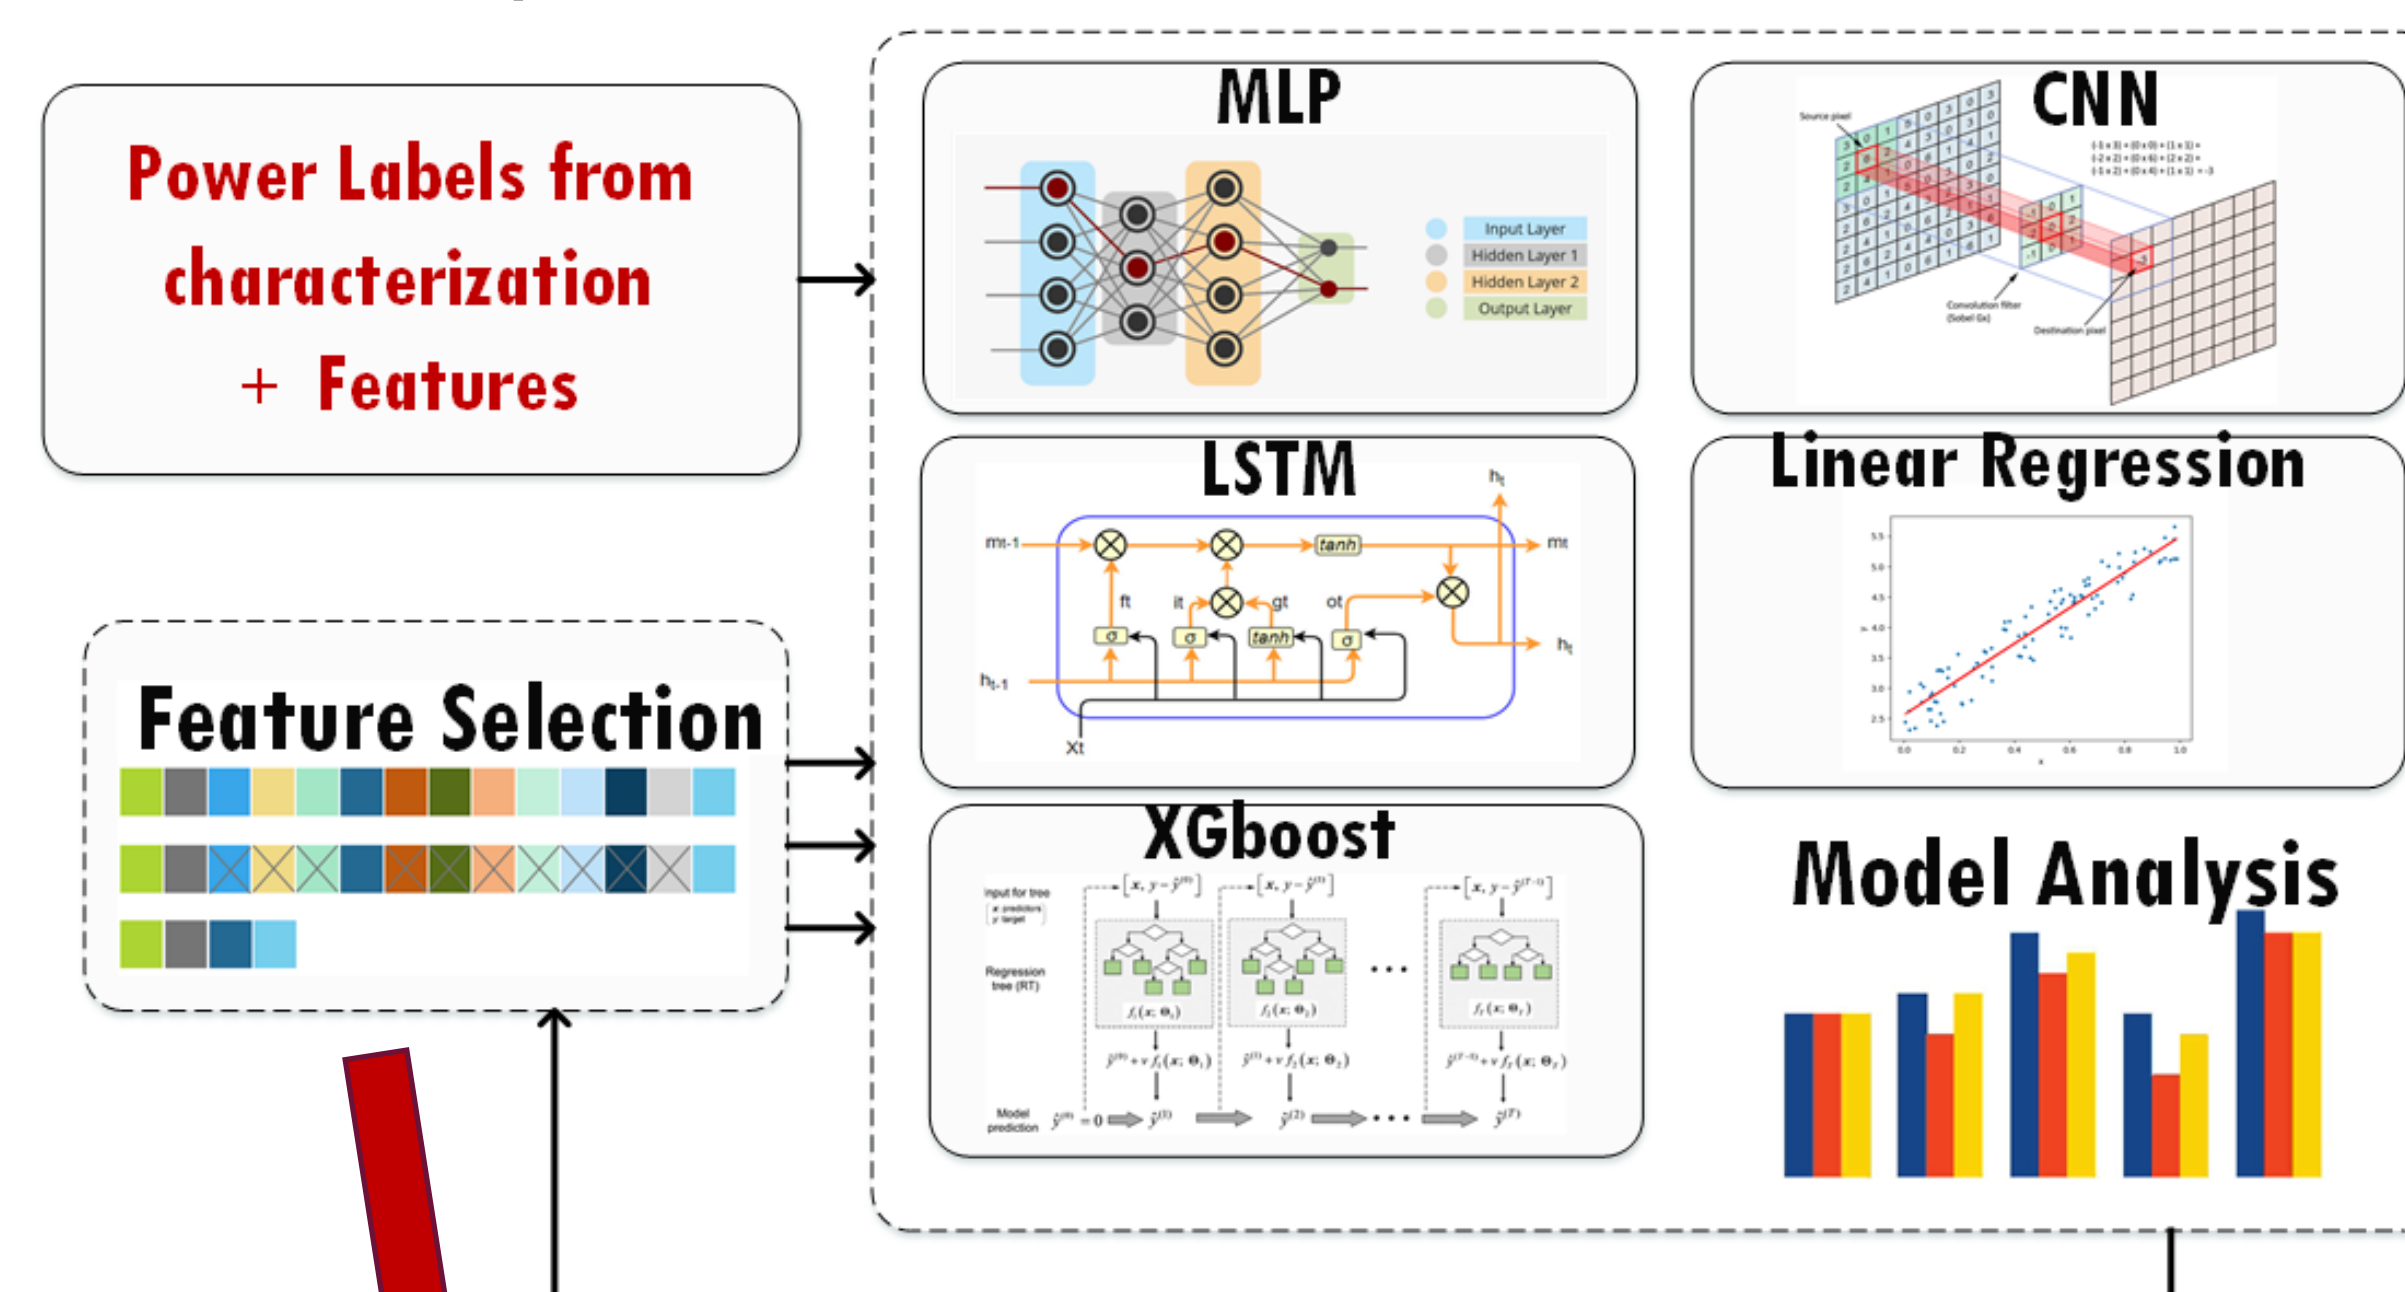
Intel's higher-end desktop CPU TDPs no longer communicate anything useful about the CPUs power consumption under load.

By Joel Hruska January 25, 2021





4- CPU ML model generation

- We propose an ML framework
 - Five different ML models
 - Feature selection mechanism
 - Selecting best model with the most important features



Feature	Number of Features	Instruction Stage	Instruction	TRF Port2	TRF Port1	Control Signals	Databus	Xgboost NRMSE	Xgboost Train Time
Scenario 1	57	*	*	*	*	*	*	5.1%	15.3
Scenario 2	59	*	*	*	*	*	*	5%	15.7
Scenario 3	3	*	*	*	*	*	*	2.77%	2.7
Scenario 4	5	*	*	*	*	*	*	2.71%	3.2
Scenario 5	61	*	*	*	*	*	*	2.52%	16.96P



5- Inference Results

- Inference evaluation in SystemC/C++
 - Running for a RISC-V-like CPU
 - Xgboost outperforms other ML methods
 - 2.5% estimation error for large MM application

