



Thermal Pressure:

Scheduler-Thermal Interactions

- Thara Gopinath, Linaro Inc



Introduction

- cpufreq cooling device can restrict the maximum frequency of a cpu to mitigate thermal events (passive response)
- Clipped max frequency implies clipped maximum possible capacity for a cpu.
- Thermal pressure - The delta between the maximum possible capacity of a cpu and maximum available capacity of a cpu due to thermal event

-



- Re-use the framework for tracking RT and DL utilization for tracking thermal pressure.
- - Provide mechanism to update thermal pressure from thermal framework to scheduler.



Where to find an initial prototype?

<https://lore.kernel.org/patchwork/cover/997754/>



Main Issue

Pelt signals – best way to track thermal utilization ?

Pros	Cons
Thermal pressure will track other scheduler signals	Thermal pressure decay tied to Pelt signal half life
	Some hacks involved in pelt library to track thermal pressure as other signals tracked by this library are "task running time based"




Test results that looks promising

hikey960 mainline kernel with debian file system.

- aobench (An occlusion renderer for benchmarking real-world floating point performance)

(Phoronix Test Suite)

	Result (secs)	Standard Error	Standard Deviation
Hikey 960 - no thermal pressure applied 1	96.54	12.4	15.46%
Hikey 960 - thermal pressure applied	162.54	20.61	31.06%



Hikey960 Android 4.9 kernel - 5 runs of PCMark with 10 mins wait in between for the chip to cool down

		Result	Standard Deviation	Standard Deviation % of average
Hikey960 - no fan - no thermal pressure applied				
pcmark	pcmark_DataManipulation	3296	1650	50.06%
pcmark	pcmark_PhotoEditingV2	11821	5939	50.24%
pcmark	pcmark_WebV2	7330	179	2.44%
pcmark	pcmark_WritingV2	5795	2929	50.54%
Hikey960- no fan - with thermal pressure applied				
pcmark	pcmark_DataManipulation	4413	204	4.62%
pcmark	pcmark_PhotoEditingV2	13749	1738	12.64%
pcmark	pcmark_WebV2	7424	245	3.30%
pcmark	pcmark_WritingV2	7067	1848	26.15%



		Result	Standard Deviation	Standard Deviation % of average
Hikey960 - with fan - no thermal pressure applied				
pcmark	pcmark_DataManipulation	5045	55	1.09%
pcmark	pcmark_PhotoEditingV2	17212	1771	10.29%
pcmark	pcmark_WebV2	7362	114	1.55%
pcmark	pcmark_WritingV2	8435	48	0.57%
Hikey960 - with fan - with thermal pressure applied				
pcmark	pcmark_DataManipulation	5076	21	0.41%
pcmark	pcmark_PhotoEditingV2	18133	81	0.45%
pcmark	pcmark_WebV2	7412	142	1.92%
pcmark	pcmark_WritingV2	8457	10	0.12%



ANYTHING ELSE ?