Hardware Counters Working Group Outbrief

Scalable Tools Workshop Granlibakken 2016

HW Counter Top Down Methodology

- Intel Experts: Michael Chynoweth, Ahmad Yasin
- Very important for understanding performance
- Descriptions exist but not for KNL or Atom
- Top down spreadsheet
 - https://download.01.org/perfmon/TMAM_Metrics.xlsx
- Collect many counters with a multiplexed run
 - is multiplexing OK?
 - fixed counters not impacted by multiplexing (fixed counter in clocks)
 - collect multiplexed data and compare with fixed counter attribution
 - if difference is greater than 10%, then multiplexed data is suspect
 - take weighted average to compare the data use PEBS to collect cycles (consumers of expensive loads)
 - Ask Stephane about how to validate multiplexing on Linux
- Fixed cost penalties: take them with grain of salt (may be overlapped)
- Understanding memory hierarchy issues on Intel architectures
 - use PEBS to track expensive loads (e.g. L3 misses)
 - use PEBS to collect cycles (consumers of expensive loads)

Intel Top-Down Methodology

TMAM	Version	3.1	Intel Confidential														
						Server	1	Server		Server		Server	r				
Key	Level1	Level2	Level3	Level4	SKL	BDX	BDW/BD	HSX	HSW	IVT	IVB	JKT/S	SNB	Locate-	Count Domain	Metric Description	Threshold
	_						W-DE					NB-EP	2	with			
FE	Fronten	d_Bound											IDQ.	UCSKL/SKX	Slots	This category represents slots fraction where	tht> 0.2
FE		Frontend_Latency					#Pipeline_	width *	IDQ_UC	PS_NO	T_DEU	VERED.	C #Pip	elin SKL/SKX,	(Slots	This metric represents slots fraction the CPU	wa:> 0.15 & P
FE			ICache_Misses		ICACH	E_168.I	FDATA_STAL	LL / CLKS	ICACHE	LIFDAT/	ICACH	E.IFETC	HN/A	SKL/SKX	(Clocks	This metric represents cycles fraction the CP	U w > 0.05 & P
FE			ITLB_Misses		ICACH	E_648.I	FTAG_STALL	/ CLKS					#ITL	B_NSKL/SKX	(Clocks	This metric represents cycles fraction the CP	U w > 0.05 & P
FE			Branch_Resteers		(INT_I	MISC.CL	EAR_RESTER	ER_CYCL	ES + 10'	BACLE	ARS.AN	n,)/cr	B #Av	R'BR_MISE	Clocks	This metric represents cycles fraction the CP	U w > 0.05 & P; \$issueBad; ~over
FE			DSB_Switches										DSB.	2MISKL/SKX	(Clocks	This metric represents cycles fraction the CP	U w > 0.05 & P
FE			MS_Switches										#MS	SVIDO, MS	1 Clocks	This metric estimates the fraction of cycles w	/her> 0.05 & P; \$issueMS
BAD		Frontend_Bandwidth											From	iten SKL/SKX,	(Slots	This metric represents slots fraction the CPU	wa:> 0.1 & (IPC > 2.0) & P
BAD	Bad_Sp	eculation											(00	PS_ISSUED.	AI Slots	This category represents slots fraction waste	d dı > 0.1; SissueBad
BAD		Branch_Mispredicts											#Mi:	spre BR_MISE	Slots	This metric represents slots fraction the CPU	has> 0.05 & P
BAD		Machine_Clears											Bad	Spi MACHIN	E Slots	This metric represents slots fraction the CPU	has> 0.05 & P
BE	Backend	Bound											1-(Frontend_B	o Slots	This category represents slots fraction where	no > 0.2
BE/Mem		Memory_Bound		_									#Me	mory_Bound	d Slots	This metric represents slots fraction the Men	noŋ> 0.2 & P
BE/Mem			L1_Bound		_		(CYCLE_AC	TIVITY.S	TALLS_	MEM_A	(#STA	LLS_ME	ET N/A	SKL/SKX	(Clocks	This metric estimates how often the CPU was	s sta(> 0.1 & P) DTLB_Load;
BE/Mem				DTLB_Load	(#Mer	n_STLB	(#Mem_ST	TLB_Hit_	Cost * I	DTLB_LO	DAD_M	ISSES.S	т (#М	lem SKL/SKX,	(Clocks	This metric represents cycles fraction where	the > 0.1
BE/Mem				FB_Full			Load_Miss	_Real_L	Load_N	Miss_Re	Load_	Miss_R	eal_ta	tency * L1D	Clocks	This metric does a *rough estimation* of how	<pre>w of > 0.3; \$issueBW; \$issueSL</pre>
BE/Mem			L2_Bound		(1 if FE	_Full <	1 CYCLE_AC	CTIVITY.S	TALLS_	L1D_M	(CYCL	E_ACTIV	V N/A	SKL/SKX	(Clocks	This metric estimates how often the CPU wa	s sta> 0.1 & P
BE/Mem			L3_Bound		(CYCL	E_ACTIN	#Mem_L3_	Hit_Fra	ction * o	CYCLE_/	CTIVIT	Y.STALL	S#Me	m_SKL/SKX	(Clocks	This metric estimates how often the CPU wa	s sta> 0.1 & P
BE/Mem			MEM_Bound		CYCLE	ACTIVI	T(1 - #Mem	L3_Hit	Fraction	n) * CYC	LE_ACT	IVITY.S	FT (1 - 1	MM SKL/SKX	(Clocks	This metric estimates how often the CPU wa	s sta> 0.1 & P
BE/Mem				MEM_Bandwidth									#OR	O_DRD_BW	Clocks	This metric estimates cycles fraction where t	he c> 0.1 & P; SissueBW
BE/Mem				MEM_Latency									FOR	O_DRD_Any	Clocks	This metric estimates cycles fraction where t	he t > 0.1 & P
BE/Mem			Store_Bound		EXE_A	CTIVITY	BOUND_ON	STORE	S/CLKS	5			RESO	DUFSKL/SKX	/(Clocks	This metric estimates how often CPU was sta	illed> 0.2 & P
BE/Core		Core_Bound					-	_					Back	end_Bound	- Slots	This metric represents slots fraction where C	ore > 0.2 & P
BE/Core			Divider		ARITH	DIVIDE	RARITH.FPU	DIV_AC	10 * AF	ND.HTIS	IDER_	JOPS /	CARIT	H.F SKL/SKX	/t Clocks	This metric represents cycles fraction where	the > 0.2 & P
BE/Core			Ports_Utilization		(#Bac	kend_B	o (#Backend	Bound	Cycles	- RESO	URCE_S	TALLS.S	SE (#Ba	ackend_Bour	n Clocks	This metric estimates cycles fraction the CPU	pei> 0.2 & P
RET	Retiring				-	_	-	-			-		UOP	S RETIRED.	I Slots	by useful work i.e. issued uops that eventual	ly (> 0.7 Microcode_Sequence
RET		Base											Reti	ring INST_RE	T Slots	This metric represents slots fraction where the	he C> 0.6 & P
RET			FP_Arith				FP_Scalar +	FP_Vec	N/A				FP S	scalar + FP_V	e Uops	This metric represents overall arithmetic floa	ting> 0.2 & P
RET				FP_Scalar			(FP_ARITH	INST P	N/A		(FP_C	OMP_C	OF (FP	COMP_OPS	Uops	This metric represents arithmetic floating-po	int > 0.1 & P
RET				FP_Vector			(FP_ARITH	INST P	N/A		(FP C	OMP O	OF (PP	COMP OPS	Uops	This metric represents arithmetic floating-po	int > 0.2 & P
RET			Other				1 - FP Arit	h –	N/A			-	1-F	P Arith	Uops	This metric represents non-floating-point (FP) uc> 0.3 & P
RET		Microcode_Sequencer		-			-						#Ref	tire IDQ.MS	Slots	This metric represents slots fraction the CPU	wa:> 0.05; \$issueMS

Gooda

- Gooda contains lists of interesting events for analyzing performance of several architectures
- List of penalties for understanding performance on different architectures
- https://github.com/David-Levinthal/gooda



Hardware Counter High Level Issues

- Intel needs constructive feedback
 - people care about performance counters
 - changing list of events makes it hard to use
- Validating events for processors is difficult
 - finite resources for validation
 - some events go private when not useful

Recommendations for Tuning

- Use auto feedback driven optimization with Intel and GCC compilers
- Data that can be collected with perf
- perf record -e branch_instruction_retired_taken -b (uses LBR) -o -
 - puts results in a pipe mode
 - use scripts to convert this to a gcov file