

Chapter 2

Validation of Hardware Events for Successful Performance Pattern Identification in High Performance Computing

Thomas Röhl, Jan Eitzinger, Georg Hager and Gerhard Wellein

Abstract Hardware performance monitoring (HPM) is a crucial ingredient of performance analysis tools. While there are interfaces like LIKWID, PAPI or the kernel interface `perf_event` which provide HPM access with some additional features, many higher level tools combine event counts with results retrieved from other sources like function call traces to derive (semi-)automatic performance advice. However, although HPM is available for x86 systems since the early 90s, only a small subset of the HPM features is used in practice. Performance patterns provide a more comprehensive approach, enabling the identification of various performance-limiting effects. Patterns address issues like bandwidth saturation, load imbalance, non-local data access in ccNUMA systems, or false sharing of cache lines. This work defines HPM event sets that are best suited to identify a selection of performance patterns on the Intel Haswell processor. We validate the chosen event sets for accuracy in order to arrive at a reliable pattern detection mechanism and point out shortcomings that cannot be easily circumvented due to bugs or limitations in the hardware.

2.1 Introduction and Related Work

Hardware performance monitoring (HPM) was introduced for the x86 architecture with the Intel Pentium in 1993 [15]. Since that time, HPM gained more and more attention in the computer science community and consequently a lot of HPM related tools were developed. Some provide basic access to the HPM registers with some additional features like LIKWID [17], PAPI [12] or the kernel interface `perf_event` [4]. Furthermore, some higher level analysis tools gather additional information by combining the HPM counts with application level traces. Popular representatives of that analysis method are HPCToolkit [1], PerfSuite [10], OpenSpeedshop [16] or Scalasca [3]. The intention of these tools is to advise the application developer with educated optimization hints. To this end, the tool

T. Röhl (✉) · J. Eitzinger · G. Hager · G. Wellein
Erlangen Regional Computing Center (RRZE), University of Erlangen-Nuremberg,
Erlangen, Germany
e-mail: thomas.roehl@fau.de

developers use performance metrics that represent a possible performance limitation, such as saturated memory bandwidth or instructions paths. The hardware metrics may be combined with information on the application level, e.g. scaling characteristics, dependence of performance on the problem size, or static code analysis, to arrive at a *signature*. A performance signature then points towards one or more *performance patterns*, as described in [19] and refined in [18]. The purpose of the patterns concept is to facilitate the identification of performance-limiting bottlenecks.

C. Guillen uses in [5] the term *execution properties* instead of performance pattern. She defines execution properties as a set of values gathered by monitoring and related thresholds. The properties are arranged in decision trees for compute- and memory-bound applications as well as trees related to I/O and other resources. This enables either a guided selection of the analysis steps to further identify performance limitation or automatic tool-based analysis. Based on the path in the decision tree, suggestions are given for what to look for in the application code. A combination of the structured performance engineering process in [18] with the decision trees in [5] defines a good basis for (partially) automated performance analysis tools.

One main problem with HPM is that none of the main vendors for x86 processors guarantees event counts to be accurate or deterministic. Although many HPM interfaces exist, only little research has been done on validating the hardware performance events. However, users tend to trust the returned HPM counts and use them for decisions about code optimization. One should be aware that HPM measurements are only guideposts until the HPM events are known to have guaranteed behavior. Moreover, analytic performance models can only be validated if this is the case.

The most extensive event validation analysis was done by Weaver et al. [20] using a self-written assembly validation code. They test determinism and overcounting for the following events: retired instructions, retired branches, retired loads and stores as well as retired floating-point operations including scalar, packed, and vectorized instructions. For validating the measurements the dynamic binary instrumentation tool *Pin* [11] was used. The main target of that work was not to identify the right events needed to construct accurate performance metrics but to find the sources of non-determinism and over/undercounting. It gives hints on how to reduce over- or undercounting and identify deterministic events for a set of architectures.

D. Zapanu et al. [21] determined the error of retired instructions and CPU cycle counts with two microbenchmarks. Since the work was released before the `perf_event` interface [4] was available for PAPI, they tested the deprecated interfaces `perfmon2` [2] and `perfctr` [13] as the basis for PAPI. They use an “empty” microbenchmark to define a default error using different counter access methods. For subsequent measurements they use a simple loop kernel with configurable iterations, define a model for the code and compare the measurement results to the model. Moreover, they test whether the errors change for increasing measurement duration and for a varying number of programmed counter registers. Finally, they give suggestions which back-end should be used with which counter access pattern to get the most accurate results.

In the remainder of this section we recommend HPM event sets and related derived metrics that represent the signature of prototypical examples picked out of the perfor-

mance patterns defined in [18]. In the following sections the accuracy of the chosen HPM events and their derived metrics is validated. Our work can be seen as a recommendation for tool developers which event sets match the selected performance patterns in the best way and how reliable they are.

2.2 Identification of Signatures for Performance Patterns

Performance patterns help to identify possible performance problems in an application. The measurement of HPM events is one part of the pattern's signature. There are patterns that can be identified by HPM measurements alone, but commonly more information is required, e.g., scaling behavior or behavior with different data set sizes. Of course, some knowledge about the micro-architecture is also required to select the proper event sets for HPM as well as to determine the capabilities of the system. For x86 systems, HPM is not part of the instruction set architecture (ISA), thus besides a few events spanning multiple micro-architectures, each processor generation defines its own list of HPM events. Here we choose the Intel Haswell EP platform (E5-2695 v3) for HPM event selection and verification. The general approach can certainly be applied to other architectures.

In order to decide which measurement results are good or bad, the characteristics of the system must be known. C. Guillen established thresholds in [5] with four different approaches: hardware characteristics, expert knowledge about hardware behavior and performance optimization, benchmarks and statistics. With decision trees but without source code knowledge it is possible to give some loose hints how to further tune the code. With additional information about the software code and run time behavior, the list of hints could be further reduced.

The present work is intended to be a referral for which HPM events provide the best information to specify the signatures of the selected performance patterns. The patterns target different behaviors of an application and/or the hardware and therefore are classified in three groups: *bottlenecks*, *hazards* and *work-related* patterns. The whole list of performance patterns with corresponding event sets for the Intel Haswell EP micro-architecture can be found at [14]. For brevity we restrict ourselves to three patterns: *bandwidth saturation*, *load imbalance* and *false sharing of cache lines*. For each pattern, we list possible signatures and shortcomings concerning the coverage of a pattern by the event set. The analysis method is comparable to the one of D. Zaparanuks et al. [21] but uses a set of seven assembly benchmarks and synthetic higher level benchmark codes that represent often used algorithms in scientific applications. But instead of comparing the raw results, we use derived metrics, combining multiple counter values, for comparison as these metric results are commonly more interesting for tool users.

2.2.1 Bandwidth Saturation

A very common bottleneck is bandwidth saturation in the memory hierarchy, notably at the memory interface but also in the L3 cache on earlier Intel designs. Proper identification of this pattern requires an accurate measurement of the data volume, i.e., the number of transferred cache lines between memory hierarchy levels. From data volume and run time one can compute transfer bandwidths, which can then be compared with measured or theoretical upper limits.

Starting with the Intel Nehalem architecture, Intel separates a CPU socket in two components, the core and the uncore. The core embodies the CPU cores and the L1 and L2 caches. The uncore covers the L3 cache as well as all attached components like memory controllers or the Intel QPI socket interconnect. The transferred data volume to/from memory can be monitored at two distinct uncore components. A CPU socket in an Intel Haswell EP machine has at most two memory controllers (iMC) in the uncore, each providing up to four memory channels. The other component is the Home Agent (HA) which is responsible for the protocol side of memory interactions.

Starting with the Intel Sandy Bridge micro-architecture, the L3 cache is segmented, with one segment per core. Still one core can make use of all segments. The data transfer volume between the L2 and L3 caches can be monitored in two different ways: One may either count the cache lines that are requested and written back by the L2 cache, or the lookups for data reads and victimized cache lines that enter the L3 cache segments. It is recommended to use the L2-related HPM events because the L3 cache is triggered by many components besides the L2 caches. Moreover, the Intel Haswell EP architecture has up to 18 L3 cache segments which all need to be configured separately. Bandwidth bottlenecks between L1 and L2 cache or L1 and registers are seldom and thus ignored in this pattern.

2.2.2 Load Imbalance

The main characterization of this pattern is that different threads have to process different working sets between synchronization points. For data-centric workloads the data volume transferred between the L1 and L2 caches for each thread may be an indicator: since the working sets have different sizes, it is likely that smaller working sets also require less data. However, the assumption that working set size is related to transferred cache lines is not expressive enough to fully identify the pattern, since the amount of required data could be the same for each thread while the amount of in-core instructions differs. Retired instructions, on the other hand, are just as unreliable as data transfers because parallelization overhead often comprises spin-waiting loops that cause abundant instructions without doing “work.” Therefore, for better classification, it is desirable to count “useful” instructions that perform the actual work the application has to do. None of the two x86 vendors provides features to filter the instruction stream and count only specific instructions in a sufficiently

flexible way. Moreover, the offered hardware events are not sufficient to overcome this shortcoming by covering most “useful” instructions like scalar/packed floating-point operations, SSE driven calculations or string related operations. Nevertheless, filtering on some instruction groups works for Intel Haswell systems, such as long-latency instructions (div, sqrt, ...) or AVX instructions. Consequently, it is recommended to measure the work instructions if possible but also the data transfers can give a first insight.

2.2.3 False Cache Line Sharing

False cache line sharing occurs when multiple cores access the same cache line while at least one is writing to it. The performance pattern thus has to identify bouncing cache lines between multiple caches. There are codes that require true cache line sharing, like producer/consumer codes, but we are referring to common HPC codes where cache line sharing should be as minimal as possible. In general, the detection of false cache line sharing is very hard when restricting the analysis space only to hardware performance measurements. The Intel Haswell micro-architecture offers two options for counting cache line transfers between private caches: There are L3 cache related μ OPs events for intra- and inter-socket transfers, but the HPM event for intra-socket movement may undercount with SMT enabled by as much as 40 % according to erratum HSW150 in [8]. The alternative is the *offcore response unit*. By setting the corresponding filter bits, the L3 hits with hitm snoops (hit a modified cache line) to other caches on the socket and the L3 misses with hitm snoops to remote sockets can be counted. The specification update [8] also lists an erratum for the offcore response unit (HSW149) but the required filter options for shared cache lines are not mentioned in it. There are no HPM events to count the transfers of shared cache lines at the L2 cache. In order to clearly identify whether a code triggers true or false cache line sharing, further information like source code analysis is required.

2.3 Useful Event Sets

Table 2.1 defines a range of HPM event sets that are best suitable for the described performance patterns regarding the HPM capabilities of the Intel Haswell EP platform. The assignment of HPM events for the pattern signatures is based on the Intel documentation [6, 9]. Some events are not mentioned in the default documentation; they are taken from Intel’s performance monitoring database [7]. Although the events were selected with due care, there is no official guarantee for the accuracy of the counts by the manufacturer. The sheer amount of performance monitoring related errata for the Intel Haswell EP architecture [8] reduces the confidence even further. But this encourages us even more to validate the chosen event sets in order to provide tool developers and users a reliable basis for their performance analysis.

Table 2.1 Desired events and available events for three performance patterns on the Intel Haswell EP micro-architecture

Pattern	Desired events	Available events
Bandwidth saturation	Data volume transferred to/from memory from/to the last level cache; data volume transferred between L2 and L3 cache	iMC:UNC_M_CAS_COUNT.RD, iMC:UNC_M_CAS_COUNT.WR, HA:UNC_H_IMC_READS.NORMAL, HA:UNC_H_BYPASS_IMC.TAKEN, HA:UNC_H_IMC_WRITES.ALL, L2_LINES_IN.ALL, L2_TRANS.L2_WB, CBOX:LLC_LOOKUP.DATA_READ, CBOX:LLC_VICTIMS.M_STATE
Load imbalance	Data volume transferred at all cache levels; number of “useful” instructions	L1D.REPLACEMENT, L2_TRANS.L1D_WB, L2_LINES_IN.ALL, L2_TRANS.L2_WB, AVX_INSTS.CALC, ARITH.DIVIDER_UOPS
False sharing of cache lines	All transfers of shared cache lines for the L2 and L3 cache; all transfers of shared cache lines between the last level caches of different CPU sockets	MEM_LOAD_UOPS_L3_ HIT_RETIRED.XSNP_HITM, MEM_LOAD_UOPS_L3_ MISS_RETIRED.REMOTE_HITM, OFFCORE_RESPONSE: LLC_HIT:HITM_OTHER_CORE, OFFCORE_RESPONSE: LLC_MISS:REMOTE_HITM

A complete list can be found at [14]

2.4 Validation of Performance Patterns

Many performance analysis tools use the HPM features of the system as their main source of information about a running program. They assume event counts to be correct, and some even generate automated advice for the developer. Previous research in the field of HPM validation focuses on singular events like retired instructions but does not verify the results for other metrics that are essential for identifying performance patterns. Proper verification requires the creation of benchmark code that has well-defined and thoroughly understood performance features and, thus, predictable event counts. Since optimizing compilers can mutilate the high level code, the feasible solutions are either to write assembly benchmarks or to perform code analysis of the assembly code created by the compiler.

The LIKWID tool suite [17] includes the `likwid-bench` microbenchmarking framework, which provides a set of assembly language kernels. They cover a variety of streaming access schemes. In addition the user can extend the framework by writing new assembly code loop bodies. `likwid-bench` takes care of loop counting, thread parallelism, thread placement, ccNUMA page placement and performance (and bandwidth) measurement. It does not, however, perform hardware event

counting. For the HPM measurements we thus use `likwid-perfctr`, which is also a part of the `LIKWID` suite. It uses a simple command line interface but provides a comprehensive set of features for the users. `Likwid-perfctr` supports almost all interesting core and uncore events for the supported CPU types. In order to relieve the user from having to deal with raw event counts, it supports *performance groups*, which combine often used event sets and corresponding formulas for computing derived metrics (e.g., bandwidths or FLOP rates). Moreover, `likwid-perfctr` provides a *Marker API* to instrument the source code and restrict measurements to certain code regions. `Likwid-bench` already includes the calls to the Marker API in order to measure only the compute kernel. We have to manually correct some of the results of `likwid-bench` to represent the obvious and hidden data traffic (mostly write-allocate transfers) that may be measured with `likwid-perfctr`.

The first performance pattern for the analysis is the bandwidth saturation pattern. For this purpose, `likwid-perfctr` already provides three performance groups called L2, L3 and MEM [17]. A separate performance group was created to measure the traffic traversing the HA. Based on the raw counts, the groups define derived metrics for data volume and bandwidth. For simplicity we use the derived metric of total bandwidth for comparison as it both includes the data volume in both directions and the run time. In Fig. 2.1 the average, minimal and maximal errors of 100 runs




Fig. 2.1 Verification tests for cache and memory traffic using a set of micro benchmarking kernels written in assembly. We show the average, minimum and maximum error in the delivered HPM counts for a collection of streaming kernels with data in L2, L3 and in memory

with respect to the exact bandwidth results are presented for seven streaming kernels and data in L2 cache, L3 cache and memory. The locality of the data in the caching hierarchy is ensured by streaming accesses to the vectors fitting only in the relevant hierarchy level. The first two kernels (*load* and *store*) perform pure loading and storing of data to/from the CPU core to the selected cache level or the memory. A combination of both is applied in the *copy* test. The last three tests are related to scientific computing and well understood. They range from the linear combination of two vectors called *daxpy* calculating $A[i] = B[i] \cdot c + A[i]$, a *stream* triad with formula $A[i] = B[i] \cdot c + C[i]$ to a vector *triad* computing $A[i] = B[i] \cdot C[i] + D[i]$.

The next pattern we look at is load imbalance. Since load imbalance requires a notion of “useful work” we have to find a way to measure floating-point operations. Unfortunately, the Intel Haswell architecture lacks HPM events to fully represent FLOP/s. For the Intel Haswell architecture, Intel has documented a HPM event `AVX_INSTS.ALL` (Event 0xC6, Umask 0x07) which captures all AVX instructions including data movement and calculations [7]. With the help of `likwid-bench` we could further refine the event to count loads (Umask 0x01), stores (Umask 0x02) and calculations (Umask 0x04) separately. Consequently, the FLOP/s performed with AVX operations can be counted. All performance patterns that require the filtering of the instruction stream for specific instructions can use the event `AVX_INSTS.CALC` for floating-point operations using the AVX vectorization extension. Due to its importance, the event is verified using the `likwid-bench` utility with assembly benchmarks that are based on AVX instructions only. Note that the use of these specific Umasks is an undocumented feature and may change with processor generations or even mask revisions. Moreover, we have found no way to count SSE or scalar floating-point instructions.

Figure 2.2 shows the minimum, maximum and average error for measuring AVX FLOP/s. The average error for all tests is below 0.07 %. As the maximal error is 0.16 % the event can be seen as sufficiently accurate for pure AVX code. Using the counter with non-AVX codes always returns 0.

Coming back to performance patterns, we now verify the load imbalance pattern using an upper triangular matrix vector multiplication code running with two threads. Since the accuracy of the cache and memory traffic related HPM events have been verified already, we use the only available floating-point operation related event `AVX_INSTS.CALC`. There is one shortcoming worth noting: If the code contains half-wide loads, the HPM event shows overcounting. The compiler frequently uses half-wide loads to reduce the probability of “split loads,” i.e., AVX loads that cross a cache line boundary if 32-byte alignment cannot be guaranteed. Experiments have shown that the event `AVX_INSTS.CALC` includes the `vinserf128` instruction as a calculation operation. In order to get reliable results, split AVX loads should be avoided. This is not a problem with `likwid-bench` as no compiler is involved and the generated assembly code is under full control. The upper triangular matrix is split so that each of the two threads operates on half of the matrix. The matrix has a size of 8192×8192 and the multiplication is performed 1000 times. The first thread processes the top rows with totally 25,167,872 elements, while the second

Fig. 2.2 Verification tests for the AVX floating point event using a set of microbenchmarking kernels with pure AVX code




Table 2.2 Verification of the load imbalance pattern using an upper triangular matrix vector multiplication code

Event/Metric	Thread 0	Thread 1	Ratio	Error
Process elements	25167872	8390656	3 : 1	
AVX floating point ops	1.26E10	4.21E09	2.991 : 1	0.29 %
L2 data volume [GByte]	406.28	115.34	3.52 : 1	17.42 %
L3 data volume [GByte]	203.06	69.74	2.912 : 1	2.94 %
Memory data volume [GByte]	112.97	37.33	3.026 : 1	0.88 %

one works on the remaining 8,390,656 elements. This distribution results in a work load imbalance for the threads of 3 : 1.

Table 2.2 lists the verification data for the code. The AVX calculation instruction count fits to a high degree the work load ratio of 3 : 1. The L2 data volume has the highest error, mainly caused by repeatedly fetching the input and output vector not included in the work load balance model. This behavior also occurs for the L3 and memory data volume but to a lesser extent as the cache lines of the input vector commonly stay in the caches. In order to get the memory data volume per core, the offcore response unit was used.

The false sharing of cache lines pattern is difficult to verify as it is not easy to write code that shows a predictable number of inter-core cache line transfers. A minimal amount of shared cache lines exist in almost every code thus HPM results unequal zero cannot be accepted as clear signature. To measure the behavior, a producer and

Table 2.3 Verification tests for false sharing of cache lines using a producer/consumer code

Amount of shared cache lines per step	Transferred cache lines according to model	Avg. amount of intra-socket transferred shared cache lines	Error [%]	Avg. amount of inter-socket transferred shared cache lines	Error [%]
2	200	328.4	64.2	404.2	102.1
4	400	568.3	42.1	607.7	51.9
8	800	897.0	12.1	985.86	23.2
16	1600	1714.2	7.1	1777.1	11.1
32	3200	2893.4	-9.6	2595.6	-18.9
64	6400	5570.8	-13.0	3512.3	-45.1
128	12,800	7350.7	-42.6	6124.9	-52.2
256	25,600	8995.3	-64.9	11471.1	-55.2
512	51,200	18224.3	-64.4	22608.3	-55.8
1024	102,400	55124.5	-46.2	45814.3	-55.3

The producer and consumer thread are located on the same CPU socket

consumer code was written, thus we verify the amount of falsely shared cache lines by using a true sharing cache line code. The producer writes to a consecutive range of memory that is read afterwards by the consumer. In the next iteration the producer uses the subsequent range of memory to avoid invalidation traffic. The memory range is aligned so that a fixed amount of cache lines is used in every step. The producer and consumer perform 100 iterations in each of the 100 runs. For synchronizing the two threads, a simple busy-waiting loop spins on a shared variable with long enough sleep times to avoid high access traffic for the synchronization variable. When using pthread conditions and a mutex lock instead, the measured values are completely unstable.

Table 2.3 shows the measurements for HPM events fitting best to the traffic caused by false sharing of cache lines. The table lists the amount of cache lines that are written by the producer thread. Since the consumer reads all these lines, the amount of transferred cache lines should be in the same range. The measurements using the events in Table 2.1 show a big discrepancy between the counts in the model and the measured transfers. For small counts of transferred cache lines, the results are likely to be distorted by the shared synchronization variable, but the accuracy should improve with increasing transfer sizes. Since the erratum HSW150 in [8] states an undercounting by as much as 40 %, the intra-socket measurements could be too low. But even when scaling up the measurements the HPM event for intra-socket cache line sharing is not accurate.

For the inter-socket false sharing, the threads are distributed over the two CPU sockets in the system. The results in Table 2.3 show similar behavior as in the intra-socket case. The HPM events for cache line sharing provide a qualitative classification for the performance pattern's signature but no quantitative one. The problem is mainly

to define a threshold for the false-sharing rate of the system and application. Further research is required to create suitable signature for this performance pattern.

2.5 Conclusion

The performance patterns defined in [18] provide a comprehensive collection for analyzing possible performance degradation on the node level. They address possible hardware bottlenecks as well as typical inefficiencies in parallel programming. We have listed suitable event sets to identify the bandwidth saturation, load imbalance, and false sharing patterns with HPM on the Intel Haswell architecture. Unfortunately the hardware does not provide all required events, such as, e.g., scalar/packed floating-point operations, or they are not accurate enough like, e.g., the sharing of cache lines at the L3 level. Moreover, a more fine-grained and correct filtering of instructions would be helpful for pattern-based performance analysis.

Using a selection of streaming loop kernels we found the error for the bandwidth-related events to be small on average ($-1\% \dots +2\%$), with a maximum undercounting of about -6% for the L3 traffic. The load imbalance pattern was verified using an upper triangular matrix vector multiplication. Although the error for the L1 to L2 cache traffic is above 15% , the results reflect the correct load imbalance of roughly $3 : 1$, indicating the usefulness of the metrics. Moreover, we have managed to identify filtered events that can accurately count AVX floating-point operations under some conditions. FLOP/s and traffic data are complementary information for identifying load imbalance. The verification of the HPM signature for the false sharing pattern failed due to large deviations from the expected event counts for the two events used. More research is needed here to arrive at a useful procedure, especially for distinguishing unwanted false cache line sharing from traffic caused by intended updates.

The remaining patterns defined in [18] need to be verified as well to provide a well-defined HPM analysis method for performance patterns ready to be included in performance analysis tools. We provide continuously updated information about suitable events for pattern identification in the Wiki on the LIKWID website.¹

Acknowledgments Parts of this work were funded by the German Federal Ministry of Research and Education (BMBF) under Grant Number 01IH13009.

References

1. Adhianto, L., Banerjee, S., Fagan, M., Krentel, M., Marin, G., Mellor-Crummey, J., Tallent, N.R.: HPCToolkit: Tools for performance analysis of optimized parallel programs. *Concurr. Comput.: Pract. Exp.D* **22**(6), 685–701 (2010)

¹<https://github.com/RRZE-HPC/likwid>.

2. Eranian, S.: Perfmon2: a flexible performance monitoring interface for Linux. In: Ottawa Linux Symposium, pp. 269–288, Citeseer (2006)
3. Geimer, M., Wolf, F., Wylie, B.J., Ábrahám, E., Becker, D., Mohr, B.: The Scalasca performance toolset architecture. *Concurr. Comput.: Pract. Exp.* **22**(6), 702–719 (2010)
4. Gleixner, T., Molnar, I.: Linux 2.6.32: perf_event.h. <http://lwn.net/Articles/310260/> (2008)
5. Guillen, C.: Knowledge-based performance monitoring for large scale HPC architectures. Dissertation p. <http://mediatum.ub.tum.de/?id=1237547> (2015)
6. Intel: Intel 64 and IA-32 Architectures Software Developer Manuals. <http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html> (2015)
7. Intel: Intel Open Source Technology Center for PerfMon. <https://download.01.org/perfmon/> (2015)
8. Intel: Intel Xeon Processor E3-1200 v3 Product Family Specification Update. <http://www.intel.com/content/dam/www/public/us/en/documents/specification-updates/xeon-e3-1200v3-spec-update.pdf> (2015)
9. Intel: Intel Xeon Processor E5 v3 Family Uncore Performance Monitoring. <https://www-ssl.intel.com/content/dam/www/public/us/en/zip/xeon-e5-v3-uncore-performance-monitoring.zip> (2015)
10. Kuftrin, R.: Perfsuite: An accessible, open source performance analysis environment for linux. In: 6th International Conference on Linux Clusters: The HPC Revolution, vol. 151, p. 05. Citeseer (2005)
11. Luk, C.K., Cohn, R., Muth, R., Patil, H., Klauser, A., Lowney, G., Wallace, S., Reddi, V.J., Hazelwood, K.: Pin: Building customized program analysis tools with dynamic instrumentation. *SIGPLAN Not.* **40**(6), 190–200 (2005). <http://doi.acm.org/10.1145/1064978.1065034>
12. Mucci, P.J., Browne, S., Deane, C., Ho, G.: PAPI: A portable interface to hardware performance counters. In: Proceedings of the Department of Defense HPCMP Users Group Conference. pp. 7–10 (1999)
13. Pettersson, M.: Linux x86 performance-monitoring counters driver (2003)
14. Roehl, T.: Performance patterns for the Intel Haswell EP/EN/EX architecture. <https://github.com/RRZE-HPC/likwid/wiki/PatternsHaswellEP> (2015)
15. Ryan, B.: Inside the Pentium. *BYTE Mag.* **18**(6), 102–104 (1993)
16. Schulz, M., Galarowicz, J., Maghrak, D., Hachfeld, W., Montoya, D., Cranford, S.: Open SpeedShop: An open source infrastructure for parallel performance analysis. *Sci. Prog.* **16**(2–3), 105–121 (2008)
17. Treibig, J., Hager, G., Wellein, G.: LIKWID: A lightweight performance-oriented tool suite for x86 multicore environments. In: Proceedings of PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures. San Diego, CA (2010)
18. Treibig, J., Hager, G., Wellein, G.: Pattern driven node level performance engineering. http://sc13.supercomputing.org/sites/default/files/PostersArchive/tech_posters/post254s2-file2.pdf (2013), sC13 poster
19. Treibig, J., Hager, G., Wellein, G.: Performance patterns and hardware metrics on modern multicore processors: Best practices for performance engineering. *Euro-Par 2012: Parallel Processing Workshops. Lecture Notes in Computer Science*, vol. 7640, pp. 451–460. Springer, Berlin (2013)
20. Weaver, V., Terpstra, D., Moore, S.: Non-determinism and overcount on modern hardware performance counter implementations. In: 2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), pp. 215–224 (2013)
21. Zapanaruks, D., Jovic, M., Hauswirth, M.: Accuracy of performance counter measurements. In: IEEE International Symposium on Performance Analysis of Systems and Software, 2009. ISPASS 2009. pp. 23–32 (2009)

Tools for High Performance Computing 2015
Proceedings of the 9th International Workshop on
Parallel Tools for High Performance Computing,
September 2015, Dresden, Germany
Knüpfer, A.; Hilbrich, T.; Niethammer, C.; Gracia, J.;
Nagel, W.E.; Resch, M.M. (Eds.)
2016, VIII, 181 p. 64 illus., 41 illus. in color., Hardcover
ISBN: 978-3-319-39588-3