Experience with low-power x86 processors for HEP usage
An initial analysis of the Intel® dual core Atom™ N330 processor

G Balazs, S Jarp, A Nowak (CERN openlab)

IT Department, CERN, 1211 Geneva 23, Switzerland

<sverre.jarp@cern.ch>

Abstract. In this paper we compare a system based on an Intel Atom N330 low-power processor to a modern Intel Xeon® dual-socket server using CERN IT's standard criteria for comparing price-performance and performance per watt. The Xeon server corresponds to what is typically acquired as servers in the LHC Computing Grid. The comparisons used public pricing information from November 2008. After the introduction in section 1, section 2 describes the hardware and software setup. In section 3 we describe the power measurements we did and in section 4 we discuss the throughput performance results. In section 5 we summarize our initial conclusions. We then go on to describe our long term vision and possible future scenarios for using such low-power processors, and finally we list interesting development directions.

1. Introduction

Since the mid-90s, CERN and the High Energy Physics community have been using PC technology for practically all computing needs. The evolution of the computing capacity has been very impressive. Compared to the Pentium Pro systems used initially, processor frequency has increased by (at least) 20x and memory capacity has followed suit. In recent years we have moved successfully from single-core to quad-core chip technology. In our community, performance is typically measured with a throughput-based benchmark procedure, which basically initiates as many SPECINT streams as there are cores in a given system adding up the individual results. The cost of the system (including overheads for racking and connectivity) is used to compute a performance-per-currency-unit number.

In recent years, power consumption of a server has become a major headache and CERN now imposes a thermal penalty of 10 CHF/Watt when comparing servers of unequal consumption. This high number, which largely exceeds the real cost of power, is directly linked to the fact that CERN has to strive to keep all its computing inside a power limit of 2.5 MW dictated by its 30 year old Computer Centre building [9]. Low-power servers can therefore easily become advantageous, even if they do not have the highest absolute performance.

This fact led us to believe that the newly announced low-power Intel Architecture (IA) processor, Atom, could be of great interest to our community. This paper discusses our findings when comparing a dual-core Atom "home-built" system to a state-of-the-art production Xeon-based server. In order to compare prices, we chose to use current Web pricing, i.e. the prices we could readily find on the Web in the moment of writing (i.e. November 2008).

2. Setup

2.1. Hardware setup of the Atom system

The sample kit received from Intel contained a D945GCLF2 mini-ITX desktop board equipped with an Atom N330 dual-core processor running at 1.6 GHz. This particular Atom model supports the “Intel64” 64-bit addressing mode whereas some previous ones only ran in 32-bit mode. Each of the two cores has two hardware threads, so the system appears as having 4 CPUs in Linux terms.

The main features of the motherboard are the following:

- 1x DIMM DDR2 667/533 memory slot supporting modules up to 2GB
- 1x PCI connector
- 2x SATA II, 1x ATA 100/66 ports
- 1x 10/100/1000 LAN connection
- 4x USB 2.0 ports (up to 8x), integrated 6 channel audio, s-video out

The motherboard provides all necessary connectivity for a desktop system, but the 2GB limitation of the memory subsystem is rather limiting. In High Energy Physics (HEP) computing we usually start a process requiring up to 2GB of memory per Linux CPU, and this would imply 8GB in total (when enabling the hardware threads). Fortunately some workloads require less memory and work is underway to reduce the memory footprint even further through multithreading or process forking [1]. Thus, we analyze a scenario in which the requirement is 512 MB per running process.

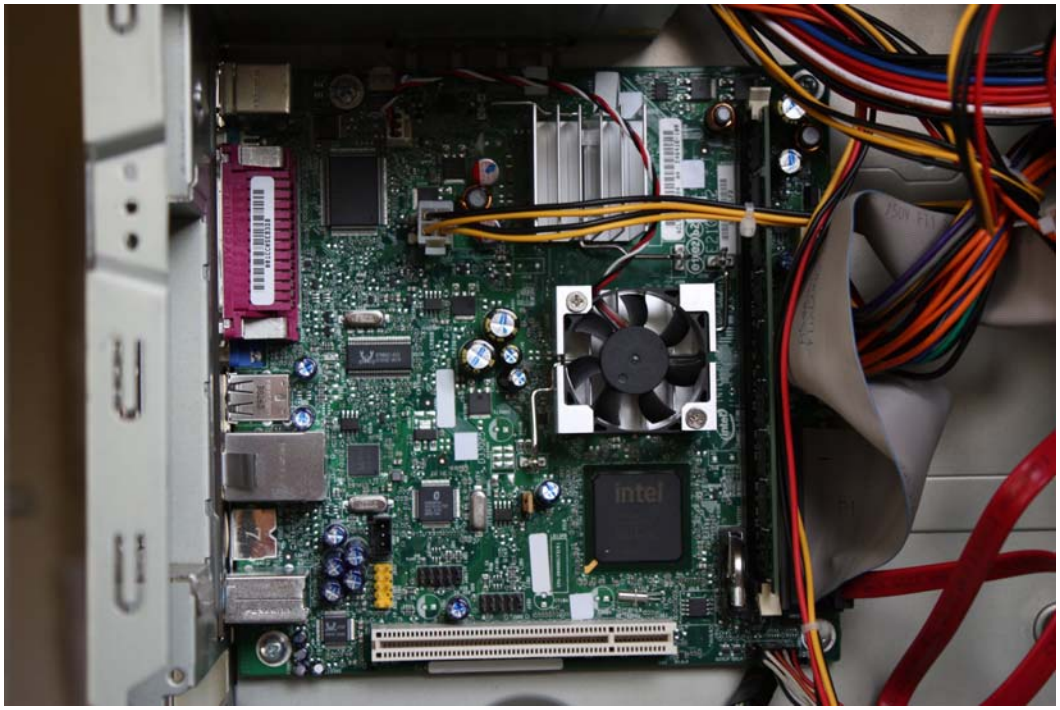

Figure 1: The Atom-based motherboard

The motherboard was installed into a standard ATX chassis with an optical drive and a 120GB IDE hard drive. The SATA II connectivity was tested and found to be satisfactory. The smallest power supply we could find (in a hurry) was a low-quality 350W version. The memory used was a 2GB 667MHz DDR2 module.

Final configuration:

CPU:	Intel Atom 330 Dual-core 1.60 GHz
Motherboard:	Intel Desktop Board D945GCLF2
Memory:	1x2GB 667MHz DDR2
Power supply:	350 Watts
Devices:	IDE HDD, DVD-ROM

2.2. System software setup

The sample kit was installed with Fedora Core 9 x86_64 (kernel: 2.6.25.14) without any particular problem. For the power efficiency tests (with disconnected drives) a Fedora Core 9 X86-64 Live CD was used, also without any issues. The installed compilers were GNU C++ 4.2.4 and 4.3.0 as well as the Intel C++ compiler (ICC 11.0, release 069).

2.3. Pricing information

In this paper the performance of the test system is being compared to a dual socket server with two 3GHz Xeon E5472 (Harpertown) CPUs, which can be considered as a typical system currently in use at CERN. One important component of the comparison is the cost of each system.

Using a Web search we found the following approximate price for an Atom N330 based computer in Swiss Francs (CHF):

Motherboard+CPU	110 CHF
2GB DDR2 memory	30 CHF
Power supply, drives	110 CHF
Total	**250 CHF**

A similar Web search indicated that a "ready-to-ship" Xeon server with two E5472 CPUs would cost about 5200 CHF. (It seemed possible to buy the components separately for about 10% less, but we did not take this into account):

2x E5472 CPU	3500 CHF
1x4GB DDR2 ECC memory	300 CHF
Other (board, PSU, drives)	1400 CHF
Total	**5200 CHF**

It must be kept in mind that the Atom system was equipped with 2GB of DDR2 memory, while the Xeon system was equipped with 4GB of DDR2 ECC FB-DIMM memory, both fulfilling the requirement of 0.5GB per process. We also take into account a second, hypothetical scenario, in which the Atom would support up to 8GB of memory, fulfilling current CERN memory requirements. It is compared to a Xeon system with the same requirement. The cost breakdown is presented below:

Atom system	
Motherboard+CPU	110 CHF
2x4GB DDR2 memory	150 CHF
Power supply, drives	110 CHF
Total	**370 CHF**

Xeon system	
2x E5472 CPU	3500 CHF
4x4GB DDR2 ECC memory	1200 CHF
Other (board, PSU, drives)	1400 CHF
Total	**6100 CHF**

We estimate the power consumption of the Atom system with 2x4 GB DDR2 memory and an optimized power supply to be at 56 Watts (25 W chipset + 8W CPU + 16W RAM + 7W HD)[1].

3. Power measurements

The power measurements were performed with a *ZES Zimmer LMG 500* power meter connected to a laptop via the RS232 interface. The test systems obtain power through a measurement adapter (LMG-MAK1) which enables the power meter to perform accurate measurements directly in the main AC circuit. The whole process is controlled centrally from a remote workstation which controls both the measured system and the power meter at the same time. The measured values are defined in Appendix A.

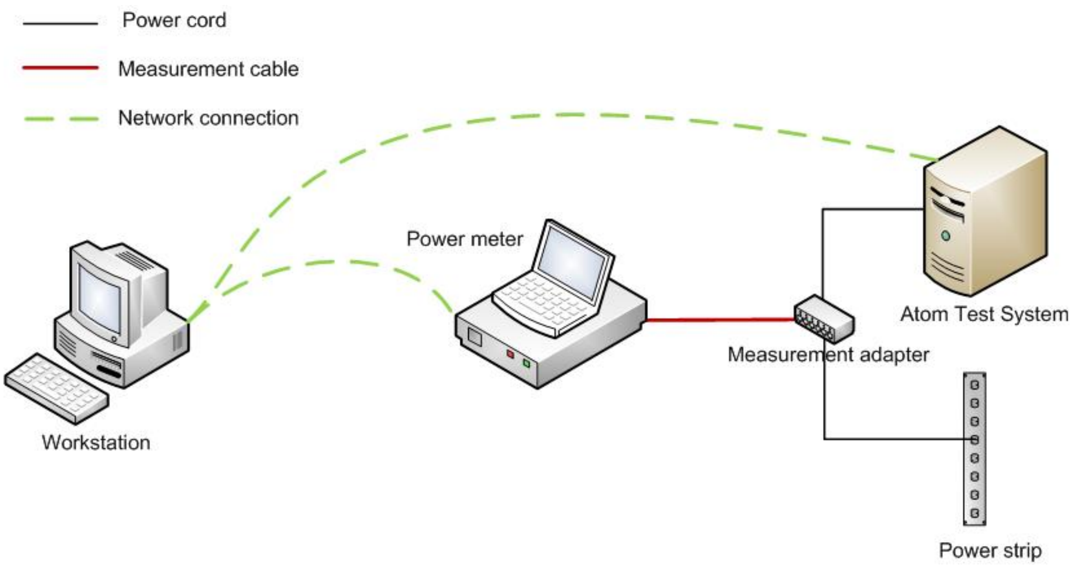

Figure 2: Power test setup

3.1. The tests

Two basic measurements were performed with the test system:

- The **Idle** test: the power consumption is measured while only the standard operating system processes are running, all components are considered to be in an idle state.

[1] Memory consumption estimation was approximate and based on work by G. Balazs [2]. Thus, more precise figures relating to real measurements might vary from our estimates.

- The **Cpuburn** test: Cpuburn is part of the standard CERN test toolkit for power measurements during the tendering process for new servers. It generates an artificial full load only on the processors while the memory subsystem remains idle.

3.2. Results

Two sets of tests of the Atom system were conducted with different hardware configurations:

3.2.1. Results with the initial Atom configuration

Configuration 1	
Memory:	1*2GB 667MHz DDR2
Power supply:	350 Watts
Devices:	HDD, DVD

Measurement Results

	Idle	Cpuburn
Active Power	46.5 W	50.7 W
Apparent Power	65.3 VA	70.1 VA
Power Factor	0.71	0.72

The ~50 Watts active power consumption for a complete system is remarkably good compared to other desktop systems. However, it should be noted that the power factor of an average consumer power supply is unimpressive and that the power consumption on the final bill will be higher.

3.2.2. *Tests to reach the lowest possible power consumption in the Atom system*.

Another measurement was performed with the HDD and the optical drive disconnected in order to get more accurate results for the motherboard and the CPU.

Configuration 2	
Memory:	1*2GB 667MHz DDR2
Power supply:	350 Watts
Devices:	None

Measurement Results

	Idle	Cpuburn
Active Power	37.6 W	38.9 W
Apparent Power	53.2 VA	55.0 VA
Power Factor	0.71	0.71

The measurement without any connected drives represents the lowest power consumption that we could reach with the test system in a desktop environment with a power supply for desktop computers. A system with an

SSD drive would have a similar (~40W) power consumption, which can probably still be decreased by several Watts using shared power sources (like in blade systems) or a power supply dedicated for low-power computers. In fact, there are other tests of the same Atom processor and motherboard that report figures ranging between 33.5 – 40 Watts for idle power consumption and 41.5 – 45 Watts under full load [7][10].

3.2.3. *Power consumption of the Xeon (Harpertown) server*.

A very detailed investigation was conducted recently [2], focused on the power consumption of systems based on the Xeon 5400-series. Among several tests, measurements with similar conditions were performed also on the dual-socket system, which we use here for comparison:

	Xeon E5472
CPUs	2xE5472 @ 3GHz
Memory:	4*4GB 667MHz FB-DIMM
Power supply	800 Watts
Devices:	HDD, DVD

Measurement Results

	Idle	Cpuburn
Active Power	191 W	301 W
Apparent Power	208 VA	316 VA
Power Factor	0.92	0.95

We measured power consumption in two scenarios: with 1x4GB and 4x4GB memory. The approximate, averaged results for active power are listed below.

Load	**1x4 GB**	**4x4 GB**
1 process	186 W	210 W
2 processes	202 W	225 W
4 processes	232 W	255 W
8 processes	265 W	290 W

4. Software performance in 64-bit mode

4.1. Scalability with test40 (Geant4)

In our initial tests we used Geant4, the software package typically used in HEP for simulating the passing of particles through matter. The particular benchmark chosen was “*test40*”, a benchmark candidate once submitted for inclusion in SPEC 2006 [3]. The compilers used to generate the binaries were GCC 4.3 and ICC 11.0. GCC is the compiler used at CERN during procurement measurements and for compiling production software. While the Intel compiler significantly shortens execution time for this benchmark, relative results (i.e. Atom N330 vs. Xeon E5472) remain similar to the GCC case. The results are not listed here for space reasons, but can be found in Appendix B of our internal openlab report [11].

As we can see from Table 1a in Appendix B in that report, a single *test40* process on the Atom executes circa 5 times slower than on a Xeon E5472. This is not so surprising, since the Atom runs at approximately 50% of the clock speed and has *in-order* processing cores, as opposed to the *out-of-order* cores in the Xeon system. The most noteworthy point here is that *test40* seems to benefit a lot from Hyper Threading. Previous attempts to take advantage of this technology at CERN with the Pentium 4 Netburst family gave unimpressive results, with the performance increasing or decreasing by around 5% [8]. In the case of the Atom, however, we see up to 50% throughput increase on this particular test. Our explanation is that since the processing cores are *in-order*, hyper threading allows for much more aggressive runtime inter-thread optimization than on Netburst-based processors.

4.2. Multi-threaded Geant4 possibilities

CERN has recently started investing in multi-threading and multi-core research [4]. One of the areas of great interest is an experimental version of Geant4 designed to work in a multi-threaded environment. Developed by a Northeastern University PhD student, Xin Dong (who was an openlab summer student in 2008), the thread safe Geant4 prototype may provide the right step towards the efficient many-core based datacenter [5].

There are several reasons why a multi-threaded version of a prominent HEP software package is of importance. First of all, the multi-threaded model should allow more processor hardware to be used efficiently inside a smaller memory footprint. The aim would be, for instance, to add less than 50 MB per additional Geant4 thread. Secondly, a byproduct of this work might be an improvement of Geant4's modularity. Finally, the experimental software should serve as a very good example to the HEP community, not only by showing that even the most complex software frameworks working on embarrassingly parallel problems can be split into threads, but also by showing significant advantages of writing software that scales with the hardware.

Since most Geant4-based programs fit into 2 GB of memory, it would be an easy task to have a multi-threaded version prepared for Atom. Combined with the low cost of the hardware and low power requirements, the price/performance per watt advantage could be significant.

4.3. Scalability with the ALICE High Level Trigger (HLT) benchmark

Online processing constitutes an important part of the workload of a HEP experiment. One of the benchmarks representative of such jobs is the CBM/ALICE High Level Trigger (HLT/trackfitting) benchmark, originally developed by physicists at the University of Heidelberg [12], and enhanced by researchers at CERN openlab [6]. The code utilizes single-precision floating point with SSE and Intel Threading Building Blocks 2.0 for workload parallelization. The memory footprint is well below 2GB and does not vary a lot with the number of threads. The compiler used was GCC, version 4.2.4. The results of the tests, conducted in the same manner as in the case of test40, are also shown in Appendix B (Table 2) in the original report [11]. The workload remains constant, while the factor defining the changing throughput is the fit time spent on a single track in the benchmark.

Similarly to the case of *test40*, Xeon performance is 5-6 times better than that of Atom. In contrast to *test40*, however, this benchmark is somewhat heavier on floating point computation; therefore one of the explanations for slightly worse figures might be the lower efficiency of the floating point execution on Atom. Again, Hyper Threading is a significant benefit, yielding an additional 50% in throughput. However, it seems that the HLT program does not scale when moving from one to two threads. The reason for this behavior is the way the Intel Threading Building Blocks (TBB) scheduler operates, scheduling work on all available hardware threads when the number of available cores is specified as "3" or "4".

4.4. Memory issues

The most serious limitation of the Atom system is the 2GB memory size. With the proliferation of multi-core processors, CERN throughput computing has become more and more memory bound, and memory along with its power consumption has now become an important factor in both software development and hardware procurements. The minimal requirement of 2GB of memory per process has led to a situation in which in today's Core based systems FB-DIMM memory consumes easily as much electrical power as the processors.

The disadvantage of slower CPUs in throughput computing is the fact that the running processes occupy memory for a longer time, so more power is spent in the overall process. Later we show certain calculations which demonstrate that even when taking this fact into account, Atom based systems might be competitive. In addition, processes which are very demanding in terms of CPU power, but not so in terms of memory, benefit significantly from this kind of model – our calculations show an advantage of up to 25% in throughput per Swiss Franc, even without taking into account potential power supply, chipset and device optimizations.

5. Initial conclusions

We believe that our test of the Atom processor shows that it has potential, even though the measurements were conducted on an unoptimized consumer system. Although absolute performance is far behind the Xeon processor, the throughput result obtained with *our benchmarks* on a dual-core dual-threaded Atom system is encouraging. However, two items need to be rectified. They are the power consumption and in particular, the memory limitation.

As already mentioned a severe limitation is the fact that only 2 GB of memory is supported and has been certified. It should be possible to double this, since the processor does support 32bit physical addressing, but 8 GB seems currently to be impossible. We also need to state that such a large memory configuration, needed to run four standard HEP processes, would increase the system cost by 50%, as shown in section 2.

The power consumption is the second issue and the low requirements of the Atom processor are frustratingly dwarfed by the other components, especially the chipset. A “professional” server would, of course, be based on more efficient components and we discuss some possibilities in the next section. Currently the Atom is no match for the Xeon server in a power constrained environment.

In conclusion, the first Atom generation is interesting, but several issues need to be addressed before it can start to move closer to a HEP production environment.

6. Long term vision and possible future scenarios

In this section, we discuss some scenarios that could improve the attractiveness of the Atom processor. Based on the calculations in tables shown in the Appendix of the original report [11], we believe that a separate Atom-based system with a power consumption of over 25W is not competitive with today’s general purpose CPUs in terms of performance per Watt. However, simply lowering the power per CPU would create new possibilities. Also, based on our Web-prices, the throughput to cost ratio of our Atom system is quite impressive, surpassing that of the Xeon by up to 25%!

Let us consider a very basic scenario in which we have a hypothetical system equipped with Atom CPUs, satisfying the 2GB per process memory requirement of CERN’s applications. In the case of a dual core 2-way Hyper Threaded Atom N330, a total of 8GB of RAM are needed for the running jobs, and we assume that a future Atom system might be able to address this amount of memory. While the CPU itself has a TDP of 8W, the power consumed by the memory depends on its speed and type. The choices in this case are DDR2 and DDR3, the latter using around 30% less power. For 4GB DDR2 modules, we have seen power consumption figures in the range of 2 Watts per GB, giving a total of about 16-17 Watts for the 8GB needed. Thus, the grand total is already at 24W, dangerously close to the 25W limit. New hope comes with DDR3 memory, which is said to be 30% more power efficient than DDR2. If that were the case, the total power consumption of the memory and CPU would be at around 19-20 Watts, leaving some 5 Watts of room for the chipset and other hardware. But it must not be forgotten that the memory and power comparison will change after we add Intel’s next generation Xeon processor, the 5500-series, and its support for DDR3 into the overall equation.

As all CPU designs evolve, it is reasonable to assume that throughput figures of the Atom would get an additional benefit from the introduction of a 4-core part, with the same 8W thermal envelope as the current N330 model. The described gain would be a 30% throughput increase compared to the N330, provided that a limit of one process per core is being kept. Given the history of chipmakers' innovation, such a scenario is not completely unrealistic. With a 4-core part with an 8W TDP and comparable performance, the competitiveness power limit would be raised to 33W, and 13W out of this could be spent to power circuits separate from the CPU and memory. Going back to the memory consumption optimization issue in software, let us consider two examples related to exploiting Hyper Threading on a 4-core part:

- a 20 to 40% reduction in memory consumption would yield about 12%-25% more throughput within the same power envelope
- a 50% reduction in memory consumption would yield 50% more throughput within the same power envelope

Another take on the memory problem might be software optimization. Although the current requirement is 2GB per process, it might be possible to have multi-threaded software such as Geant4 as a unique user process exploiting several software threads inside a 2GB memory. In such a case, the reduction in power consumption would be advantageous, both for a smaller Atom-based server and for a large Xeon-based one. We can easily imagine that a single process would be allowed to take all CPUs in an Atom system. If the same thing would be allowed (and configured) in a larger Xeon server it would almost be a revolution and have a broad impact on the way HEP optimizes its throughput processing.

Pooling of resources, be it I/O control, auxiliary circuits or even cooling apparatus, usually reduces overall power consumption of the system. By combining today's blade or microblade technology with the power efficiency of the Atom, one could easily design a product very suitable for throughput computing. A blade or microblade system with multi-core Atom boards would easily achieve very good power consumption figures, while delivering high computational performance, and be ideal for many of CERN's computing tasks. Given that blade systems are powered by highly efficient power supplies and cooled by a collective cooling system, energy savings in this domain could be substantial. As with other many-CPU and many-core systems, inter-core communication could become a serious bottleneck. In the case of multi-CPU boards, however, the communications overhead could be reduced and perhaps even solved with today's existing backplane solutions.

Compiler optimization should also be pushed on in-order systems, such as the Atom. *In-order* processing relies on excellent compiler back-ends, which, contrary to *out-of-order* architectures, must account for each processor cycle and plan in advance the placement of each instruction with precision. Today, ICC does a much better scaling job on the Xeon, presumably because it gets assistance from the *out-of-order* microarchitecture. As far as we know, there are some optimization opportunities for the Atom implemented currently in ICC, enabled with the *–xsse3_atom* flag. The results we obtained show, surprisingly, a very slight decrease in performance when using this switch.

Other possibilities could include models similar to the SiCorTex model[13] and expansion card scenarios, however the former is difficult to develop and limited in the scope of applications, and the latter is already being taken care of by the engineers working on the Larrabee chip [14].

One final fact is that new processor generations will come on both the Xeon and Atom side. The Xeon 5500-series CPUs have just been released in a server version [15] and have definitely brought significant improvements to the microarchitecture, most notably a greatly expanded reorder buffer in the *out-of-order* engine. But, even more importantly, it features Hyper Threading, just like the Atom processor, affecting results significantly. It therefore seems appropriate to rerun such a comparison with the "year 2009" flavours of the Atom and the Xeon processors.

7. Development directions

The following points are seen as interesting for future development in the domain of the Atom family:

- Support of at least 4GB of energy-optimized memory
- Mainboard and chipset power reduction (i.e. through removing unneeded components)
- Efficient (low power) power supplies and/or pooled power supplies
- Microblade technology
- Compiler optimizations for:
    - in-order execution
    - low power operation
- Hyper Threading comparison of future Xeon systems and future Atom systems
- Increased core density processors

References

[1] P Mato Vila – “Multi-Processors and memory Requirements”, July 2008
[2] G Balazs – “Power measurements of server computers”, CERN 2008 (unpublished)
[3] http://www.spec.org/
[4] X Dong – “Multi-threaded Geant4 with Shared Detector”, CERN 2008 (unpublished)
[5] X Dong et al – “First Results in a Thread-Parallel Geant4”, Northeastern Uni., Boston, USA, 2008
[6] H Bjerke – Review of a parallel High Level Trigger benchmark (using multithreading and/or SSE), CERN 2008 (https://edms.cern.ch/file/893336/1/HB_presentation.pdf)
[7] E Kucharik – “Dual-core Atom 330 tested”, Fudzilla 2008 (http://www.fudzilla.com/index.php?option=com_content&task=view&id=9775&Itemid=40)
[8] Private communication, H Meinhard (CERN/IT/FIO)
[9] A Hirstius, S Jarp, A Nowak – “Strategies for increasing data centre power efficiency”, CERN, 2008
[10] P Schmid, A Roos – “Efficiency: Core 2 Nukes Atom On The Desktop”, Tom’s Hardware 2008 (http://www.tomshardware.com/reviews/intel-atom-efficiency,2069-11.html)
[11] G Balasz, S Jarp, A Nowak – “Is the Atom processor ready for High Energy Physics”, CERN openlab Internal Report, Nov. 2008 (http://openlab-mu-internal.web.cern.ch/openlab-mu-internal/Documents/ 2_Technical_Documents/Technical_Reports/2008/CERN Atom 330 analysis.pdf)
[12] I.Kisel/GSI: “Fast SIMDized Kalman filter based track fit” (http://www-linux.gsi.de/~ikisel/reco/CBM/ DOC-2007-Mar-127-1.pdf)
[13] http://www.sicortex.com/
[14] http://www.intel.com/technology/visual/microarch.htm
[15] http://www.intel.com/pressroom/archive/releases/20090330corp_sm.htm

Appendix A. Measuring values related to power

Active Power (P): The component of electric power that performs work, often referred to as 'real' power. That part of the electric power can be used by the actual power consumers. It is measured in Watts (W).
Apparent Power (S): The product of the voltage (in volts) and the current (in amperes). This part of the power represents what is being drawn from the electrical circuit. It is measured in volt-amperes (VA).
Power Factor: The ratio between the Active power and the Apparent power. The power factor is a number between 0 and 1 representing the power efficiency of the power supply in our case.

Since the overall electric power consumption is always charged based on the Apparent Power that is measured at the end of the subscriber’s power line, the results for the Apparent Power consumption are used in CERNs tendering process. For deeper analysis and calculations to examine the power consumption of the different devices in the computers the Active Power results are used.