# Is the Atom processor ready for High Energy Physics?

## An initial analysis of the dual core Atom N330 processor

***Gyorgy Balazs, Sverre Jarp, Andrzej Nowak (CERN openlab)***

***CERN, November 2008***

# *Executive Summary*

This paper compares an amateur-built Atom N330 single-socket server with a standard dual-socket Xeon Core 2 Quad server to gauge the potential of the Atom system for running High Energy Physics applications. The Xeon server corresponds roughly to what is bought in standard acquisitions for the Large Hadron Collider Computing Grid, however the test scenario includes using memory-optimized software.

Two benchmarks are being used; one simulation benchmark from offline and one trigger benchmark from online. On the 1.6GHz Atom server the offline benchmark, which is a Monte-Carlo simulation program, obtains 1/5 of the throughput performance of the 3.0 GHz Xeon server. When running four parallel processes on the Atom server and eight on the Xeon, we get a throughput ratio of 13.3 (5*8/3) in favor of the Xeon. This may sound unimpressive, but by comparing Web-based pricing, we find that the Atom server is at least 20 times cheaper. Of course, this pricing comparison is very approximate, but this is nevertheless seen as an encouragement.

The power consumption is less encouraging, since we find a ratio of “only” 5 in favor of the Atom system. This is mainly due to the high consumption of the Atom chipset. The Atom system is therefore currently no match for the Harpertown in a power constrained environment.

Another big worry is the fact that the Atom system is only certified with a 2GB memory size. We believe that this could be increased to 4GB since the Atom processor supports 32bit physical addressing, but some (not all) of the LHC software frameworks currently need 8GB when running four processes in parallel. The online benchmark runs a bit more slowly that the offline one, but could probably fit four processes into the available memory.

Our vision is a rack full of Atom blades, but before reaching such a stage, more work is needed by both sides. We must continue to encourage the LHC experiments to reduce the memory footprint of their offline applications to 1GB or less – a target which should be reachable by at least two of the four experiments. There is work underway to

demonstrate that memory can also be saved by converting the applications to use multithreading. Intel would have to improve on at least two fronts; the memory limit must be lifted to 4GB (or more) and the Atom chipset must be made more power efficient.

In the final section we suggest ways to improve even further: improved compiler optimization, additional cores and so on.

# 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 Harpertown-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

## Hardware setup

The sample kit received from Intel contained a D945GCLF2 mini-ITX desktop board equipped with an Intel 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.

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/1 slot 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.

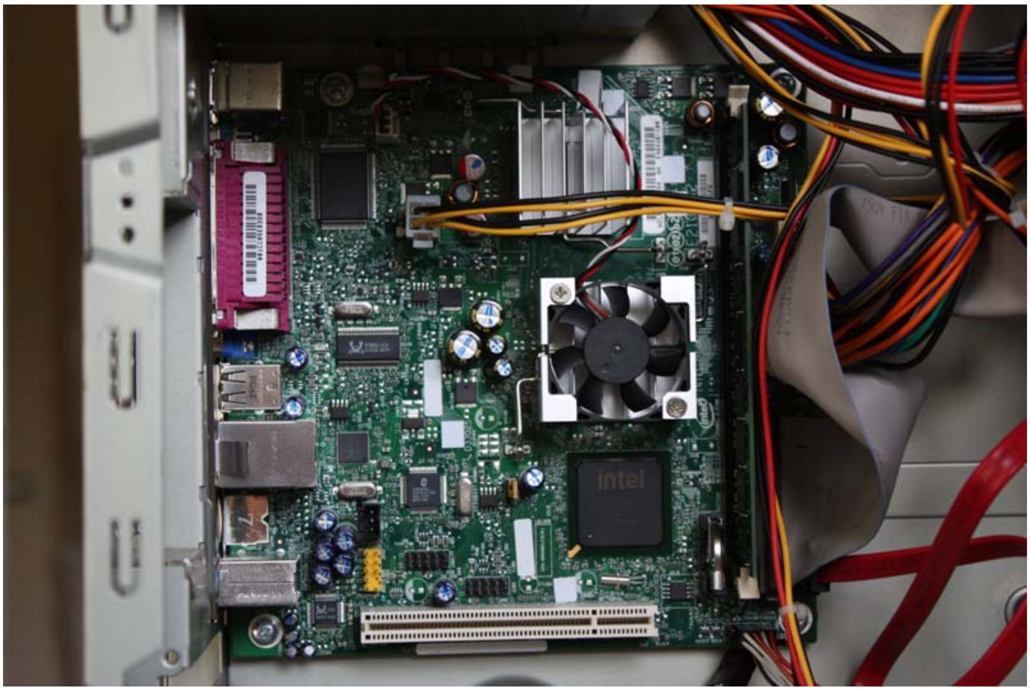

***Figure 1: The Atom-based motherboard***

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 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 |

## 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.

## Pricing information

In this paper the performance of the test system is being compared to a dual socket server with 3GHz 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" Harpertown based 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 Harpertown 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 <u>current</u> CERN memory requirements. It is compared with a Harpertown 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** |

| Harpertown 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 the 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.

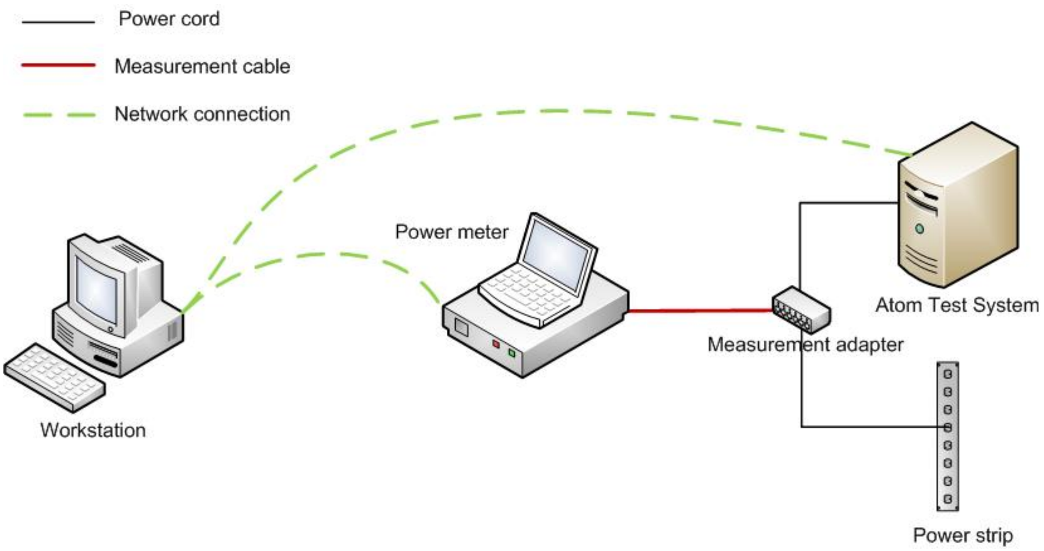

***Figure 2: Power test setup***

---

[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 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 idle state.

- **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.

## Results:

Two sets of tests were conducted with different hardware configurations:

### Tests with the initial 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.

### Test to reach the lowest power consumption

Another measurement was performed with the HDD and 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 CPU and motherboard, which report figures ranging between 33.5-40 Watts for idle power consumption and 41.5-45 Watts under full load [7][10].

**Power consumption of a Harpertown based system**

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

**Harpertown 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 have measured power consumption in two scenarios: with 1x4GB and 4x4GB memory. The approximate, averaged results for active power are below. We use these numbers for our calculations presented in Appendix B.

| **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]*

## 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 the Intel compiler 11 (ICC 11.0, release 069). 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 vs. Harpertown) remain similar to the GCC case. The following table contains a throughput comparison of the dual socket quad-core Harpertown system at 3GHz with our test Atom board. 100% of throughput is defined as a single thread running on Atom. The power measured should be considered approximate (+/- 5%).

As we can see from Table 1a in Appendix B, a single *test40* process on the Atom executes circa 5 times slower than on a Harpertown. 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 Harpertown 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.

## 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 this year), 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.

## Scalability with the Alice High Level Trigger 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, developed by physicists at the University of Heidelberg, 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 shown in Appendix B. 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*, Harpertown performance is 5-6 times better than that of Atom, as one can see in Table 3 in Appendix B. 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 TBB scheduler operates, scheduling work on all available hardware threads when the number of available cores is specified as "3" or "4".

**The memory issues**

The most important 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 show 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. Two items need to be rectified, however. 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 is currently 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 problem and the low needs of the Atom processor are frustratingly dwarfed by the other components, including 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 Harpertown in a power constrained environment.

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

# *6 Long term vision and possible 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, 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 Harpertown 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, "Nehalem", and its support for DDR3 into the overall equation.

As all CPU designs evolve, it is reasonable to throughput figures of the Atom would get an additional benefit from introducing a 4-core part, with the same 8W TDP 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 Harpertown, presumably because it gets assistance from the out-of-order microarchitecture. As far as we know, there are few optimization opportunities for the Atom implemented currently in ICC, enabled with the –xL/sse3_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 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.

One final fact is that new processor generations will come on both the Xeon and Atom side on a frequent basis. The Nehalem family, due to be publicly available in the near future, will bring certain improvements to the microarchitecture, most notably the out-of-order engine will benefit from a greatly expanded reorder buffer. More importantly, however, it will feature Hyper Threading, which could profoundly affect the scalability results of tests.

# *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] Pere Mato Vila – “Multi-Processors and memory Requirements”, July 2008

[2] Gyorgy Balazs – “Power measurements of server computers”, CERN 2008 (unpublished)

[3] http://www.spec.org

[4] Xin Dong – “Multi-threaded Geant4 with Shared Detector”, CERN 2008 (unpublished)

[5] Xin Dong, Gene Cooperman – “First Results in a Thread-Parallel Geant4”, Northeastern University, Boston, USA, 2008

[6] Håvard 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] Eliot 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] Dr Andreas Hirstius, Sverre Jarp, Andrzej Nowak – “Strategies for increasing data centre power efficiency”, CERN, 2008

[10] Patrick Schmid, Achim Roos – “Efficiency: Core 2 Nukes Atom On The Desktop”, Tom’s Hardware 2008 (http://www.tomshardware.com/reviews/intel-atom-efficiency,2069-11.html)

# *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.

# *Appendix B – Data tables*

Lifetime power consumption is calculated based on a standard 3 year lifetime period over which each Watt consumed costs about 2.5 CHF. ***Note: The “Throughput in 2.5 MW” figure cannot be compared across scenarios.***

### Scenario 1: Atom 330 (2GB) vs. Xeon E5472 (4GB)

| | SETUP | USER TIME | | ACTIVE | ADVANTAGE | | |
|---|---|---|---|---|---|---|---|
| | #proc | Runtime AVG (us) | % of 1 proc | POWER (W) | Workload | Throughput | Throughput per Watt |
| **ATOM 330 @ 1.6 GHz Fedora 9, GCC 4.3, 2GB RAM** | 1 | 156 | 100% | 47 W | 100% | **100%** | **100%** |
| | 2 | 157 | 100% | 48 W | 200% | 199% | 195% |
| | 3 | 192 | 123% | 49 W | 300% | 244% | 234% |
| | 4 | 207 | 132% | 50 W | 400% | 302% | 287% |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.3, 4GB RAM** | 1 | 32 | 21% | 186 W | 100% | 488% | 123% |
| | 2 | 32 | 21% | 202 W | 200% | 973% | 227% |
| | 4 | 32 | 21% | 232 W | 400% | 1944% | 394% |
| | 8 | 32 | 21% | 265 W | 800% | 3891% | 690% |

*Table 1a: “test40” Harpertown vs. Atom comparison*

| | SETUP | ACTIVE | ADVANTAGE | COST | | Throughput index | Reality |
|---|---|---|---|---|---|---|---|
| | #proc | POWER (W) | Throughput per Watt | System cost / unit | Power cost | TP / CHF | TP in 2.5 MW |
| **ATOM 330 @ 1.6 GHz Fedora 9, GCC 4.3, 2GB RAM** | 1 | 47 W | **100%** | CHF 250 | CHF 118 | **272 (100%)** | 5,319,149 % |
| | 2 | 48 W | 195% | CHF 250 | CHF 120 | 538 (198%) | 10,383,387 % |
| | 3 | 49 W | 234% | CHF 250 | CHF 123 | 654 (240%) | 12,436,224 % |
| | 4 | 50 W | 287% | CHF 250 | CHF 124 | 808 (297%) | **15,261,575 %** |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.3, 4GB RAM** | 1 | 186 W | 123% | CHF 5,200 | CHF 465 | 86 (31%) | 6,552,419 % |
| | 2 | 202 W | 227% | CHF 5,200 | CHF 505 | 170 (62%) | 12,048,007 % |
| | 4 | 232 W | 394% | CHF 5,200 | CHF 580 | 336 (123%) | 20,947,470 % |
| | 8 | 265 W | 690% | CHF 5,200 | CHF 663 | 663 (243%) | **36,706,422 %** |

*Table 1b: “test40” Harpertown vs. Atom comparison, power and cost*

| | SETUP | USER TIME | | ACTIVE | ADVANTAGE | | |
|---|---|---|---|---|---|---|---|
| | #proc | Runtime AVG (us) | % of 1 proc | POWER (W) | Workload | Throughput | Throughput per Watt |
| **ATOM 330 @ 1.6 GHz Fedora 9, GCC 4.2.4, 2GB RAM** | 1 | 2.595 | **100%** | 47 W | 100% | **100%** | **100%** |
| | 2 | 1.295 | 50% | 48 W | 100% | 200% | 196% |
| | 3 | 0.83125 | 32% | 49 W | 100% | 312% | 299% |
| | 4 | 0.835 | 32% | 50 W | 100% | 311% | 295% |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.2.4, 4GB RAM** | 1 | 0.470 | 18% | 186 W | 100% | 552% | 140% |
| | 2 | 0.230 | 9% | 202 W | 100% | 1128% | 263% |
| | 4 | 0.115 | 4% | 232 W | 100% | 2257% | 457% |
| | 8 | 0.062 | 2% | 265 W | 100% | 4185% | 742% |

***Table 1c: "tbb" Harpertown vs. Atom comparison***

| | SETUP | ACTIVE | ADVANTAGE | COST | | Throughput index | Reality |
|---|---|---|---|---|---|---|---|
| | #proc | POWER (W) | Throughput per Watt | System cost / unit | Power cost | TP / CHF | TP in 2.5 MW |
| **ATOM 330 @ 1.6 GHz Fedora 9, GCC 4.2.4, 2GB RAM** | 1 | 47 W | **100%** | CHF 250 | CHF 118 | **272 (100%)** | 5,319,149 % |
| | 2 | 48 W | 196% | CHF 250 | CHF 120 | 541 (199%) | 10,436,776 % |
| | 3 | 49 W | 299% | CHF 250 | CHF 123 | 838 (307%) | 15,927,574 % |
| | 4 | 50 W | 295% | CHF 250 | CHF 124 | 831 (305%) | **15,695,881 %** |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.2.4, 4GB RAM** | 1 | 186 W | 140% | CHF 5,200 | CHF 465 | 97 (35%) | 7,421,071 % |
| | 2 | 202 W | 263% | CHF 5,200 | CHF 505 | 197 (72%) | 13,963,625 % |
| | 4 | 232 W | 457% | CHF 5,200 | CHF 580 | 390 (143%) | 24,315,967 % |
| | 8 | 265 W | 742% | CHF 5,200 | CHF 663 | 713 (262%) | **39,485,697 %** |

***Table 1d: "tbb" Harpertown vs. Atom comparison, power and cost***

## Scenario 2: Atom 330 (2x4GB) vs. Xeon E5472 (4x4GB)

| | SETUP | USER TIME | | ACTIVE | ADVANTAGE | | |
|---|---|---|---|---|---|---|---|
| | #proc | Runtime AVG (us) | % of 1 proc | POWER (W) | Workload | Throughput | Throughput per Watt |
| **Atom X @ 1.6 GHz Fedora 9, GCC 4.3, 2x4GB RAM** | 1 | 156 | **100%** | 53 W | 100% | **100%** | **100%** |
| | 2 | 157 | 100% | 54 W | 200% | 199% | 196% |
| | 3 | 192 | 123% | 55 W | 300% | 244% | 235% |
| | 4 | 207 | 132% | 56 W | 400% | 302% | 286% |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.3, 4x4GB RAM** | 1 | 32 | 21% | 210 W | 100% | 488% | 123% |
| | 2 | 32 | 21% | 225 W | 200% | 973% | 229% |
| | 4 | 32 | 21% | 255 W | 400% | 1944% | 404% |
| | 8 | 32 | 21% | 290 W | 800% | 3891% | 711% |

*Table 2a: "test40" Harpertown vs. Atom comparison*

| | SETUP | ACTIVE | ADVANTAGE | COST | | Throughput index | Reality |
|---|---|---|---|---|---|---|---|
| | #proc | POWER (W) | Throughput per Watt | System cost / unit | Power cost | TP / CHF | TP in 2.5 MW |
| **Atom X @ 1.6 GHz Fedora 9, GCC 4.3, 2x4GB RAM** | 1 | 53 W | **100%** | CHF 370 | CHF 133 | **199 (100%)** | 4,716,981 % |
| | 2 | 54 W | 196% | CHF 370 | CHF 135 | 394 (198%) | 9,229,677 % |
| | 3 | 55 W | 235% | CHF 370 | CHF 138 | 480 (241%) | 11,079,545 % |
| | 4 | 56 W | 286% | CHF 370 | CHF 140 | 592 (297%) | **13,490,142 %** |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.3, 4x4GB RAM** | 1 | 210 W | 123% | CHF 6,100 | CHF 525 | 73 (36%) | 5,803,571 % |
| | 2 | 225 W | 229% | CHF 6,100 | CHF 563 | 146 (73%) | 10,816,433 % |
| | 4 | 255 W | 404% | CHF 6,100 | CHF 638 | 288 (144%) | 19,058,091 % |
| | 8 | 290 W | 711% | CHF 6,100 | CHF 725 | 570 (286%) | **33,542,075 %** |

*Table 2b: "test40" Harpertown vs. Atom comparison, power and cost*

|  | SETUP | USER TIME |  | ACTIVE | ADVANTAGE |  |  |
|---|---|---|---|---|---|---|---|
|  | #proc | Runtime AVG (us) | % of 1 proc | POWER (W) | Workload | Throughput | Throughput per Watt |
| **Atom X @ 1.6 GHz Fedora 9, GCC 4.2.4, 2x4GB RAM** | 1 | 2.595 | **100%** | 53 W | 100% | **100%** | **100%** |
|  | 2 | 1.295 | 50% | 54 W | 100% | 200% | 197% |
|  | 3 | 0.83125 | 32% | 55 W | 100% | 312% | 301% |
|  | 4 | 0.835 | 32% | 56 W | 100% | 311% | 294% |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.2.4, 4x4GB RAM** | 1 | 0.470 | 18% | 210 W | 100% | 552% | 139% |
|  | 2 | 0.230 | 9% | 225 W | 100% | 1128% | 266% |
|  | 4 | 0.115 | 4% | 255 W | 100% | 2257% | 469% |
|  | 8 | 0.062 | 2% | 290 W | 100% | 4185% | 765% |

*Table 2c: "tbb" Harpertown vs. Atom comparison*

|  | SETUP | ACTIVE | ADVANTAGE | COST |  | Throughput index | Reality |
|---|---|---|---|---|---|---|---|
|  | #proc | POWER (W) | Throughput per Watt | System cost / unit | Power cost | TP / CHF | TP in 2.5 MW |
| **Atom X @ 1.6 GHz Fedora 9, GCC 4.2.4, 2x4GB RAM** | 1 | 53 W | **100%** | CHF 370 | CHF 133 | **199 (100%)** | 4,716,981 % |
|  | 2 | 54 W | 197% | CHF 370 | CHF 135 | 396 (199%) | 9,277,134 % |
|  | 3 | 55 W | 301% | CHF 370 | CHF 138 | 615 (309%) | 14,190,021 % |
|  | 4 | 56 W | 294% | CHF 370 | CHF 140 | 609 (306%) | **13,874,038 %** |
| **Harpertown @ 3.0 GHz SLC 4.7, GCC 4.2.4, 4x4GB RAM** | 1 | 210 W | 139% | CHF 6,100 | CHF 525 | 83 (41%) | 6,572,948 % |
|  | 2 | 225 W | 266% | CHF 6,100 | CHF 563 | 169 (85%) | 12,536,232 % |
|  | 4 | 255 W | 469% | CHF 6,100 | CHF 638 | 334 (168%) | 22,122,762 % |
|  | 8 | 290 W | 765% | CHF 6,100 | CHF 725 | 613 (308%) | **36,081,758 %** |

*Table 2d: "tbb" Harpertown vs. Atom comparison, power and cost*