

Energy Efficiency, Performance and Scalability of XML Parsing Using SIMD

ABSTRACT

XML is a data format designed for documents as well as the representation of data structures. The simplicity and generality of the rules make it widely used in web services and database systems. Traditional XML parsers have been built around the byte-at-a-time model, in which they process every character token in the file in a sequential fashion. Unfortunately, the byte-at-time sequential model is a performance barrier in demanding applications, and is also energy-inefficient, making poor use of the wide registers and other parallelism features in modern processors.

This paper assesses the energy and performance of a new approach to XML parsing based on parallel bit stream technology. This method first converts the character streams into sets of parallel bitstreams and then exploits SIMD operations prevalent on modern CPUs. Our first generation Parabix1 parser then uses bitscan instructions over these streams to make multibyte moves in an otherwise sequential approach. Our second generation Parabix2 technology further parallelizes our parsers by replacing much of the sequential bit scanning with a parallel scanning approach based on bitstream addition. We evaluate Parabix1 and Parabix2 against two widely-used XML parsers, James Clark's Expat and Apache's Xerces-C on three generations of x86 machines, including the new Intel Sandy Bridge. We show that Parabix2's speedup is $2\times$ – $8\times$ over Expat and Xerces. In stark contrast to the energy expenditures necessary to realize performance gains through multicore parallelism, we also show that our Parabix parsers deliver energy savings directly in proportion to performance gains. We also assess the scalability advantages of SIMD processor improvements the different Intel machine generations, culminating with an evaluation of the 256-bit AVX technology in SandyBridge vs. the now legacy 128-bit SSE technology.

1. Introduction

Extensible Markup Language (XML) is a core technology standard of the World-Wide Web Consortium (W3C) that provides a common framework for encoding and communicating structured information of all kinds. In applications ranging from Office Open XML in Microsoft Office to NDFD XML of the NOAA National Weather Service, from KML in Google Earth to Castor XML in the Martian Rovers, from ebXML for e-commerce data interchange to RSS for news feeds from web sites everywhere, XML plays a ubiquitous role in providing a common framework for data interoperability world-wide and beyond. As XML 1.0 editor Tim Bray is quoted in the W3C celebration of XML at 10 years, "there is essentially no computer in the world, desk-top, hand-held, or back-room, that doesn't process XML sometimes."

With all this XML processing, a substantial literature has arisen addressing XML processing performance in general and the performance of XML parsers in particular. Nicola and John specifically identified XML parsing as a threat to database performance and outlined a number of potential directions for potential performance improvements [19]. The nature of XML APIs was found to have a significant affect on performance with event-based SAX (Simple API for XML) parsers avoiding the tree construction costs of the more flexible DOM (Document Object Model) parsers [20]. The

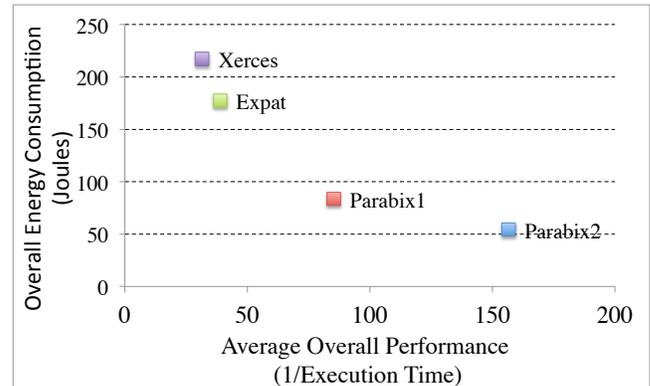


Figure 1: Energy vs. Performance for Four XML Parsers

commercial importance of XML parsing spurred developments of hardware-based approaches including the development of a custom XML chip [17] as well as FPGA-based implementations [12]. As promising as these approaches may be for particular niche applications, however, it is still likely that the bulk of the world's XML processing workload will be carried out on commodity processors using software-based solutions.

To accelerate XML parsing performance in software, most recent work has focused on parallelization. The use of multicore parallelism for chip multiprocessors has attracted the attention of several groups [18, 21, 22], while SIMD (single-instruction multiple data) parallelism has been of interest to Intel in designing new SIMD instructions [16] as well as to our group in developing parallel bit stream technology [5, 7, 8]. Each of these approaches has shown considerable performance benefits over traditional sequential parsing following the byte-at-a-time model.

With a focus on performance, however, relatively less attention has been paid to reducing energy consumption. For example, in addressing performance through multicore parallelism, one generally has to pay an energy price for performance gains because of the increased processing required for synchronization. A focus on reduction of energy consumption is a key topic in this paper, in which we study the energy and performance characteristics of several XML parsers across three generations of x86-64 processor technology. The parsers we consider are the widely used byte-at-a-time parsers Expat and Xerces as well as our own Parabix1 and Parabix2 parsers. A compelling result is that the performance benefits of parallel bit stream technology translate directly and proportionately to substantial energy savings. Figure 1 is an energy-performance scatter plot showing the results we obtain for the four parsers.

The remainder of this paper is organized as follows. Section 2 presents background material on XML parsing and traditional parsing methods. Section 3 then reviews parallel bit stream technology as applied to XML parsing in our Parabix1 and Parabix2 parsers. Section 4 then introduces our methodology and approach for the performance and energy study tackled in the remainder of the paper. Section 5 presents a detailed performance evaluation on a Core i3 processor as our primary evaluation platform, addressing a number of microarchitectural issues including cache misses, branch mispredictions, SIMD instruction counts and so on. Sec-

tion 6 then looks at scalability and performance gains through three generations of Intel architecture culminating with performance assessment on our one week-old Sandy Bridge test machine. Section 7 looks specifically at issues in applying the new 256-bit AVX technology to parallel bit stream technology and notes that the major performance benefit seen so far is a result of the change to 3-operand instruction form. Section 8 concludes the paper with a discussion of ongoing work and further research directions.

2. Background

This section provides a brief overview of XML and traditional and parallel XML processing technology. Section 3 describes the key design and performance aspects of both generations of the Parabix parallel XML processing technology.

2.1 XML

In 1998, the W3C officially adopted XML as a standard. XML is a platform-independent data interchange format. The defining characteristics of XML are that it can represent virtually any type of information through the use of self-describing markup tags and can easily store semi-structured data in a descriptive fashion. XML markup encodes a description of an XML document’s storage layout and logical structure. Because XML was intended to be human-readable, XML markup tags are often verbose by design [4]. For example, a typical XML file could be:

```
<?xml version="1.0"?>
<Products>
  <Product ID="0001">
    <ProductName Language="English">Widget</ProductName>
    <ProductName Language="French">Bitoniau</ProductName>
    <Company>ABC</Company>
    <Price>$19.95</Price>
  </Product>
  ...
</Products>
```

Figure 2: Simple XML Document

XML files can be classified as “document-oriented” or “data-oriented” [13]. Document-oriented XML is designed for human readability, such as shown in Figure 2; data-oriented XML files are intended to be parsed by machines and omit any “human-friendly” formatting techniques, such as the use of whitespace and descriptive “natural language” naming schemes. Although the XML specification itself does not distinguish between “XML for documents” and “XML for data” [4], the latter often requires the use of an XML parser to extract the information within. The role of an XML parser is to transform the text-based XML data into an application-ready format.

2.2 Traditional XML Parsers

Traditional XML parsers process XML sequentially a single byte-at-a-time. Following this approach, an XML parser processes a source document serially, from the first to the last byte of the source file. Each character of the source text is examined in turn to distinguish between the XML-specific markup, such as an opening angle bracket ‘<’, and the content held within the document. The current character that the parser is processing is referred to as its cursor position. As the parser moves the cursor through the source document, the parser alternates between markup scanning, and data validation and processing operations. At each processing step, the parser scans the source document and either locates the expected markup, or reports an error condition and terminates. In other words, traditional XML parsers are complex finite-state machines that use byte comparisons to transition between data and metadata states. Each state transition indicates the context in which to interpret the subsequent characters. Unfortunately, textual data tends

to consist of variable-length items in generally unpredictable patterns [7]; thus any character could be a state transition until deemed otherwise.

Expat and Xerces-C are popular byte-at-a-time sequential parsers. Both are C/C++ based and open-source. Expat was originally released in 1998; it is currently used in Mozilla Firefox and Open Office [10]. Xerces-C was released in 1999 and is the foundation of the Apache XML project [15]. For example, the main loop of Xerces-C well-formedness scanner contains:

```
XXXXXXXXXX  XERCES CODE  XXXXXXXXXXXX
```

The major disadvantage of the byte-at-a-time sequential approach to XML parsing is that each character incurs at least one conditional branch. The cumulative effect of branch mispredictions penalties are known to degrade parsing performance in proportion to the markup density of the source document [8] (i.e., the proportion of XML-markup vs. XML-data).

2.3 Parallel XML Parsing

In general, parallel XML acceleration methods comes in one of two forms: multithreaded approaches and SIMD-based techniques. Multithreaded XML parsers take advantage of multiple cores via number of strategies. Approaches include preparing the XML file to locate key partitioning points [22] and speculative P-DFAs [22]. Once divided, the XML workload is processed independently across the available cores. SIMD XML parsers leverage the SIMD registers to overcome the performance limitations of the byte-at-a-time sequential processing paradigm as well as inherent data dependent branch misprediction rates [7]. SIMD instructions allow the processor to perform the same operation on multiple pieces of data simultaneously. To our knowledge, the only SIMD-based XML parsers are Parabix1 and Parabix2, both of which were designed and developed by Cameron et al. [8]. We discuss both versions of Parabix in Section 3.

2.4 SIMD Operations

3. Parabix

3.1 Parabix1

At a high level, Parabix1 processes source XML in a functionally equivalent manner as a traditional processor. That is, Parabix1 moves sequentially through the source document, maintaining a single cursor position throughout the parsing process. Where Parabix1 differs from the traditional parser is that it scans for key markup characters using a series of basis bitstreams. A bitstream is simply a sequence of 0s and 1s, where there is one such bit in the bitstream for each character in a source data stream. A basis bitstream is a bitstream that consists of only transposed textual XML data. In other words, a source character consisting of M bits can be represented with M bitstreams and by utilizing M SIMD registers of width W , it is possible to scan through W characters in parallel. The register width W varies between 64-bit for MMX, 128-bit for SSE, and 256-bit for AVX. Figure 3 presents an example of how we represent 8-bit ASCII characters using eight bitstreams. $B_0 \dots B_7$ are the individual bitstreams. The 0 bits in the bitstreams are represented by periods, so that the 1 bits stand out.

In order represent the byte-oriented character data as parallel bitstreams, the source data is first loaded in sequential order and converted into its transposed representation through a series of packs, shifts, and bitwise operations. Using the SIMD capabilities of current commodity processors, this transposition of source data to bitstreams incurs an amortized overhead of about 1 CPU cycle per byte for transposition [8]. When parsing, we need to consider multiple properties of characters at different stages during the

```

source data ▷ <t1>abc</t1><t2/>
B0          ..1.1.1.1.1.1...1.
B1          ...1.11.1.1..1..111
B2          11.1...111.111.11
B3          1..1...11..11..11
B4          1111...1.111111.1
B5          111111111111111111
B6          .1..111..1...1...
B7          .....

```

Figure 3: Parallel Bitstream Example

process. Using the basis bitstreams, it is possible to combine them using bitwise logic in order to compute character-class bitstreams; that is, streams that identify the positions at which characters belonging to a specific character class occur. For example, a ASCII character is an open angle bracket ‘<’ if and only if $B_2 \wedge \dots \wedge B_5 = 1$ and the other basis bitstreams are 0 at the same position within the basis bitstreams. Once these character-class bitstreams are created, bit-scan operations, common to commodity processors, can be used for sequential markup scanning and data validation operations. A common operation in all XML parsers is identifying the start tags (<) and their accompanying end tags (either “/>” or “>” depending whether the element tag is an empty element tag or not, respectively).

```

source data ▷ <t1>abc<t1/><t2/>
M0 = 1      1.....
M1 = advance(M0) .1.....
M2 = bitscan('>') ...1.....
M3 = advance(M2) ....1.....
M4 = bitscan('<') .....1.....
M5 = bitscan('/') .....1.....
M6 = advance(M5) .....1.....
M7 = bitscan('<') .....1...
M8 = bitscan('/') .....1.
M9 = advance(M8) .....1

```

Figure 4: Parabix1 Start and End Tag Identification

Unlike traditional parsers, these sequential operations are accelerated significantly since bit scan operations can perform up to W finite state transitions per clock cycle. This approach has recently been applied to Unicode transcoding and XML parsing to good effect, with research prototypes showing substantial speed-ups over even the best of byte-at-a-time alternatives [7–9].

3.2 Parabix2

In Parabix2, we replace the sequential single-cursor parsing using bit scan instructions with a parallel parsing method using bitstream addition. Unlike the single-cursor approach of Parabix1 (and conceptually of all sequential XML parsers), Parabix2 processes multiple cursors in parallel. For example, using the source data from Figure 4, Figure 5 shows how Parabix2 identifies and moves each of the start tag markers forwards to the corresponding end tag. Like Parabix1, we assume that N (the name chars) has been computed using the basis bitstreams and that

In general, the set of bit positions in a marker bitstream may be considered to be the current parsing positions of multiple parses taking place in parallel throughout the source data stream. A further aspect of the parallel method is that conditional branch statements used to identify syntax error at each each parsing position are eliminated. Although we do not show it in the prior examples,

```

source data ▷ <t1>abc<t1/><t2/>
N = name chars .11.111.11...11..
M0 = [<]       1.....1.....
M1 = advance(M0) .1.....1.....
M2 = scant0('/', '>') ...1.....1.....
M3 = scant0('>')   ...1.....1.....

```

Figure 5: Parabix2 Start and End Tag Identification

error bitstreams can be used to identify any well-formedness errors found during the parsing process. Error positions are gathered and processed in as a final post processing step. Hence, Parabix2 offers additional parallelism over Parabix1 in the form of multiple cursor parsing as well as significantly reduces branch misprediction penalty.

4. Methodology

In this section, we describe our methodology for the measurements and investigation of XML parsing energy consumption and performance. In brief, for each of the XML parsers under study we propose to measure and evaluate the energy consumption required to carry out XML well-formedness checking, under a variety of workloads, and as executed on three different Intel cores.

To begin our study, we propose to first investigate each of the XML parsers in terms of the PMCs hardware events as listed in the following subsection. Based on the recommendation of previous proposals [1–3], we have chosen several key hardware performance events for which the authors indicate have a strong correlation to energy consumption. We also measure other runtime counts such as the number of SIMD instructions and bitwise operations using the PIN binary instrumentation framework. From these data, we hope to gain insight into the XML parser execution characteristics and compare and contrast different industrial parsers.

The foundational work by Bellosa in [1] as well as more recent work in [2, 3] show that hardware-usage patterns has a significant impact in the energy consumption of a particular application; [1–3] further show that there is a strong correlation between specific performance events and energy usage—but the authors of each differ slightly in opinion as to which performance monitoring counters¹ (PMCs) to use.

The following subsections describe the XML parsers under study, XML workloads, the hardware architectures, PMC hardware events selected for measurement, and the energy measurement set up. We analyze the performance of the different parsers based on the hardware performance counter measurements and contrast their energy consumption measurements based on direct measurement.

4.1 Parsers

The XML parsing technologies selected for this study are the Parabix2, Xerces-C++, and Expat XML parsers. Parabix2 [14] (parallel bit streams for XML) is the second generation Parabix parser. Parabix2 is an open-source XML parser that leverages the SIMD capabilities of modern commodity processors; it employs the new parallelization techniques using parallel parsing with bit stream addition to deliver dramatic performance improvements over traditional byte-at-a-time parsing technology. Xerces-C++ version 3.1.1 (SAX) [15] is a validating open source XML parser

¹Performance monitoring counters (PMCs) are special-purpose registers that are included in most modern microprocessors; they store the running count of specific hardware events, such as retired instructions, cache misses, branch mispredictions, and arithmetic-logic unit operations to name a few. They can be used to capture information about any program at run-time, under any workload, at a very fine granularity.

written in C++ by the Apache project. Expat version 2.0.1 [10] is a non-validating XML parser library written in C.

4.2 Workloads

Distinguishing between “document-oriented” XML and “data-oriented” XML is a popular way to describe the two basic classes of XML documents. Data-oriented XML is used as an interchange format. Document-oriented XML is used to impose structure on information that rarely fits neatly into a relational database—particularly information intended for publishing. Data-oriented XML are characterized by a higher markup density. Markup density is defined as the ratio of the total markup contained within an XML file to the total XML document size. This metric may have substantial influence on the performance of XML parsing. As such we choose workloads with a spectrum of markup densities.

Table 1 shows the document characteristics of the XML input files selected for this performance study. The jawiki.xml and dewiki.xml XML files represent document-oriented XML inputs, containing three-byte and four-byte UTF8 sequence. The remaining files are data-oriented inputs and consist of only ASCII characters. [8]

4.3 Platform Hardware

Intel Core 2.

The Intel Core 2 is a Conroe based processor produced by Intel. Table 2 gives the hardware description of the Intel Core 2 machine selected.

Processor	Intel Core 2 Duo processor 6400 (2.13GHz)
L1 Cache	32KB I-Cache, 32KB D-Cache
L2 Cache	2MB
Front Side Bus	1066 MHz
Memory	2GB
Hard disk	80GB SCSI
Max TDP	65W

Table 2: Core 2

Intel Core i3.

The Intel Core i3 is a Nehalem based processor produced by Intel. The intent of this processor is to serve as an example low end server processor. Table 3 gives the hardware description of the Intel Core i3 machine selected.

Processor	Intel Clarkdale I3-530 (2.93GHz)
L1 Cache	32KB I-Cache, 32K D-Cache
L2 Cache	256KB
L3 Cache	4-MB
Front Side Bus	1333 MHz
Memory	4GB
Hard disk	SCSI 1TB
Max TDP	73W

Table 3: Core i3

Intel Core i5.

The Intel Core i5 is a Sandy Bridge based processor produced by Intel. Table 4 gives the hardware description of the Intel Core i3 machine selected.

4.4 PMC Hardware Events

Each of the hardware events selected relates to the energy consumption due to one or more hardware units. For example, total

Processor	Intel Core I5-2300 (2.80GHz)
L1 Cache	192 KB
L2 Cache	4 X 256KB
L3 Cache	6-MB
Front Side Bus	1333 MHz
Memory	6GB DDR3
Hard disk	SATA 1TB
Max TDP	95W

Table 4: Sandy Bridge

branch miss predictions corresponds to the use of the branch misprediction unit.

Initial PMC hardware event set:

- Processor Cycles
- Branch Instructions
- Branch Mispredictions
- Integer Instructions
- SIMD Instructions
- Cache Misses

4.5 Energy Measurement

To measure energy we use a Fluke i410 current clamp applied on the 12V wires that supply power to the processor sockets. The clamp detects the magnetic field created by the flowing current and converts it into voltage levels (1mV per 1A current). The voltage levels are then monitored by an Agilent 34410a multimeter at the granularity of 100 samples per second. This measurement captures the power to the processor package, including cores, caches, Northbridge memory controller, and the quick-path interconnects. [11].

5. Baseline Evaluation on Corei3

5.1 Cache behavior

Core i3 has a three level cache hierarchy. The miss penalty for each level is about 4 cycles, 11 cycles, and 36 cycles. Figure 6, Figure 7 and Figure 8 show the L1, L2 and L3 data cache misses of all the four parsers. Although XML parsing is not a memory intensive application, the cost of cache miss for Expat and Xerces can be about half cycle per byte while the performance of Parabix is hardly affected by cache misses. Cache miss isn’t just a problem for performance but also energy consumption. L1 cache miss cost about 8.3nJ; L2 cache miss cost about 19nJ; L3 cache miss cost about 40nJ. With a 1GB input file, Expat would consume more than 0.6J and Xerces would consume 0.9J on cache misses alone.

5.2 Branch Mispredictions

Despite years of improvement, branch misprediction is still a significant bottleneck of performance. The penalty of a branch misprediction is generally more than 10 CPU cycles. As shown in Figure 10, the cost of branch mispredictions for Expat can be more than 7 cycles per byte, which is as much as the processing time of Parabix2 on the same workload.

Reducing the branch misprediction rate is difficult for text-based applications due to the variable-length nature of syntactic elements. Therefore, the alternative solution of reducing branches becomes more attractive. However, the traditional byte-at-a-time method of XML parsing usually involves large amount of inevitable branches. As shown in Figure 9, Xerces can have an average of 13 branches for each byte it processed on the high markup density

File Name	dewiki.xml	jawiki.xml	roads.gml	po.xml	soap.xml
File Type	document	document	data	data	data
File Size (kB)	66240	7343	11584	76450	2717
Markup Item Count	406792	74882	280724	4634110	18004
Markup Density	0.07	0.13	0.57	0.76	0.87

Table 1: XML Document Characteristics

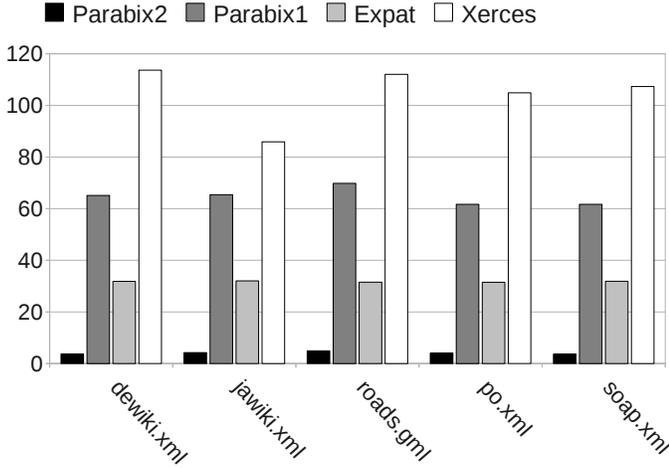


Figure 6: L1 Data Cache Misses on Core i3 (y-axis: Cache Misses per KByte)

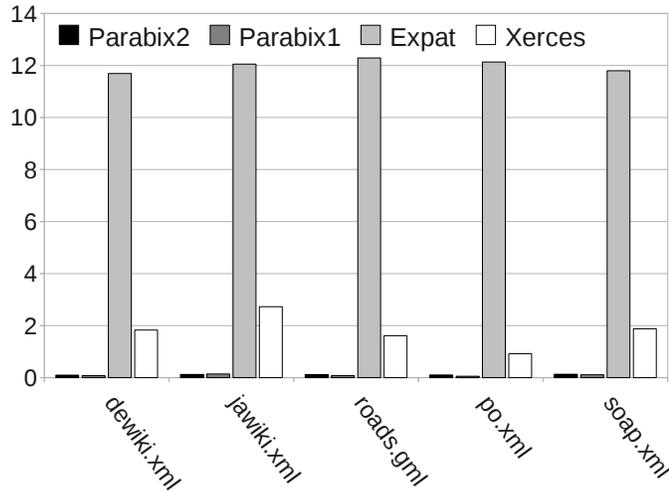


Figure 7: L2 Data Cache Misses on Core i3 (y-axis: Cache Misses per KByte)

file. Parabix substantially eliminate the branches by using parallel bit streams. Parabix1 still have a few branches for each block of 128 bytes (SSE) due to the sequential scanning. But with the new parallel scanning technique, Parabix2 is essentially branch-free as shown in the Figure 9. As a result, Parabix2 has minimal dependency on the markup density of the workloads.

5.3 SIMD/Total Instructions

Parabix gains its performance by using parallel bitstreams, which are mostly generated and calculated by SIMD instructions. The ratio of executed SIMD instructions over total instructions indicates the amount of parallel processing we were able to achieve. We use Intel pin, a dynamic binary instrumentation tool, to gather

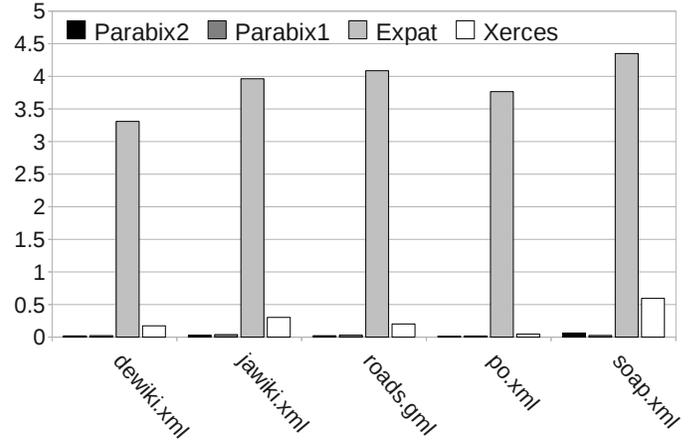


Figure 8: L3 Cache Misses on Core i3 (y-axis: Cache Misses per KByte)

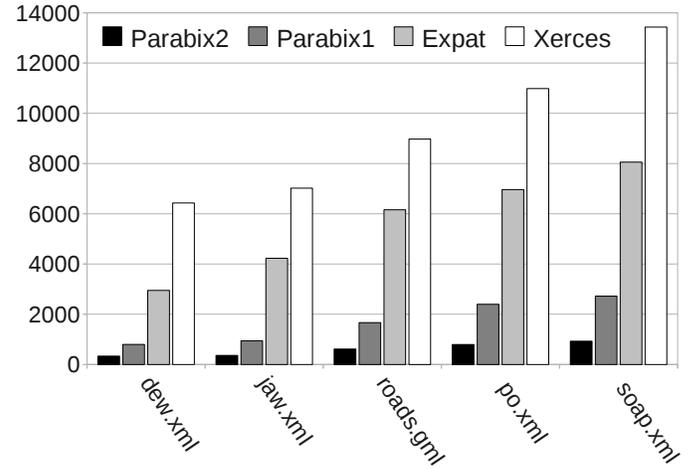


Figure 9: Branches on Core i3 (y-axis: Branches per KByte)

instruction mix. Then we adds up all the vector instructions that have been executed. Figure 11 and Figure 12 show the percentage of SIMD instructions of Parabix1 and Parabix2 (Expat and Xerce do not use any SIMD instructions). For Parabix1, 18% to 40% of the executed instructions consists of SIMD instructions. By using bistream addition for parallel scanning, Parabix2 uses 60% to 80% SIMD instructions. Although the ratio decrease as the markup density increase for both Parabix1 and Parabix2, the decreasing rate of Parabix2 is much lower and thus the performance degradation caused by increasing markup density is smaller.

5.4 CPU Cycles

Figure 13 shows the result of the overall performance evaluated as CPU cycles per thousands input bytes. Parabix1 is 1.5 to 2.5 times faster on document-oriented input and 2 to 3 times faster on data-oriented input compared with Expat and Xerces. Parabix2 is

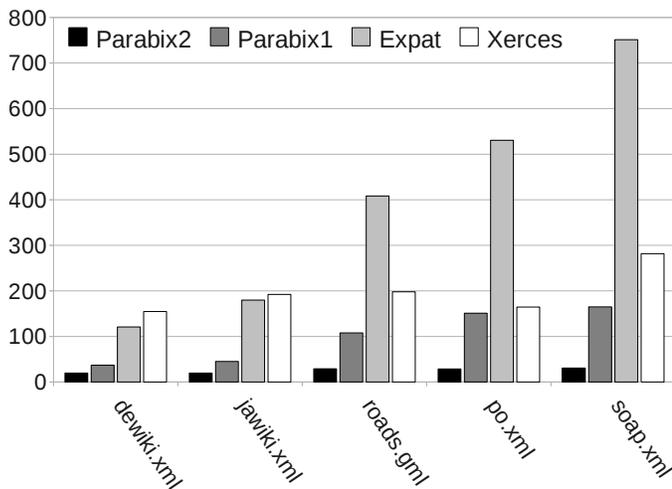


Figure 10: Branch Mispredictions on Core i3 (y-axis: Branch Mispredictions per KByte)

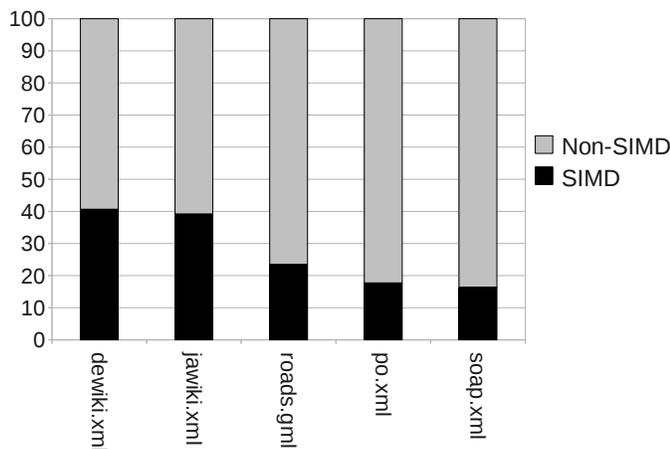


Figure 11: Parabix1 SIMD Instruction Ratio (y-axis: percent)

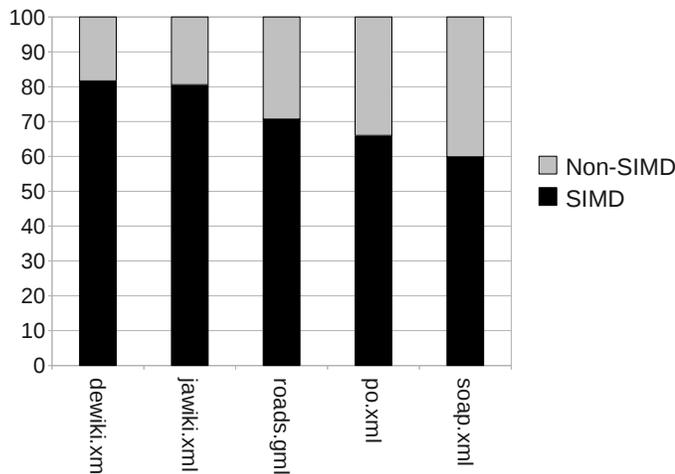


Figure 12: Parabix2 SIMD Instruction Ratio (y-axis: percent)

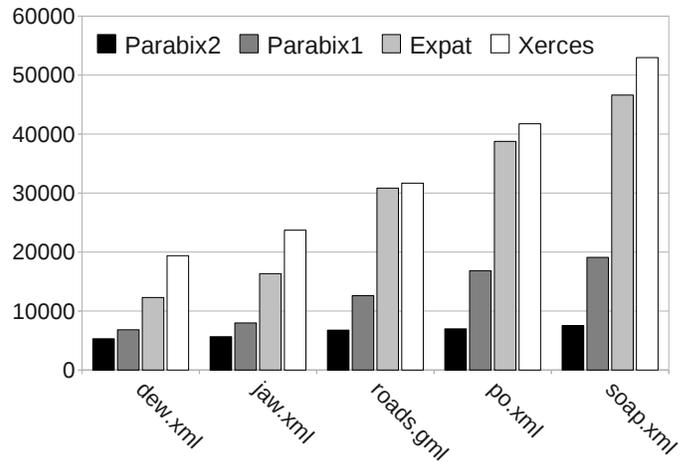


Figure 13: Processing Time on Core i3 (y-axis: Total CPU Cycles per KByte)

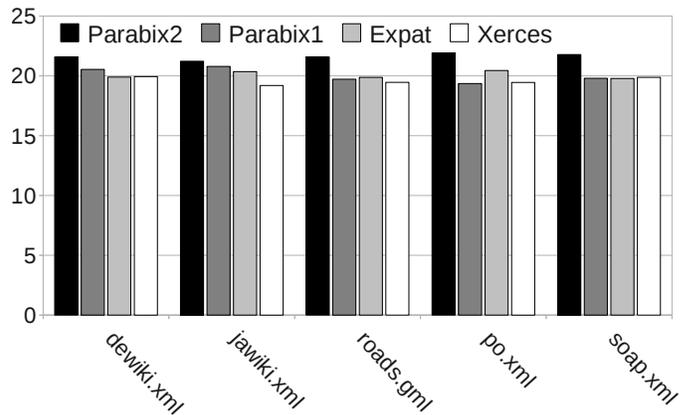


Figure 14: Average Power on Core i3 (watts)

2.5 to 4 times faster on document-oriented input and 4.5 to 7 times faster on data-oriented input. Traditional parsers can be dramatically slowed down by higher markup density while Parabix with parallel processing is less affected. The comparison is not entirely fair for Xerces that transcodes input into UTF-16, which typically takes several cycles per byte. However, transcoding using parallel bitstreams can be much faster and it takes less than a cycle per byte to transcode ASCII files such as road.gml, po.xml and soap.xml [6].

5.5 Power and Energy

There is a growing concern of power consumption and energy efficiency. Chip producers not only work on improving the performance but also have worked hard to develop power efficient chips. We studied the power and energy consumption of Parabix in comparison with Expat and Xerces on corei3.

Figure 14 shows the average power consumed by the four different parsers. The average power of corei3-530 is about 21 watts. This model released by Intel last year has a good reputation for power efficiency. Parabix2 dominated by SIMD instructions uses only about 5% higher power than the other parsers.

The more interesting trend is energy, Figure 15 shows the energy consumption of the four different parsers. Although Parabix2 needs slight higher power, its processing time is much shorter and therefore consumes much less energy. Parabix2 consumes 50 to 75 nJ per byte while Expat and Xerces consumes 80nJ to 320nJ and 140nJ to 370nJ per byte separately.

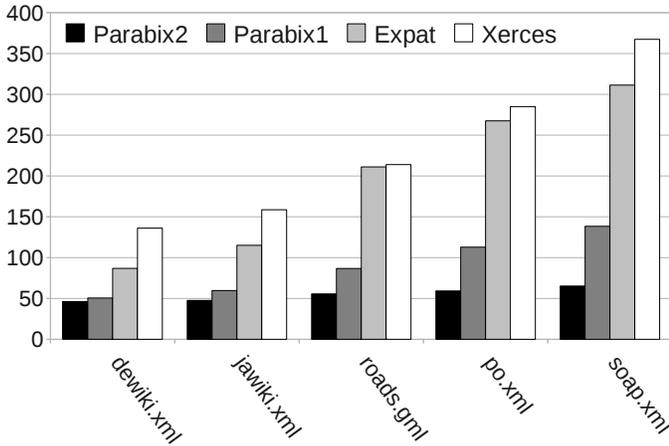


Figure 15: Energy Consumption on Core i3 (μJ per KByte)

6. Scalability

6.1 Performance

Figure 16 (a) shows the performance of Parabix2 on three different cores: core2, corei3 and sandybridge. The average processing time of the five workloads, which is evaluated as CPU cycles per thousand bytes, is divided up by bitstream parsing and byte space postprocessing. Bitstream parsing, mainly consists of SIMD instructions, is able to achieve 17% performance improvement moving from core2 to corei3; 22% performance improvement moving from corei3 to sandybridge, which is relatively stable compared to postprocessing, which gains 18% to 31% performance moving from core2 to corei3; 0 to 17% performance improvement moving from corei3 to sandybridge.

As comparison, we also measured the performance of Expat on all the three cores, which is shown in Figure 16 (b). The performance improvement is less than 5% by running Expat on corei3 instead of core2 and it is less than 10% by running on sandybridge instead of corei3.

Parabix2 scales much better than Expat and is able to achieve an overall performance improvement up to 26% simply by running the same code on a newer core. Further improvement on sandybridge with AVX will be discussed in the next section.

6.2 Power and Energy

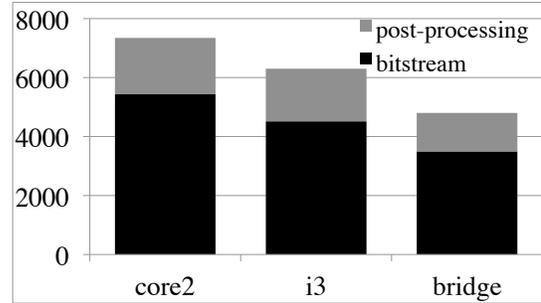
The newer processors are not only designed to have better performance but also more energy-efficient. Figure 17 shows the average power when running Parabix2 on core2, corei3 and sandybridge with different input files. On core2, the average power is about 32 watts. Corei3 saves 30% of the power compared with core2. Sandybridge saves 25% of the power compared with corei3 and consumes only 15 watts.

The energy consumption is further improved by better performance, which means a shorter processing time, as we moved to the newer cores. As a result, Parabix2 on sandybridge cost 72% to 75% less energy than Parabix2 on core2.

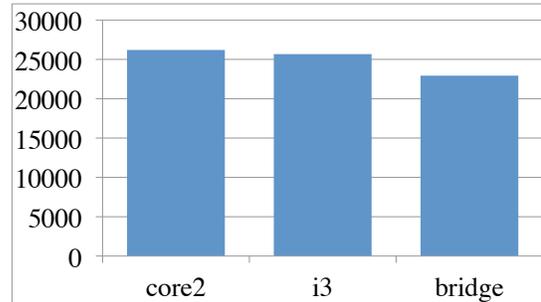
7. Scaling Parabix2 for AVX

Technology

Parabix2 was originally developed for 128-bit SSE2 technology widely available on all 64-bit Intel and AMD processors. In this section, we discuss the scalability and performance of Parabix2 to take advantage of the new 256-bit AVX (Advanced Vector Extensions) technology that has just become commercially available in the latest Intel processors based on the Sandy Bridge microarchi-



(a) Parabix2



(b) Expat

Figure 16: Processing Time of Parabix and Expat (y-axis: Total CPU Cycles per KByte)

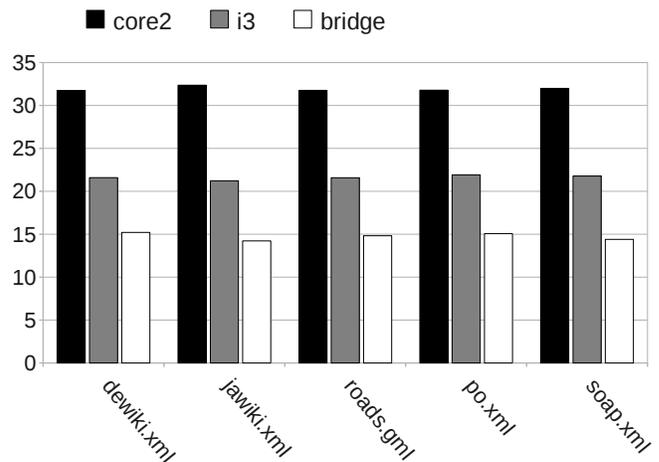


Figure 17: Average Power of Parabix2 (watts)

ture.

7.1 Three Operand Form

In addition to the introduction of 256-bit operations, AVX technology also makes a change in the structure of the base SSE instructions, moving from a destructive 2-operand form long used with SSE technologies to a nondestructive 3-operand form. In the 2-operand form, one register is used as both a source and destination register, equivalent to the assignment $a = a [\text{op}] b$. Thus, whenever the subsequent instructions used the value of both a and b , one of them had to be copied beforehand, or reconstituted or reloaded afterwards in order to recover the value. With 3-operand form, output may be directed to a third register independent of the source operands, as reflected by the assignment $c = a [\text{op}] b$. By avoiding the copying or reconstituting of operand values, a considerable reduction in instruction count may be possible. AVX technology

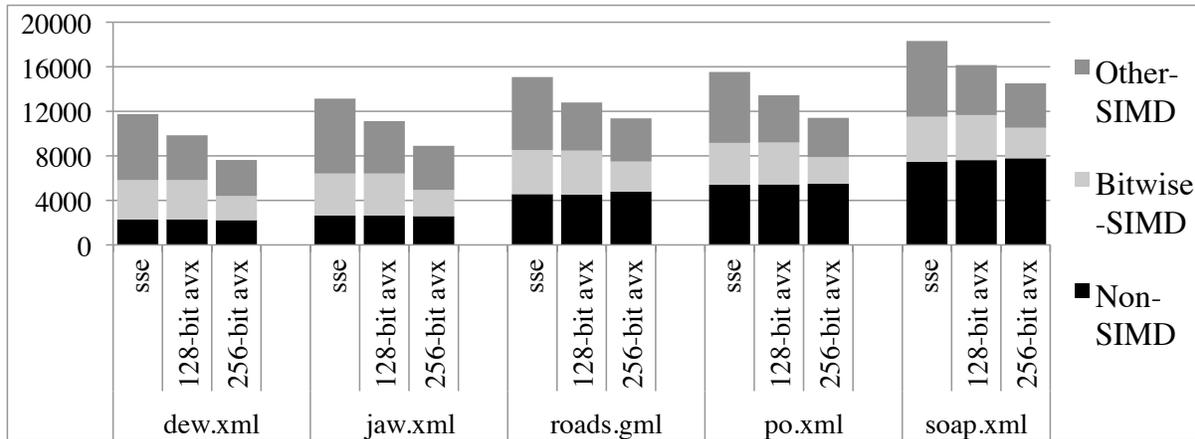


Figure 19: Parabix2 Instruction Counts (y-axis: Instructions per KByte)

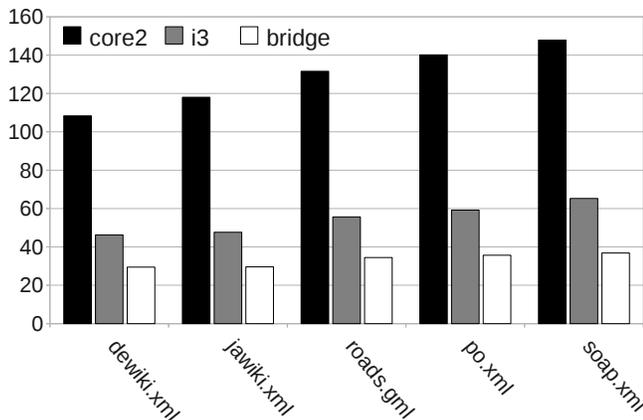


Figure 18: Energy consumption of Parabix2 (nJ/B)

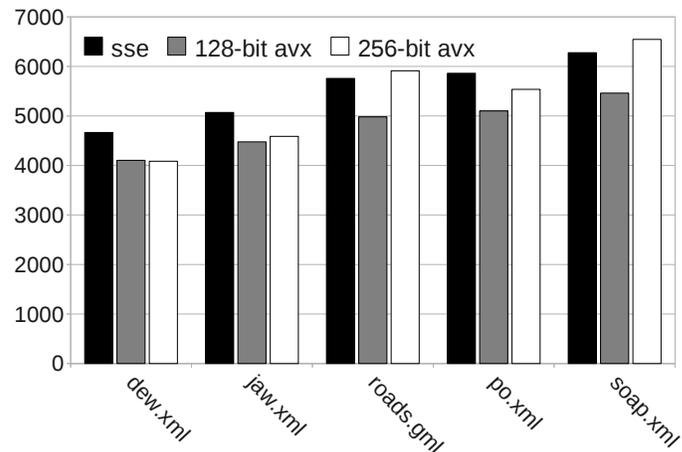


Figure 20: Parabix2 Performance (y-axis: CPU cycles per KB)

makes available the 3-operand form both with the new 256-bit operations as well as with base 128-bit operations of SSE.

7.2 256-bit Operations

With the introduction of 256-bit SIMD registers with AVX technology, one might ideally expect up to a 50% reduction in the instruction count for the SIMD workload of Parabix2. However, in the Sandy Bridge implementation, Intel has focused on implementing floating point operations as opposed to the integer based operations. That is, 256-bit SIMD is available for loads, stores, bitwise logic and floating operations, while SIMD integer operations and shifts are only available in 128-bit form. Nevertheless, with loads, stores and bitwise logic comprising a major portion of the Parabix2 SIMD instruction mix, a substantial reduction in instruction count and consequent performance improvement was anticipated.

7.3 Performance Results

We implemented two versions of Parabix2 using AVX technology. The first was simply the recompilation of the existing Parabix2 source code to take advantage of the 3-operand form of AVX instructions while retaining a uniform 128-bit SIMD processing width. The second involved rewriting core library functions for Parabix2 to use 256-bit AVX operations wherever possible and to simulate the remaining operations using pairs of 128-bit operations.

Figure 19 shows the reduction in instruction counts achieved in these two versions. For each workload, the base instruction count of the Parabix2 binary compiled in SSE-only mode is shown with the caption “sse,” the version obtained by simple recompi-

lation with AVX-mode enabled is labeled “avx 128-bit,” and the version reimplemented to use 256-bit operations wherever possible is labelled “avx 256-bit.” The instruction counts are divided into three classes. The “non-SIMD” operations are the general purpose instructions that use neither SSE nor AVX technology. The “bitwise SIMD” class comprises the bitwise logic operations, that are available in both 128-bit form and 256-bit form. The “other SIMD” class comprises all other SIMD operations, primarily comprising the integer SIMD operations that are available only at 128-bit widths even with 256-bit AVX technology.

Note that, in each workload, the number of non-SIMD instructions remains relatively constant with each workload. As may be expected, however, the number of “bitwise SIMD” operations remains the same for both SSE and 128-bit while dropping dramatically when operating 256-bits at a time. Ideally one may expect up to a 50% reduction in these instructions versus the 128-bit AVX. The actual reduction measured was 32%–39% depending on workload. Because some bitwise logic is needed in implementation of simulated 256-bit operations, the full 50% reduction in bitwise logic was not achieved.

The “other SIMD” class shows a substantial “30%-35%” reduction with AVX 128-bit technology compared to SSE. This reduction is due to eliminated copies or reloads when SIMD operations are compiled using 3-operand AVX form versus 2-operand SSE form. A further 10%–20% reduction is observed with Parabix2 version

rewritten to use 256-bit operations.

While the successive reductions in SIMD instruction counts are quite dramatic with the two AVX implementations of Parabix2, the performance benefits are another story. As shown in Figure 20, the benefits of the reduced SIMD instruction count are achieved only in the AVX 128-bit version. In this case, the benefits of 3-operand form seem to fully translate to performance benefits. Bizarrely, perhaps, the performance of Parabix2 in the 256-bit AVX implementation does not improve significantly and actually degrades for files with higher markup density. We believe that this is primarily due to the current AVX implementation in Sandy Bridge, with significant latency in many of the 256-bit instructions in comparison to their 128-bit counterparts. If these latency issues can be addressed in future AVX implementations, further substantial performance and energy benefits could be realized in XML parsing with Parabix2

8. Conclusion

This paper has examined energy efficiency and performance characteristics of four XML parsers considered over three generations of Intel processor architecture and shown that parsers based on parallel bit stream technology have dramatically better performance, energy efficiency and scalability than traditional byte-at-a-time parsers widely deployed in current software. Based on a novel application of the short vector SIMD technology commonly found in commodity processors of all kinds, parallel bit stream technology scales well with improvements in processor SIMD capabilities. With the recent introduction of the first generation of Intel processors that incorporate AVX technology, the change to 3-operand form SIMD operations has delivered a substantial benefit for the Parabix2 parsers simply through recompilation. Restructuring of Parabix2 to take advantage of the 256-bit SIMD capabilities also delivered a substantial reduction in instruction count, but without corresponding performance benefits in the first generation of AVX implementations.

There are many directions for further research. These include compiler and tools technology to automate the low-level programming tasks inherent in building parallel bit stream applications, widening the research by applying the techniques to other forms of text analysis and parsing, and further investigation of the interaction between parallel bit stream technology and processor architecture. Two promising avenues include investigation of GPGPU approaches to parallel bit stream technology and the leveraging of the intraregister parallelism inherent in this approach to also take advantage of the intrachip parallelism of multicore processors.

9. References

- [1] F. Bellosa. The case for event-driven energy accounting. Technical Report TR-14-01-07, University of Erlangen, Department of Computer Science, June 2001.
- [2] R. Bertran, M. Gonzalez, X. Martorell, N. Navarro, and E. Ayguade. Decomposable and responsive power models for multicore processors using performance counters. In *Proceedings of the 24th ACM International Conference on Supercomputing, ICS '10*, pages 147–158, New York, NY, USA, 2010. ACM.
- [3] W. Bircher and L. John. Complete system power estimation: A trickle-down approach based on performance events. In *Performance Analysis of Systems Software, 2007. ISPASS 2007. IEEE International Symposium on*, pages 158–168, Apr. 2007.
- [4] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau. Extensible markup language (XML) 1.0 (fifth edition). W3C Recommendation, 2008.
- [5] R. Cameron, K. Herdy, and E. Amiri. Parallel bit stream technology as a foundation for XML parsing performance. In *International Symposium on Processing XML Efficiently: Overcoming Limits on Space, Time, or Bandwidth*, Aug. 2009.
- [6] R. D. Cameron. A case study in SIMD text processing with parallel bit streams: UTF-8 to UTF-16 transcoding. In *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming, PPOPP '08*, pages 91–98, New York, NY, USA, 2008. ACM.
- [7] R. D. Cameron, E. Amiri, K. S. Herdy, D. Lin, T. C. Shermer, and F. P. Popowich. Parallel parsing with bitstream addition: An XML case study. Technical Report TR 2010-11, Simon Fraser University, School of Computing Science, October 2010.
- [8] R. D. Cameron, K. S. Herdy, and D. Lin. High performance XML parsing using parallel bit stream technology. In *CASCON '08: Proceedings of the 2008 conference of the center for advanced studies on collaborative research*, pages 222–235, New York, NY, USA, 2008. ACM.
- [9] R. D. Cameron and D. Lin. Architectural support for SWAR text processing with parallel bit streams: the inductive doubling principle. In *ASPLOS '09: Proceeding of the 14th international conference on Architectural support for programming languages and operating systems*, pages 337–348, New York, NY, USA, 2009. ACM.
- [10] J. Clark. The Expat XML Parser. <http://expat.sourceforge.net/>.
- [11] F. Corporation. Fluke Clamp Meters. <http://www.fluke.com/>.
- [12] Z. Dai, N. Ni, and J. Zhu. A 1 cycle-per-byte XML parsing accelerator. In *FPGA '10: Proceedings of the 18th Annual ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, pages 199–208, New York, NY, USA, 2010. ACM.
- [13] B. DuCharme. Documents vs. data, schemas vs. schemas. In *XML 2004*, Washington D.C., 2004.
- [14] R. D. C. et al. Parabix2. <http://parabix.costar.sfu.ca/>.
- [15] A. S. Foundation. Xerces C++ Parser. <http://xerces.apache.org/xerces-c/>.
- [16] Z. Lei. XML parsing accelerator with Intel streaming SIMD extensions 4 (Intel SSE4). <http://software.intel.com/en-us/articles/xml-parsing-accelerator-with-intel-streaming-simd-extensions-4-intel-sse4/>, 2008.
- [17] M. Leventhal and E. Lemoine. The XML chip at 6 years. In *International Symposium on Processing XML Efficiently: Overcoming Limits on Space, Time, or Bandwidth*, Aug. 2009.
- [18] X. Li, H. Wang, T. Liu, and W. Li. Key elements tracing method for parallel XML parsing in multi-core system. *Parallel and Distributed Computing Applications and Technologies, International Conference on*, 0:439–444, 2009.
- [19] Matthias Nicola and Jasmi John. XML Parsing: A Threat to Database Performance. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, New Orleans, Louisiana, 2003.
- [20] Perkins, E. and Kostoulas, M. and Heifets, A. and Matsa, M. and Mendelsohn, N. Performance Analysis of XML APIs. In *XML 2005*, Atlanta, Georgia, Nov. 2005.
- [21] B. Shah, P. Rao, B. Moon, and M. Rajagopalan. A data parallel algorithm for XML DOM parsing. In Z. BellahsÁlne, E. Hunt, M. Rys, and R. Unland, editors, *Database and XML Technologies*, volume 5679 of *Lecture Notes in Computer Science*, pages 75–90. Springer Berlin / Heidelberg, 2009.
- [22] Y. Zhang, Y. Pan, and K. Chiu. Speculative p-DFAs for parallel XML parsing. In *High Performance Computing (HiPC), 2009 International Conference on*, pages 388–397, Dec. 2009.