

Boosting the Efficiency of Text Processing on Commodity

Processors: The Parabix Story

Paper ID ****

Abstract

In modern applications text files are employed widely. For example, XML files provide data storage in human readable format and are ubiquitous in applications ranging from database systems to mobile phone SDKs. Traditional text processing tools are built around a byte-at-a-time processing model where each character token of a document is examined. The byte-at-a-time model is highly challenging for commodity processors. It includes many unpredictable input-dependent branches which cause pipeline squashes and stalls. Furthermore, typical text processing tools perform few operations per processed character and experience high cache miss rates. Overall, parsing text in important domains like XML processing requires high performance motivating hardware designers to adopt ASIC solutions.

In this paper, we enable text processing applications to effectively use commodity processors. We introduce Parabix (Parallel Bitstream) technology, a software toolkit and execution framework that allows applications to exploit modern SIMD instructions extensions for high performance text processing. Parabix enables the application developer to write constructs assuming unlimited SIMD data parallelism and Parabix's bitstream translator generates code based on machine specifics (e.g., SIMD register widths). The key insight into efficient text processing in Parabix is the data organization. Parabix transposes the sequence of character bytes into sets of 8 parallel bit streams which then enables us to operate on multiple characters with single bit-parallel SIMD operators. We demonstrate the features and efficiency of Parabix with a XML parsing application. We evaluate a Parabix-based XML parser against two widely used XML parsers, Expat and Apache's Xerces, and across three generations of x86 processors, including the new Intel SandyBridge. We show that Parabix's speedup is $2\times$ – $7\times$ over Expat and Xerces. We observe that Parabix overall makes efficient use of intra-core parallel hardware on commodity processors and supports significant gains in energy. Using Parabix, we assess the scalability advantages of SIMD processor improvements across Intel processor generations, culminating with a look at the latest 256-bit AVX technology in SandyBridge versus the now legacy 128-bit SSE technology. We also examine Parabix on mobile platforms using ARM processors with Neon SIMD extensions. Finally, we partition the XML program into pipeline stages and demonstrate that thread-level parallelism exploits SIMD units scattered across the different cores and improves performance ($2\times$ on 4 cores) at same energy levels as the single-thread version.

1 Introduction

We have now long since reached the limit to classical Dennard voltage scaling that enabled us to keep all of transistors afforded by Moore's law active. This has already resulted in a rethink of the way general-purpose processors are built: processor frequencies have remained stagnant over the last 5 years with

#****

HPCA 2012 Submission #****. CONFIDENTIAL REVIEW COPY. DO NOT DISTRIBUTE.

#****

0056
0057
0058
0059
0060
0061
0062
0063
0064
0065
0066
0067
0068
0069
0070
0071
0072
0073
0074
0075
0076
0077
0078
0079
0080
0081
0082
0083
0084
0085
0086
0087
0088
0089
0090
0091
0092
0093
0094
0095
0096
0097
0098
0099
0100
0101
0102
0103
0104
0105
0106
0107
0108
0109
0110
0111

the capability to boost core speeds on Intel multicores only if other cores on the chip are shut off. Chip makers strive to achieve energy efficient computing by operating at more optimal core frequencies and aim to increase performance with a larger number of cores. Unfortunately, given the limited levels of parallelism that can be found in applications [5], it is not certain how many cores can be productively used in scaling our chips [14]. This is because exploiting parallelism across multiple cores tends to require heavyweight threads that are difficult to manage and synchronize.

The desire to improve the overall efficiency of computing is pushing designers to explore customized hardware [18, 21] that accelerate specific parts of an application while reducing the overheads present in general-purpose processors. They seek to exploit the transistor bounty to provision many different accelerators and keep only the accelerators needed for an application active while switching off others on the chip to save power consumption. While promising, given the fast evolution of languages and software, its hard to define a set of fixed-function hardware for commodity processors. Furthermore, the toolchain to create such customized hardware is itself a hard research challenge. We believe that software, applications, and runtime models themselves can be refactored to significantly improve the overall computing efficiency of commodity processors.

In this paper, we tackle the infamous “thirteenth dwarf” (parsers/finite state machines) that is considered to be the hardest application class to parallelize [1] and show how Parabix, a novel software architecture, tool chain and run-time environment can indeed be used to dramatically improve parsing efficiency on commodity processors. Based on the concept of transposing byte-oriented character data into parallel bit streams for the individual bits of each byte, the Parabix framework exploits the SIMD extensions on commodity processors (SSE/AVX on x86, Neon on ARM) to process hundreds of character positions in an input stream simultaneously. To achieve transposition, Parabix exploits sophisticated SIMD instructions that enable data elements to be packed and unpacked from registers in a regular manner which improves the overall cache access behavior of the application resulting in significantly fewer misses and better utilization. Parabix also dramatically reduces branches in parsing code resulting in a more efficient pipeline and substantially improves register/cache utilization which minimizes energy wasted on data transfers.

We apply Parabix technology to the problem of XML parsing and develop several implementations for different computing platforms. XML is a particularly interesting application; it is a standard of the web

0112 consortium that provides a common framework for encoding and communicating data. XML provides crit-
0113 ical data storage for applications ranging from Office Open XML in Microsoft Office to NDFD XML of
0114 the NOAA National Weather Service, from KML in Google Earth to Castor XML in the Martian Rovers,
0115 as well as ubiquitous XML data in Android phones. XML parsing efficiency is important for multiple
0116 application areas; in server workloads the key focus is on overall transactions per second, while in appli-
0117 cations in network switches and cell phones, latency and energy are of paramount importance. Traditional
0118 software-based XML parsers have many inefficiencies including considerable branch misprediction penal-
0119 ties due to complex input-dependent branching structures as well as poor use of memory bandwidth and
0120 data caches due to byte-at-a-time processing and multiple buffering. XML ASIC chips have been around
0121 for over 6 years, but typically lag behind CPUs in technology due to cost constraints. Our focus is how
0122 much we can improve performance of the XML parser on commodity processors with Parabix technology.
0123
0124
0125
0126
0127
0128
0129
0130
0131
0132
0133

0134 In the end, as summarized by Figure 1 our Parabix-based XML parser improves the performance by and
0135 energy efficiency several-fold compared to widely-used software parsers and approaching the performance
0136 of performance of ASIC XML parsers.¹
0137
0138
0139

0140 Overall we make the following contributions in this paper.

0141
0142 1) We introduce the Parabix architecture, tool chain and run-time environment and describe how it may
0143 be used to produce efficient XML parser implementations on a variety of commodity processors. While
0144 studied in the context of XML parsing, the Parabix framework can be widely applied to many problems in
0145 text processing and parsing.
0146
0147
0148
0149

0150 2) We compare our Parabix XML parsers against conventional parsers and assess the improvement in
0151 overall performance and energy efficiency on each platform. We are the first to compare and contrast
0152 SSE/AVX extensions across multiple generation of Intel processors and show that there are performance
0153 challenges when using newer generation SIMD extensions, possibly due to their memory interface. We
0154 compare the ARM Neon extensions against the x86 SIMD extensions and comment on the latency of
0155 SIMD operations across these architectures.
0156
0157
0158
0159
0160
0161

0162 3) Finally, building on the SIMD parallelism of Parabix technology, we multithread the Parabix XML
0163 parser to to enable the different stages in the parser to exploit SIMD units across all the cores. This further
0164
0165

0166 ¹The actual energy consumption of the XML ASIC chips is not published by the companies.
0167

0168
0169
0170
0171
0172
0173
0174
0175
0176
0177
0178
0179
0180
0181
0182
0183
0184
0185
0186
0187
0188
0189
0190
0191
0192
0193
0194
0195
0196
0197
0198
0199
0200
0201
0202
0203
0204
0205
0206
0207
0208
0209
0210
0211
0212
0213
0214
0215
0216
0217
0218
0219
0220
0221
0222
0223




Figure 1: XML Parser Technology Energy vs. Performance

improves performance while maintaining the energy consumption constant with the sequential version.

The remainder of this paper is organized as follows. Section 2 presents background material on XML parsing and provides insight into the inefficiency of traditional parsers on mainstream processors. Section 3 describes the Parabix architecture, tool chain and run-time environment. Section 4 describes the application of the Parabix framework to the construction of an XML parser meeting enforcing all the well-formedness rules of the XML specification. Section 5 then describes the overall methodology of our performance and energy study. Section 6 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, and SIMD instruction counts. Section 7 examines scalability and performance gains through three generations of Intel architecture. Section 8 examines the extension of the Parabix technology to take advantage of Intel’s new 256-bit AVX technology, while Section 9 investigates the applications of this technology on mobile platforms using ARM processors with Neon SIMD extensions. Section 10 then looks at the multithreading of the Parabix XML parser using pipeline parallelism. Section 11 concludes the paper.

2 Background

2.1 XML

In 1998, the W3C officially adopted XML as a standard. 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

0224 easily store semi-structured data in a descriptive fashion. XML markup encodes a description of an XML
0225 document's storage layout and logical structure. Because XML was intended to be human-readable, XML
0226 markup tags are often verbose by design [6].
0227
0228
0229

0230 XML files can be classified as "document-oriented" or "data-oriented" [13]. Document-oriented
0231 XML is designed for human readability, such as shown in Figure 2; data-oriented XML files are intended
0232 to be parsed by machines and omit "human-friendly" formatting techniques, such as the use of whitespace
0233 and descriptive "natural language" naming schemes. Although the XML specification itself does not
0234 distinguish between "XML for documents" and "XML for data" [6], the latter often requires the use of an
0235 XML parser to extract the information within. The role of an XML parser is to transform the text-based
0236 XML data into application ready data.
0237
0238
0239
0240
0241
0242
0243

```
0244 <?xml version="1.0"?>  
0245 <Products>  
0246   <Product ID="0001">  
0247     <ProductName Language="English">Widget</ProductName>  
0248     <ProductName Language="French">Bitoniau</ProductName>  
0249     <Company>ABC</Company>  
0250     <Price>$19.95</Price>  
0251   </Product>  
</Products>
```

0252 Figure 2: Example XML Document
0253
0254
0255

0256 2.2 Traditional XML Parsers 0257

0258 Traditional XML parsers process XML sequentially a single byte-at-a-time. Following this approach,
0259 an XML parser processes a source document serially, from the first to the last byte of the source file. Each
0260 character of the source text is examined in turn to distinguish between the XML-specific markup, such
0261 as an opening angle bracket 'i', and the content held within the document. The current character that
0262 the parser is processing is commonly referred to using the concept of a current cursor position. As the
0263 parser moves the cursor through the source document, the parser alternates between markup scanning, and
0264 data validation and processing operations. At each processing step, the parser scans the source document
0265 and either locates the expected markup, or reports an error condition and terminates. In other words,
0266 traditional XML parsers operate as complex finite-state machines that use byte comparisons to transition
0267 between data and metadata states. Each state transition indicates the context in which to interpret the
0268 subsequent characters. Unfortunately, textual data tends to consist of variable-length items sequenced in
0269
0270
0271
0272
0273
0274
0275
0276
0277
0278
0279

0280 generally unpredictable patterns [8]; thus any character could be a state transition until deemed otherwise.

0281 Expat and Xerces-C are popular byte-a-time sequential parsers. Both are C/C++ based and open-source.
0282
0283 Expat was originally released in 1998; it is currently used in Mozilla Firefox and provides the core func-
0284 tionality of many additional XML processing tools [11]. Xerces-C was released in 1999 and is the foun-
0285 dation of the Apache XML project [17].
0286
0287
0288
0289

0290 A major disadvantage of the sequential byte-at-a-time approach to XML parsing is that each XML
0291 character incurs at least one conditional branch. The cumulative effect of branch mispredictions penal-
0292 ties are known to degrade XML parsing performance in proportion to the markup density of the source
0293 document [9] (i.e., the proportion of XML-markup to XML-data).
0294
0295
0296
0297

0298 2.3 Parallel XML Parsing

0300 In general, parallel XML acceleration methods come in one of two forms: multithreaded approaches
0301 and SIMD-based techniques. Multithreaded XML parsers take advantage of multiple cores via number of
0302 strategies. Common strategies include preparsing the XML file to locate key partitioning points [22] and
0303 speculative p-DFAs [22]. SIMD XML parsers leverage the SIMD registers to overcome the performance
0304 limitations of the sequential byte-at-a-time processing model and its inherently data dependent branch
0305 misprediction rates. Further, data parallel SIMD instructions allow the processor to perform the same
0306 operation on multiple pieces of data simultaneously. The Parabix1 and Parabix2 parsers studied in this
0307 paper fall under the SIMD classification. The Parabix parser versions studied are described in further
0308 detail in Section 3.
0309
0310
0311
0312
0313
0314
0315
0316
0317
0318
0319

0320 3 Parabix

0321 This section presents an overview of the SIMD-based parallel bit stream text processing framework,
0322 *Parabix*. In Section 4, we discuss one of its implementations, *Parabix-XML*. A more comprehensive study
0323 of Parabix and Parabix-XML, can be found in the technical report “Parallel Parsing with Bitstream Addi-
0324 tion: An XML Case Study” [8].
0325
0326
0327
0328
0329

0330 The fundamental difference between the Parabix framework and traditional text processing models is
0331 in how Parabix represents the source data. Rather than viewing it as a stream of bytes or characters and
0332 performing per-byte comparisons to determine when to transition between parsing states, Parabix views
0333
0334
0335

0336 data as a set of *bit streams* and uses a mixture of boolean-logic and integer-based math on those streams
0337
0338 to effectively parse many bytes in parallel. In this paper, we use the notation \wedge , \vee and \neg to denote the
0339
0340 AND, OR and NOT boolean operations. Bit streams are the foundation of the Parabix technology. In order
0341
0342 to understand how it is possible to perform complex state-transition-based text processing operations in
0343
0344 parallel, understanding what bit streams are, how they are created, and — more importantly — how they
0345
0346 are used by Parabix, is critical.
0347

0348 3.1 What are Bit Streams?

0349
0350 A bit stream is simply a sequence of 1s and 0s. The significance of each bit value is dependent on the
0351
0352 type of bit stream. We view bit streams in one of two ways: n 1-bit values for boolean-logic operations
0353
0354 and 1 n -bit value for integer-based math. For simplicity, assume that n is infinitely long w.r.t. the size
0355
0356 of the source data. In reality, each bit stream is divided into a series of w -bit blocks, where w is equal to
0357
0358 the width of the SIMD registers within the system (e.g., 64-bits for MMX, 128-bits for SSE/NEON, and
0359
0360 256-bits for AVX). We discuss how these $\frac{n}{w}$ bit stream segments can be used as infinitely-long bit stream
0361
0362 in Section ??.
0363
0364

0365 The first type of bit streams used in Parabix are referred to as *basis bit streams*, which contain the input
0366
0367 source data. Parabix uses the basis bit streams to construct *character-class bit streams* in which each 1 bit
0368
0369 indicates the presense of a significant character (or class of characters) in the parsing process. Character-
0370
0371 class bit streams are used to compute *lexical bit streams* and *error bit streams*, which Parabix uses to
0372
0373 process and validate the source document. The remainder of this section will discuss each type of bit
0374
0375 stream.
0376

0377 **Basis Bit Streams:** To construct the basis bit streams, the source data is first loaded in sequential order
0378
0379 and then transposed — through a series of SIMD pack, shift, and bitwise operations — so that Parabix
0380
0381 can efficiently produce the character-class bit streams. Essentially, when the source data is in basis bit
0382
0383 stream form, the k -th bit of i -th character in the source text is in the i -th (bit) position of the k -th basis
0384
0385 bit stream, b_k . Using the SIMD capabilities of current commodity processors, the transposition process
0386
0387 incurs an amortized cost of approximately 1 cycle per byte [9]. The size of k is dependent on the code
0388
0389 unit size of the text encoding format of the source document. A code unit is simply a fixed number of bits
0390
0391

(or bytes) used to represent a specific character (code point). Some encoding formats, such as UNICODE, may use a multiple code units to express all of its possible code points. How multi-code-unit characters can be parsed efficiently goes beyond the scope of this paper. The most dominant format in data-oriented XML documents is ASCII, which uses a 8-bit code unit to represent all of the 128 code points within it. In Figure 3, we show how the ASCII string “b7<A” is represented as 8 basis bit streams, $b_{0..7}$. The bits used to construct b_7 have been highlighted in this example.

STRING								
ASCII	0110001 0	0011011 1	0011111 0	0100000 1				

b_0	b_1	b_2	b_3	b_4	b_5	b_6	b_7
0	1	1	0	0	0	1	0
0	0	1	1	0	1	1	1
0	0	1	1	1	1	0	0
0	1	0	0	0	0	0	1

Figure 3: Example 7-bit ASCII Basis Bit Streams

Character-class Bit Streams: Typically, as text parsers process input data, they locate specific characters to determine if and when to transition between data and metadata parsing. For example, in XML, any opening angle bracket character, ‘<’, may indicate that we are starting a new markup tag. Traditional byte-at-a-time parsers find these characters by comparing the value of each code point with a set of known code points and branching appropriately when one is found, typically using an if or switch statement. Using this method to perform multiple transitions in parallel is non-trivial and may require fairly sophisticated algorithms to do so correctly.

Character-class bit streams allow us to perform up to 128 “comparisons” in parallel in a single operation (using Intel SSE/ARM NEON) by using a series of boolean-logic operations to merge multiple basis bit streams into a single character-class stream that marks the positions of key characters with a 1. For example, a character is an ‘<’ if and only if $\neg(b_0 \vee b_1) \wedge (b_2 \wedge b_3 \wedge b_4 \wedge b_5) \wedge \neg(b_6 \vee b_7) = 1$. Classes of characters can be found with similar formulas. For example, a character is a number [0–9] if and only if $\neg(b_0 \vee b_1) \wedge (b_2 \wedge b_3) \wedge \neg(b_4 \wedge (b_5 \vee b_6))$. An important observation here is that a range of characters can sometimes take fewer operations and require fewer basis bit streams to compute than individual characters. Finding optimal solutions to all character-classes is non-trivial and goes beyond the scope of this paper.

0448
0449
0450
0451
0452
0453
0454
0455
0456
0457
0458
0459
0460
0461
0462
0463
0464
0465
0466
0467
0468
0469
0470
0471
0472
0473
0474
0475
0476
0477
0478
0479
0480
0481
0482
0483
0484
0485
0486
0487
0488
0489
0490
0491
0492
0493
0494
0495
0496
0497
0498
0499
0500
0501
0502
0503

Lexical and Error Bit Streams: To perform lexical analysis on the input data, Parabix computes lexical and error bit streams from the character-class bit streams using a mixture of both boolean logic and integer math. Lexical bit streams typically mark multiple current parsing positions. Unlike the single-cursor approach of traditional text parsers, these allow Parabix to process multiple cursors in parallel. Error bit streams are often the byproduct or derivative of computing lexical bit streams and can be used to identify any well-formedness issues found during the parsing process. The presense of a 1 bit in an error stream tends to mean that the lexical stream cannot be trusted to be completely accurate and Parabix may need to perform some sequential parsing on that section to determine the cause and severity of the error.

To form lexical bit streams, we have to introduce a few new operations: Advance and ScanThru. The Advance operator accepts one input parameter, c , which is typically viewed as a bit stream containing multiple cursor bits, and returns $c + c$ — effectively moves each cursor one position to the “right”. ScanThru accepts two input parameters, c and m ; any bit that is in both c and m is moved to first subsequent 0-bit in m by calculating $(c + m) \wedge \neg m$. For example, in Figure 4 suppose we have the regular expression $\langle [a-zA-Z]^+ \rangle$ and wish to find all instances of it in the source text. We begin by constructing the character classes C_0 , which consists of all letters, C_1 , which contains all ‘>’s, and C_2 , which marks all ‘<’s. In L_0 the position of every ‘<’ is advanced by one to locate the first character of each token. By computing E_0 , the parser notes that “<>” does not match the expected pattern. To find the end positions of each token, the parser calculates L_1 by moving the cursors in L_0 through the letter bits in C_0 . L_1 is then validated to ensure that each token ends with a ‘>’ and discovers that “<error]” too fails to match the expected pattern. With additional post bit-stream processing, the erroneous cursor positions in L_0 and L_1 can be removed or ignored; the details of which go beyond the scope of this paper.

source text	<a><valid> <string> <>ignored<>error]
$C_0 = [a-zA-Z]$.1..11111...111111....1111111..111111.
$C_1 = [>]$..1.....1.....1...1.....1.....
$C_2 = [<]$	1..1.....1.....1.....1.....1.....
$L_0 = Advance(C_2)$.1..1.....1.....1.....1.....1.....
$E_0 = L_0 \wedge \neg C_0$1.....
$L_1 = ScanThru(L_0, C_0)$..1.....1.....1...1.....1.....1.....
$E_1 = L_1 \wedge \neg C_1$1.....

Figure 4: Lexical Parsing in Parabix

0504 Using this parallel bit stream approach, conditional branch statements used to identify key positions
0505 and/or syntax errors at each each parsing position are mostly eliminated, which, as Section XXX shows,
0506 minimizes branch misprediction penalties. Accurate parsing and parallel lexical analysis is done through
0507 processor-friendly equations that requires no speculation nor multithreading.
0508
0509
0510
0511

0512 3.2 Parabix Tool Chain

0513 To support the Parabix framework, we are developing tool technology to automate various aspects of
0514 parallel bit stream programming. At present, our tool chain consists of two compilers: a character class
0515 compiler (ccc) and a unbounded bit stream to C/C++ block-at-a-time processing compiler (Pablo).
0516
0517
0518
0519
0520

0521 The character class compiler is used to automatically produce bit stream logic for all the individual
0522 characters (e.g., delimiters) and character classes (e.g., digits, letters) used in a particular application.
0523 Input is specified using a character class syntax adapted from the standard regular expression notations.
0524 Output is a minimized set of three-address bitwise operations, such as $a = b \& c$, to compute each of the
0525 character classes from the basis bit streams.
0526
0527
0528
0529
0530

0531 For example, Figure 5 shows the input and output produced by the character class compiler for the
0532 example of $[0-9]$ discussed in the previous section. The output operations may be viewed as operations
0533 on a single block of input at a time, or may be viewed as operations on unbounded bitstreams as supported
0534 by the Pablo compiler.
0535
0536
0537
0538

```
0539     INPUT:  digit = [0-9]  
0540     OUTPUT: temp1 = (basis_bits.bit_0 | basis_bits.bit_1)  
0541                temp2 = (basis_bits.bit_2 & basis_bits.bit_3)  
0542                temp3 = (temp2 &~ temp1)  
0543                temp4 = (basis_bits.bit_5 | basis_bits.bit_6)  
0544                temp5 = (basis_bits.bit_4 & temp4)  
0545                digit = (temp3 &~ temp5)  
0546  
0547  
0548
```

0549 Figure 5: Character Class Compiler Input/Output

0550
0551
0552 Pablo is a compiler that abstracts away the details of programming parallel bit stream code in terms
0553 of finite SIMD register widths and application buffer sizes. Input to Pablo is a language for expressing
0554 bitstream operations on unbounded bitstreams. The operations include bitwise logic, the `Advance` and
0555 `ScanThru` operations described in the previous subsection as well as `if` and `while` control structures. Pablo
0556
0557
0558
0559

0560
0561
0562
0563
0564
0565
0566
0567
0568
0569
0570
0571
0572
0573
0574
0575
0576
0577
0578
0579
0580
0581
0582
0583
0584
0585
0586
0587
0588
0589
0590
0591
0592
0593
0594
0595
0596
0597
0598
0599
0600
0601
0602
0603
0604
0605
0606
0607
0608
0609
0610
0611
0612
0613
0614
0615

translates these operations to block-at-a-time code in C/C++, where the block size is the register width for SIMD operations on the selected target architecture. The key functionality of Pablo is to arrange for block-to-block carry bit propagation to implement the long bitstream shift and addition operations required by Advance and ScanThru.

3.3 Parabix Run-Time Libraries

The Parabix architecture also includes run-time libraries that support a machine-independent view of basic SIMD operations, as well as a set of core function libraries. For machine-independence, we program all operations using an abstract SIMD machine based on the Inductive Doubling Instruction Set Architecture (IDISA) [10]. Originally developed for 128-bit AltiVec operations on Power PC as well as 64-bit MMX and 128-bit SSE operations on Intel, we have recently extended our library support to include the new 256-bit AVX operations on Intel as well as the 128-bit NEON operations on the ARM architecture. Further details are provided in later sections.

4 The Parabix XML Parser

4.1 Parser Structure

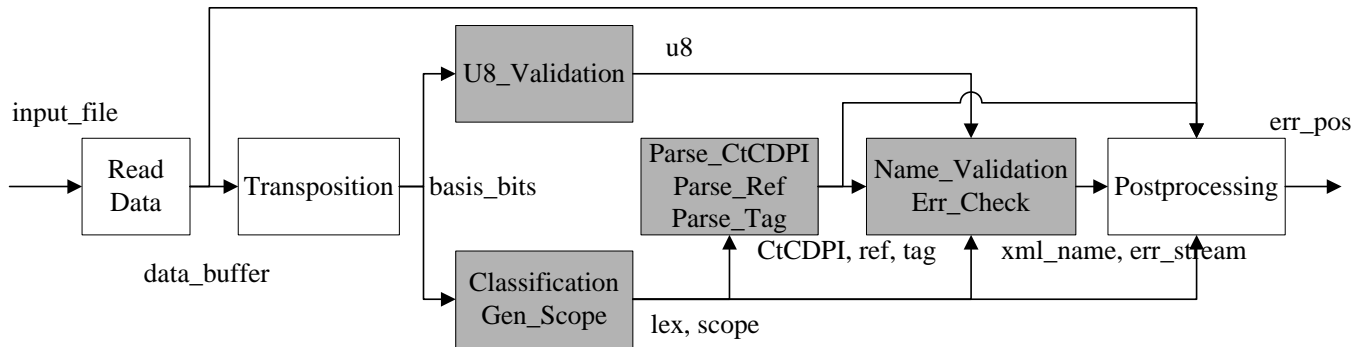


Figure 6: Parabix2 Structure

Figure 6 shows the overall structure of the Parabix XML parser set up for well-formedness checking. The input file is processed using 11 functions organized into 7 modules. In the first module, the Read_Data function loads data blocks from an input file to data_buffer. The data is then transposed to eight parallel basis bitstreams (basis_bits) in the Transposition module. The eight bitstreams are used in the Classification function to generate all the XML lexical item streams (lex) as well as in the U8_Validation module

#****

HPCA 2012 Submission #****. CONFIDENTIAL REVIEW COPY. DO NOT DISTRIBUTE.

#****

0616 to validate UTF-8 characters. The lexical item streams and scope streams (scope) that are generated
0617 in Gen_Scope function are supplied to the parsing module, which consists three functions, Parse_CtCDPI,
0618 Parse_Ref and Parse_tag. These functions deal with the parsing of comments, CDATA sections, processing
0619 instructions, references and tags. After this, information is gathered by Name_Validation and Err_Check
0620 functions, producing name check streams and error streams. These are then passed to the final module for
0621 Postprocessing. All the possible errors that cannot be conveniently detected by bitstreams are checked in
0622 this last module. The final output reports any well-formedness error detected and its position within the
0623 input file.
0624
0625
0626
0627
0628
0629
0630

0631 Within this structure, all functions in the four shaded modules consist entirely of parallel bit stream
0632 operations. Of these, the Classification function consists of XML character class definitions that are gen-
0633 erated using ccc, while much of the U8_Validation similarly consists of UTF-8 byte class definitions that
0634 are also generated by ccc. The remainder of these functions are programmed using our unbounded bit-
0635 stream language following the logical requirements of XML parsing. All the functions in the four shaded
0636 modules are then compiled to low-level C/C++ code using our Pablo compiler. This code is then linked in
0637 with the general Transposition code available in the Parabix run-time library, as well as the hand-written
0638 Postprocessing code that completes the well-formed checking.
0639
0640
0641
0642
0643
0644
0645
0646
0647

0648 5 Methodology

0649 In this section we describe our methodology for the measurements and investigation of XML parser
0650 energy consumption and performance. In brief, for each of the four XML parsers under study we propose
0651 to measure and evaluate the energy consumption required to carry out XML well-formedness checking,
0652 under a variety of workloads, and as executed on three different Intel processors.
0653
0654
0655
0656
0657
0658
0659

0660 To begin our study we propose to first investigate each of the XML parsers in terms of the Performance
0661 Monitoring Counter² (PMC) hardware events listed in the PMC Hardware Events subsection. Based on
0662 the findings of previous work [2–4] we have chosen several key hardware performance events for which
0663 the authors indicate a strong correlation with energy consumption. In addition, we measure the runtime
0664
0665
0666
0667

0668 ²Performance Monitoring Counters are special-purpose registers available with most modern microprocessors. PMCs store
0669 the running count of specific hardware events, such as retired instructions, cache misses, branch mispredictions, and arithmetic-
0670 logic unit operations. PMCs can be used to capture information about any program at run-time and under any workload at a
0671 fine granularity.

0672 counts of SIMD instructions and bitwise operations using the Intel Pin binary instrumentation framework.
0673
0674 Based on these data we gain further insight into XML parser execution characteristics and compare and
0675
0676 contrast each of the Parabix parser versions against the performance of standard industry parsers.
0677

0678 The foundational work by Bellosa in [2] as well as more recent work in [3,4] demonstrate that hardware-
0679
0680 usage patterns have a significant impact on the energy consumption characteristics of an application [2–4].
0681
0682 Further, the authors demonstrate a strong correlation between specific PMC events and energy usage.
0683
0684 However, each author differs slightly in their opinion of the exact set of PMCs to use.
0685

0686 The following subsections describe the XML parsers under study, XML workloads, the hardware archi-
0687
0688 tectures, PMC hardware events selected for measurement, and the energy measurement instrumentation
0689
0690 set up. We analyze the performance of each of the XML parsers under study based on PMC hardware
0691
0692 event counts and contrast their energy consumption measurements based on direct measurements.
0693
0694

0695 5.1 Parsers

0696
0697 The XML parsing technologies selected for this study are the Parabix1, Parabix2, Xerces-C++, and
0698
0699 Expat XML parsers. Parabix1 (parallel bit Streams for XML) is our first generation SIMD and Parallel Bit
0700
0701 Stream technology based XML parser [15]. Parabix1 leverages the processor built-in *bitscan* operation
0702
0703 for high-performance XML character scanning as well as the SIMD capabilities of modern commodity
0704
0705 processors to achieve high performance. Parabix2 [16] represents the second generation of the Parabix1
0706
0707 parser. Parabix2 is an open-source XML parser that also leverages Parallel Bit Stream technology and the
0708
0709 SIMD capabilities of modern commodity processors. However, Parabix2 differs from Parabix1 in that it
0710
0711 employs new parallelization techniques, such as a multiple cursor approach to parallel parsing together
0712
0713 with bit stream addition techniques to advance multiple cursors independently and in parallel. Parabix2
0714
0715 delivers dramatic performance improvements over traditional byte-at-a-time parsing technology. Xerces-
0716
0717 C++ version 3.1.1 (SAX) [17] is a validating open source XML parser written in C++ by the Apache
0718
0719 project. Expat version 2.0.1 [11] is a non-validating XML parser library written in C.
0720

0721 5.2 Workloads

0722
0723 Markup density is defined as the ratio of the total markup contained within an XML file to the total
0724
0725 XML document size. This metric has substantial influence on the performance of traditional recursive
0726
0727

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

descent XML parser implementations. We use a mixture of document-oriented and data-oriented XML files in our study to provide workloads with a full 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 and contain the three-byte and four-byte UTF-8 sequence required for the UTF-8 encoding of Japanese and German characters respectively. The remaining data files are data-oriented XML documents and consist entirely of single byte 7-bit encoded ASCII characters.

5.3 Platform Hardware

Intel Core2 Intel Core2 processor, code name Conroe, produced by Intel. Table ?? gives the hardware description of the Intel Core2 machine.

Processor	Core2 Duo (2.13GHz)	i3-530 (2.93GHz)	Sandybridge (2.80GHz)
L1 D Cache	32KB	32KB	32KB
L2 Cache	Shared 2MB	256KB/core	256KB/core
L3 Cache	—	4MB	6MB
Bus or QPI	1066Mhz Bus	1333Mhz QPI	1333Mhz QPI
Memory	2GB	4GB	6GB
Max TDP	65W	73W	95W

Table 2: Platform Hardware Specs

Intel Core-i3 processor, code name Nehalem, produced by Intel. The intent of the selection of this processor is to serve as an example of a low end server processor. Table ?? gives the hardware description of the Intel Core-i3 machine. Intel Core-i5 processor, code name SandyBridge produced by Intel. Table ?? gives the hardware description of the Intel Core-i3 machine. Each of the hardware events selected relates to performance and energy features associated with one or more hardware units. For example, total branch mispredictions relate to the branch predictor and branch target buffer capacity.

The set of PMC events used included in this study are as follows. Processor Cycles, Branch Instructions, Branch Mispredictions, Integer Instructions, SIMD Instructions and Cache Misses.




Figure 7: Cache Misses per kB of input data.

5.4 Energy Measurement

We measure energy consumption using the 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 [12].

6 Baseline Evaluation on Core-i3

6.1 Cache behavior

Core-i3 has a three level cache hierarchy. The approximate miss penalty for each cache level is 4, 11, and 36 cycles respectively. Figure 7(a), Figure 7(b) and Figure ?? show the L1, L2 and L3 data cache misses for each of the parsers. Although XML parsing is non memory intensive application, cache misses for the Expat and Xerces parsers represent a 0.5 cycle per XML byte cost whereas the performance of the Parabix parsers remains essentially unaffected by data cache misses. Cache misses not only consume additional CPU cycles but increase application energy consumption. L1, L2, and L3 cache misses consume approximately 8.3nJ, 19nJ, and 40nJ respectively. As such, given a 1GB XML file as input, Expat and Xerces would consume over 0.6J and 0.9J respectively due to cache misses alone.

6.2 Branch Mispredictions

Despite improvements in branch prediction, branch misprediction penalties contribute significantly to XML parsing performance. On modern commodity processors the cost of a single branch misprediction




Figure 8: Branch characteristics on the Core-i3 per kB of input data.

is commonly cited as over 10 CPU cycles. As shown in Figure 8(b), the cost of branch mispredictions for the Expat parser can be over 7 cycles per XML byte—this cost alone is equal to the average total cost for Parabix2 to process each byte of XML.

In general, reducing the branch misprediction rate is difficult in text-based XML parsing applications. This is due in part to the variable length nature of the syntactic elements contained within XML documents, a data dependent characteristic, as well as the extensive set of syntax constraints imposed by the XML 1.0 specification. As such, traditional byte-at-a-time XML parsers generate a performance limiting number of branch mispredictions. As shown in Figure 8(a), Xerces averages up to 13 branches per XML byte processed on high density markup.

The performance improvement of Parabix1 in terms of branch mispredictions results from the veritable elimination of conditional branch instructions in scanning. Leveraging the processor built-in *bit scan* operation together with parallel bit stream technology Parabix1 can scan up to 64 bytes of source XML with a single *bit scan* instruction. In comparison, a byte-at-a-time parser must process a conditional branch instruction per XML byte scanned.

As shown in Figure 8(a), Parabix2 processing is almost branch free. Utilizing a new parallel scanning technique based on bit stream addition, Parabix2 exhibits minimal dependence on source XML markup density. Figure 8(a) displays this lack of data dependence via the constant number of branch mispredictions shown for each of the source XML files.

6.3 SIMD Instructions vs. Total Instructions

Parabix achieves performance via parallel bit stream technology. In Parabix XML processing, parallel bit streams are both computed and predominately operated upon using the SIMD instructions of commodity processors. The ratio of retired SIMD instructions to total instructions provides insight into the relative degree to which Parabix achieves parallelism over the byte-at-a-time approach.

Using the Intel Pin tool, we gather the dynamic instruction mix for each XML workload, and classify instructions as either vector (SIMD) or non-vector instructions. Figures ?? and 9(b) show the percentage of SIMD instructions for Parabix1 and Parabix2 respectively. For Parabix1, 18% to 40% of the executed instructions are SIMD instructions. Using bit stream addition to scan XML characters in parallel, the Parabix2 instruction mix is made up of 60% to 80% SIMD instructions. Although the resulting ratios are (negatively) proportional to the markup density for both Parabix1 and Parabix2, the degradation rate of Parabix2 is much lower and thus the performance penalty incurred by increasing the markup density is reduced.


6.4 CPU Cycles

Figure 9(a) shows overall parser performance evaluated in terms of CPU cycles per kilobyte. Parabix1 is 1.5 to 2.5 times faster on document-oriented input and 2 to 3 times faster on data-oriented input than the Expat and Xerces parsers respectively. Parabix2 is 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 by dense markup, while Parabix2 is generally unaffected. The results presented are not entirely fair to the Xerces parser since it first transcodes input from UTF-8 to UTF-16 before processing. In Xerces, this transcoding requires several cycles per byte. However, transcoding using parallel bit streams is significantly faster and requires less than a single cycle per byte. [7].

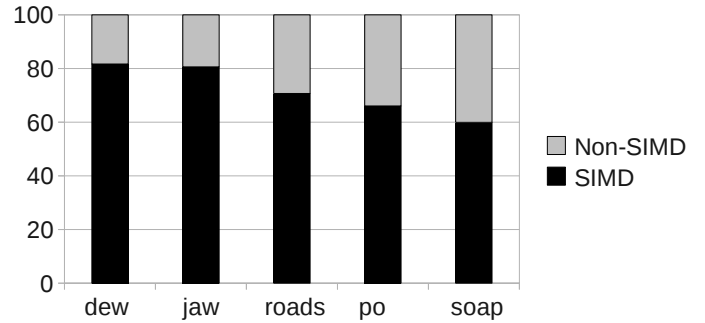
6.5 Power and Energy

In response to the growing industry concerns on power consumption and energy efficiency, chip producers work hard to not only improve performance but also achieve high energy efficiency in processors design. We study the power and energy consumption of Parabix in comparison with Expat and Xerces on Core-i3. The average power of Core-i3 530 is about 21 watts. This Intel model has a good reputation for


0952
0953
0954
0955
0956
0957
0958
0959
0960
0961
0962
0963
0964
0965
0966
0967
0968
0969
0970
0971
0972
0973
0974
0975
0976
0977
0978
0979
0980
0981
0982
0983
0984
0985
0986
0987
0988
0989
0990
0991
0992
0993
0994
0995
0996
0997
0998
0999
1000
1001
1002
1003
1004
1005
1006
1007




(a) Performance : # Cycles/kb



(b) SIMD Instruction Breakdown. Y Axis : % SIMD Instruction/kb



(c) Avg. Power (Watts)

(d) Energy Consumption (μ J per kB)

power efficiency. Figure 9(c) shows the average power consumed by each parser. Parabix2, dominated by SIMD instructions, uses approximately 5% additional power.

As shown in Figure 9(d), a comparison of energy efficiency demonstrates a more interesting result. Although Parabix2 requires slightly more power (per instruction), the processing time of Parabix2 is significantly lower, and therefore Parabix2 consumes substantially less energy than the other parsers. Parabix2 consumes 50 to 75 nJ per byte while Expat and Xerces consume 80nJ to 320nJ and 140nJ to 370nJ per byte respectively.

7 Scalability

7.1 Performance

Figure 9 (a) demonstrates the average XML well-formedness checking performance of Parabix2 for each of the workloads and as executed on each of the processor cores — Core2 Core-i3 and SandyBridge. Processing time is shown in terms of bit stream based operations executed in ‘bit-space’ and postprocessing operations executed in ‘byte-space’. In the Parabix2 parser, bit-space parallel bit stream parser oper-

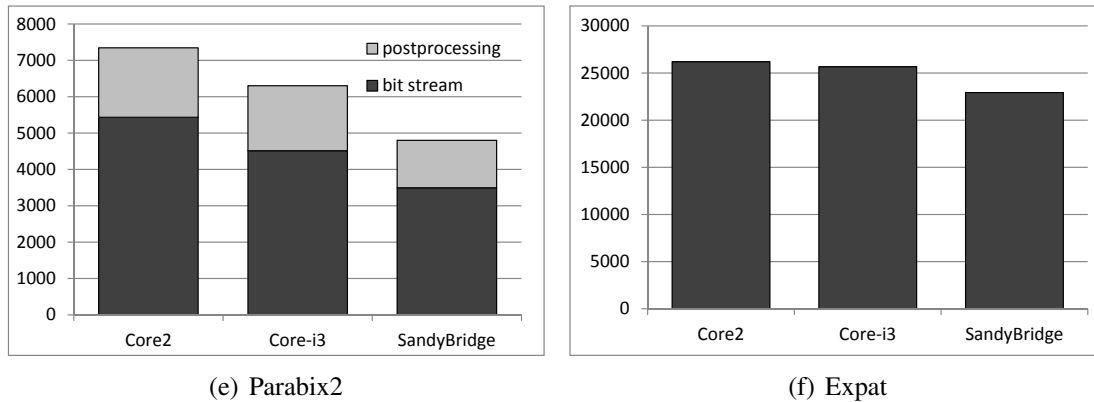


Figure 9: Average Performance Parabix vs. Expat (y-axis: CPU Cycles per kB)

ations consist primarily of SIMD instructions; byte-space operations consist of byte comparisons across arrays of values. Executing Parabix2 on Core-i3 over Core2 results in an average performance improvement of 17% in bit stream processing whereas migrating Parabix2 from Core-i3 to SandyBridge results in a 22% average performance gain. Bit space measurements are stable and consistent across each of the source inputs and cores. Postprocessing operations demonstrate data dependent variance. Performance gains from 18% to 31% performance are observed in migrating Parabix2 from Core2 to Core-i3; 0% to 17% performance from Core-i3 to SandyBridge. For the purpose of comparison, Figure 9 (b) shows the performance of the Expat parser on each of the processor cores. A performance improvement of less than 5% is observed when executing Expat on Core-i3 over Core2 and less than 10% on SandyBridge over Core-i3.

Overall, Parabix2 scales better than Expat. Simply executing identical Parabix2 object code on SandyBridge results in an overall performance improvement up to 26%. Additional performance aspects of Parabix2 on SandyBridge with AVX instructions are discussed in the following sections.

7.2 Power and Energy

Figure 10(a) shows the average power consumption of Parabix2 over each workload and as executed on each of the processor cores — Core2, Core-i3 and SandyBridge. Average power consumption on Core2 is 32 watts. Execution on Core-i3 results in 30% power saving over Core2. SandyBridge saves 25% of the power compared with Core-i3 and consumes only 15 watts.

In XML parsing we observe energy consumption is dependent on processing time. That is, a reduction in

1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119

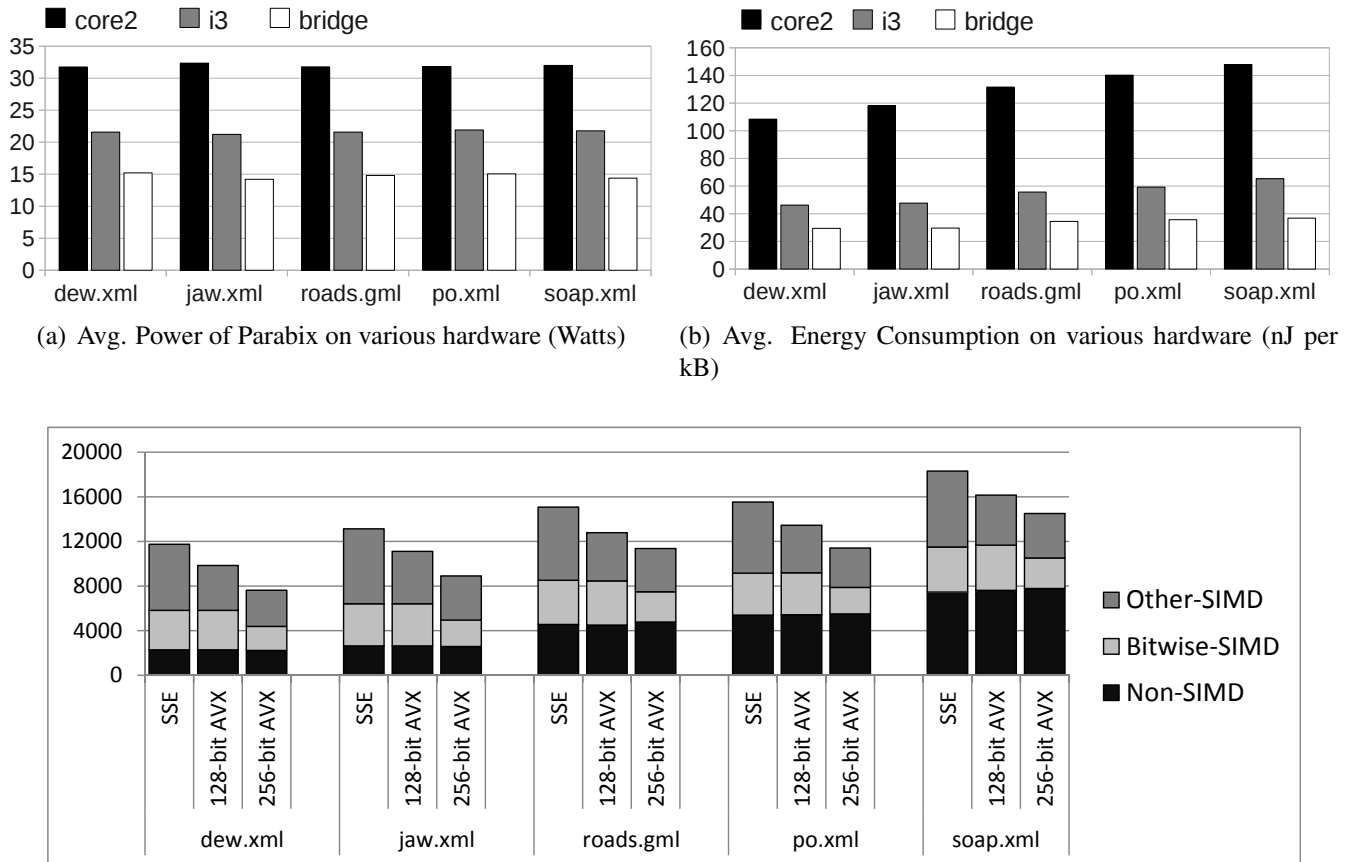


Figure 10: Parabix2 Instruction Counts (y-axis: Instructions per kB)

processing time results in a directly proportional reduction in energy consumption. With newer processor cores comes improvements in application performance. As a result, Parabix2 executed on SandyBridge consumes 72% to 75% less energy than Parabix2 on Core2.

8 Scaling Parabix2 for AVX

In this section, we discuss the scalability and performance advantages of our 256-bit AVX (Advanced Vector Extensions) Parabix2 port. Parabix2 originally targetted the 128-bit SSE2 SIMD technology available on all modern 64-bit Intel and AMD processors but has recently been ported to AVX. AVX technology is commercially available on the latest the SandyBridge microarchitecture Intel processors.

8.1 Three Operand Form

In addition to the widening of 128-bit operations to 256-bit operations, AVX technology uses a non-destructive 3-operand instruction format. Previous SSE implementations used a destructive 2-operand

1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175




Figure 11: Parabix2 Performance (y-axis: CPU Cycles per kB)

instruction format. In the 2-operand format a single register is used as both a source and destination register. For example, $a = a [op] b$. As such, 2-operand instructions that require the value of both a and b , must either copy an additional register value beforehand, or reconstitute or reload a register value afterwards to recover the value. With the 3-operand format, output may now be directed to the third register independently of the source operands. For example, $c = a [op] b$. By avoiding the copying or reconstituting of operand values, a considerable reduction in instruction count in the form of reduced load and store instructions is possible. AVX technology makes available the 3-operand form for both the new 256-bit operations as well as the base 128-bit SSE operations.

8.2 256-bit AVX Operations

With the introduction of 256-bit SIMD registers, and under ideal conditions, one would anticipate a corresponding 50% reduction in the SIMD instruction count of Parabix2 on AVX. However, in the Sandy-Bridge AVX implementation, Intel has focused primarily on floating point operations as opposed to the integer based operations. 256-bit SIMD is available for loads, stores, bitwise logic and floating operations, whereas SIMD integer operations and shifts are only available in the 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 but not achieved.

8.3 Performance Results

We implemented two versions of Parabix2 using AVX technology. The first was simply the recompilation of the existing Parabix2 source code written to take advantage of the 3-operand form of AVX

#****

#****

1176 instructions while retaining a uniform 128-bit SIMD processing width. The second involved rewriting
1177 the core library functions of Parabix2 to leverage the 256-bit AVX operations wherever possible and to
1178 simulate the remaining operations using pairs of 128-bit operations.
1179
1180

1182 Figure 10 shows the reduction in instruction counts achieved in these two versions. For each workload,
1183 the base instruction count of the Parabix2 binary compiled in SSE-only mode is shown with the caption
1184 “sse,” the version obtained by simple recompilation with AVX-mode enabled is labeled “128-bit avx,”
1185 and the version reimplemented to use 256-bit operations wherever possible is labelled “256-bit avx.” The
1186 instruction counts are divided into three classes. The “non-SIMD” operations are the general purpose
1187 instructions that use neither SSE nor AVX technology. The “bitwise SIMD” class comprises the bitwise
1188 logic operations, that are available in both 128-bit form and 256-bit form. The “other SIMD” class com-
1189 prises all other SIMD operations, primarily comprising the integer SIMD operations that are available only
1190 at 128-bit widths even with 256-bit AVX technology.
1191
1192
1193
1194
1195
1196
1197
1198
1199

1200 Note that, in each workload, the number of non-SIMD instructions remains relatively constant with each
1201 workload. As may be expected, however, the number of “bitwise SIMD” operations remains the same for
1202 both SSE and 128-bit while dropping dramatically when operating 256-bits at a time. Ideally one one may
1203 expect up to a 50% reduction in these instructions versus the 128-bit AVX. The actual reduction measured
1204 was 32%–39% depending on workload. Because some bitwise logic is needed in implementation of
1205 simulated 256-bit operations, the full 50% reduction in bitwise logic was not achieved.
1206
1207
1208
1209
1210
1211

1212 The “other SIMD” class shows a substantial 30%-35% reduction with AVX 128-bit technology com-
1213 pared to SSE. This reduction is due to eliminated copies or reloads when SIMD operations are compiled
1214 using 3-operand AVX form versus 2-operand SSE form. A further 10%–20% reduction is observed with
1215 Parabix2 version rewritten to use 256-bit operations.
1216
1217
1218
1219

1220 While the successive reductions in SIMD instruction counts are quite dramatic with the two AVX im-
1221 plementations of Parabix2, the performance benefits are another story. As shown in Figure 11, the benefits
1222 of the reduced SIMD instruction count are achieved only in the AVX 128-bit version. In this case, the ben-
1223 efits of 3-operand form seem to fully translate to performance benefits. Based on the reduction of overall
1224 Bitwise-SIMD instructions we expected a 11% improvement in performance. Instead, perhaps bizzarely,
1225 the performance of Parabix2 in the 256-bit AVX implementation does not improve significantly and actu-
1226
1227
1228
1229
1230
1231

ally degrades for files with higher markup density (average 10%). Dewiki.xml, on which bitwise-SIMD instructions reduced by 39%, saw a performance improvement of 8%. We believe that this is primarily due to the intricacies of the first generation AVX implementation in SandyBridge, with significant latency in many of the 256-bit instructions in comparison to their 128-bit counterparts. The 256-bit instructions also have different scheduling constraints that seem to reduce overall SIMD throughput. 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.

9 Parabix on Mobile Platforms


The Samsung Galaxy Tab GT-P1000M device houses a Samsung S5PC110 ARM Cortex-A8 1Ghz single-core, dual-issue, superscalar microprocessor. It includes a 32kB L1 data cache and a 512kB L2 shared cache. In addition to the standard feature set found in such low-power 32-bit microprocessors, the S5PC110 includes the ARM NEON general-purpose SIMD engine. ARM NEON makes available a 128-bit SIMD instruction set similar in functionality to Intel SSE3 instruction set. In this section, we present our performance comparison of a NEON-based port of Parabix2 versus the Expat parser, and executed on the Samsung Galaxy Tab GT-P1000M hardware. Xerces is excluded from this portion of our study due to the complexity of the cross-platform build process in porting native C/C++ applications to the Android platform.

9.1 Performance Results

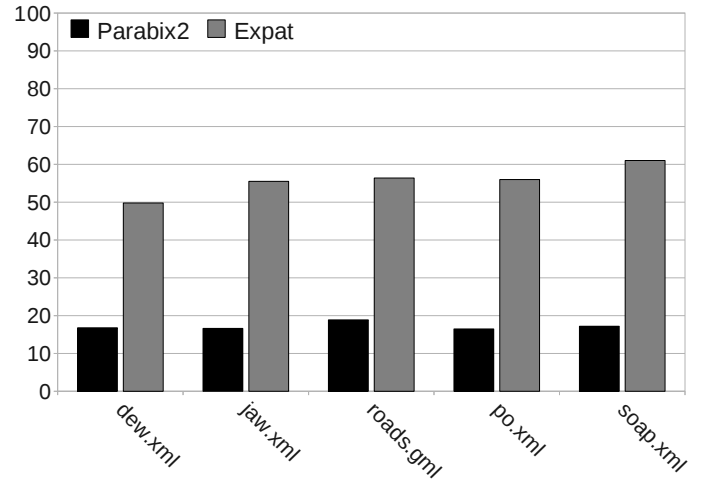
Migration of Parabix2 to the Android platform began with the retargetting of a subset of the Parabix2 IDISA SIMD library for ARM NEON. This library code was cross-compiled for Android using the Android NDK. The Android NDK is a companion tool to the Android SDK that allows developers to build performance-critical portions of applications in native code. The majority of the Parabix2 SIMD functionality ported directly. However, for a small subset of the SIMD functions of Parabix2 NEON equivalents did not exist. In such cases we simply simulated logical equivalencies using the available the instruction set.

A comparison of Figure 12(a) and Figure 9(a) demonstrates that the performance of both Parabix2 and Expat degrades substantially on Cortex-A8. This result was expected given the comparably performance

1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343



(a) ARM Neon Performance



(b) Performance ARM Neon vs Core i3 SSE.

limited Cortex-A8 hardware architecture. Surprisingly on Cortex-A8 Expat outperforms Parabix2 on each of the lower markup density workloads, dew.xml and jaw.xml. On the remaining higher-density workloads, Parabix2 performs only moderately better than Expat. The higher latency of the NEON instructions on Cortex-A8 is the likely contributor to this loss in performance. A more interesting aspect of this result is demonstrated in a comparison of Figure 12(b) and Figure 12(b). These figure demonstrate that the relative performance of each parser degrades in a relatively constant manner. That is, compared to the Core-i3, on the GT-P1000M, Parabix2 and Expat operate at approximately 17.2% and 55.7% efficiency respectively. Figure 12(b) shows that the baseline cost of Parabix2 operations implemented using the NEON instruction set—and thereby the baseline cost of Parabix2—is substantially higher on the Cortex-A8 processor. Given that Parabix2 was not designed with the limitations of the Cortex-A8 in mind, in the future a careful analysis of the cost of each instruction provided in the ARMv7 ISA may allow us to better utilize the hardware resources provided. In particular, future performance enhancement to ARM NEON could result in substantial overall improvement in Parabix2 execution time.

10 Multi-threaded Parabix

The general problem of addressing performance through multicore parallelism is the increasing energy cost. As discussed in previous sections, Parabix, which applies SIMD-based techniques can not only achieves better performance but consumes less energy. Moreover, using multiple cores, we can further improve the performance of Parabix while keeping the energy consumption at the same level.

1344 A typical approach to parallelizing software, data parallelism, requires nearly independent data, How-
1345 ever, the nature of XML files makes them hard to partition nicely for data parallelism. Several approaches
1346 have been used to address this problem. A preparsing phase has been proposed to help partition the XML
1347 document [19]. The goal of this preparsing is to determine the tree structure of the XML document so
1348 that it can be used to guide the full parsing in the next phase. Another data parallel algorithm is called
1349 ParDOM [20]. It first builds partial DOM node tree structures for each data segments and then link them
1350 using preorder numbers that has been assigned to each start element to determine the ordering among
1351 siblings and a stack to manage the parent-child relationship between elements.
1352
1353
1354
1355
1356
1357
1358
1359

1360 Data parallelism approaches introduce a lot of overheads to solve the data dependencies between seg-
1361 ments. Therefore, instead of partitioning the data into segments and assigning different data segments to
1362 different cores, we propose a pipeline parallelism strategy that partitions the process into several stages
1363 and let each core work with one single stage.
1364
1365
1366
1367

1368 The interface between stages is implemented using a circular array, where each entry consists of all ten
1369 data structures for one segment as listed in Table 3. Each thread keeps an index of the array (I_N), which is
1370 compared with the index (I_{N-1}) kept by its previous thread before processing the segment. If I_N is smaller
1371 than I_{N-1} , thread N can start processing segment I_N , otherwise the thread keeps reading I_{N-1} until I_{N-1}
1372 is larger than I_N . The time consumed by continuously loading the value of I_{N-1} and comparing it with I_N
1373 will be later referred as stall time. When a thread finishes processing the segment, it increases the index
1374 by one.
1375
1376
1377
1378
1379
1380
1381

1382 Figure 12 demonstrates the XML well-formedness checking performance of the multi-threaded Parabix
1383 in comparison with the single-threaded version. The multi-threaded Parabix is more than two times faster
1384 and runs at 2.7 cycles per input byte on the SandyBridge machine.
1385
1386
1387

1388 Figure 12(d) shows the average power consumed by the multi-threaded Parabix in comparison with
1389 the single-threaded version. By running four threads and using all the cores at the same time, the power
1390 consumption of the multi-threaded Parabix is much higher than the single-threaded version. However, the
1391 energy consumption is about the same, because the multi-threaded Parabix needs less processing time. In
1392 fact, as shown in Figure 12(e), parsing soap.xml using multi-threaded Parabix consumes less energy than
1393 using single-threaded Parabix.
1394
1395
1396
1397
1398
1399

1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455

		Data Structures									
		data_buffer	basis_bits	u8	lex	scope	ctCDPI	ref	tag	xml_names	err_streams
size		128	128	496	448	80	176	112	176	16	112
Stage1	read_data transposition classification	write read	write read		write						
Stage2	validate_u8 gen_scope parse_CtCDPI parse_ref		read	write	read read read	write read read	write read	write			write
Stage3	parse_tag validate_name gen_check			read read	read read read	read read read	read read read	read	write read read	write read	write write
Stage4	postprocessing	read			read		read	read			read

Table 3: Relationship between Each Pass and Data Structures




Figure 12: Multithreaded Parabix

11 Conclusion

In this paper we presented Parabix a software runtime framework for exploiting SIMD data units found on commodity processors for text processing. The Parabix framework allows to focus on exposing the parallelism in their application assuming an infinite resource abstract SIMD machine without worrying about or having to change code to handle processor specifics (e.g., 128 bit SIMD SSE vs 256 bit SIMD on AVX). We applied Parabix technology to a widely deployed application; XML parsing and demonstrate the efficiency gains that can be obtained on commodity processors. Compared to the conventional XML parsers, Expat and Xerces, we achieve $2\times$ — $7\times$ improvement in performance and average $x\times$ improvement in energy. We achieve high compute efficiency with an overall $?\times$ reduction in branches, $?\times$ reduction in branch mispredictions, $?$ processing upto 128 characters with a single operation. We used the Parabix framework and XML parsers to study the features of the new 256 bit AVX extension in Intel processors. We find that while the move to 3-operand instructions deliver significant benefit the wider operations in some cases have higher overheads compared to the existing 128 bit SSE operations. We

also compare Intel's SIMD extensions against the ARM Neon. Note that Parabix allowed us to perform these studies without having to change the application source. Finally, we parallelized the Parabix XML parser to take advantage of the SIMD units in every core on the chip. We demonstrate that the benefits of thread-level-parallelism are complementary to the fine-grain parallelism we exploit; parallelized Parabix achieves a further $2\times$ improvement in performance.

References

- [1] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick. The landscape of parallel computing research: A view from Berkeley. Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, Dec 2006. 2
- [2] F. Bellosa. The case for event-driven energy accounting. Technical Report TR-I4-01-07, University of Erlangen, Department of Computer Science, June 2001. 12, 13
- [3] 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. 12, 13
- [4] 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. 12, 13
- [5] G. Blake, R. G. Dreslinski, T. Mudge, and K. Flautner. Evolution of thread-level parallelism in desktop applications. In *Proceedings of the 37th annual international symposium on Computer architecture*, ISCA '10, 2010. 2
- [6] 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
- [7] 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. 17
- [8] 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. 6
- [9] 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. 6, 7
- [10] 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. 11
- [11] J. Clark. The Expat XML Parser. <http://expat.sourceforge.net/>. 6, 13
- [12] F. Corporation. Fluke Clamp Meters. <http://www.fluke.com/>. 15
- [13] B. DuCharme. Documents vs. data, schemas vs. schemas. In *XML 2004*, Washington D.C., 2004. 5
- [14] H. Esmailzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger. Dark silicon and the end of multicore scaling. In *Proceeding of the 38th annual international symposium on Computer architecture*, ISCA '11, 2011. 2
- [15] R. D. C. et. al. Parabix1. <http://parabix.costar.sfu.ca/>. 13
- [16] R. D. C. et. al. Parabix2. <http://parabix.costar.sfu.ca/>. 13
- [17] A. S. Foundation. Xerces C++ Parser. <http://xerces.apache.org/xerces-c/>. 6, 13
- [18] R. Hameed, W. Qadeer, M. Wachs, O. Azizi, A. Solomatnikov, B. C. Lee, S. Richardson, C. Kozyrakis, and M. Horowitz. Understanding sources of inefficiency in general-purpose chips. In *Proceedings of the 37th annual international symposium on Computer architecture*, ISCA '10, 2010. 2
- [19] W. Lu, Y. Pan, , and K. Chiu. A parallel approach to xml parsing. *The 7th IEEE/ACM International Conference on Grid Computing*, 2006. 25

1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567

- [20] B. Shah, P. R. Rao, B. Moon, and M. Rajagopalan. A data parallel algorithm for xml dom parsing. In *Proceedings of the 6th International XML Database Symposium on Database and XML Technologies, XSym '09*, pages 75–90, Berlin, Heidelberg, 2009. Springer-Verlag. 25
- [21] G. Venkatesh, J. Sampson, N. Goulding, S. Garcia, V. Bryksin, J. Lugo-Martinez, S. Swanson, and M. B. Taylor. Conservation cores: reducing the energy of mature computations. In *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems, ASPLOS '10*, 2010. 2
- [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. 6