

Bitwise Data Parallelism in Regular Expression Matching

Anonymous Authors

Institutions

emails

Abstract

New parallel algorithms for the classical grep (global regular expression print) problem are introduced together with implementations using commodity SIMD and GPU technologies. Building on the bitwise data parallelism underlying previous work in Unicode transcoding and XML parsing, the new algorithms add the important element of nondeterminism for tackling the full generality of regular expression processing. On widely-deployed commodity hardware using 128-bit SSE2 SIMD technology, our algorithm implementations can substantially outperform traditional grep implementations based on NFAs, DFAs or backtracking. The algorithms are also designed to scale with the availability of additional parallel resources such as the wider SIMD facilities (256-bit) of Intel AVX2 or future 512-bit extensions. Our GPU implementations show further acceleration limited only by data transfer speed.

Categories and Subject Descriptors Theory of computation [*Formal languages and automata theory*]: Regular languages; Computer systems organization [*Parallel architectures*]: Single instruction, multiple data

Keywords regular expression matching, grep, parallel bit stream technology

1. Introduction

The use of regular expressions to search texts for occurrences of string patterns has a long history and remains a pervasive technique throughout computing applications today. The origins of regular expression matching date back to automata theory developed by Kleene in the 1950s [10]. Thompson [23] is credited with the first construction to convert regular expressions to nondeterministic finite automata (NFA). Following Thompson's approach, a regular expression of length m is first converted to an NFA with $O(m)$ nodes. It is then possible to search a text of length n using the NFA in worst case $O(mn)$ time. Often, a more efficient choice is to convert the NFA into a DFA. A DFA has only a single active state at any time in the matching process and hence it is possible to search a text of length n in $O(n)$ time. However, it is well known that the conversion of an NFA to an equivalent DFA may result in state explosion. That is, the number of resultant DFA states may increase exponentially. In [3] a new approach to text searching was proposed

based on bit-parallelism [2]. This technique takes advantage of the intrinsic parallelism of bitwise operations within a computer word. Given a w -bit word, the Shift-Or algorithm [3] algorithm uses the bit-parallel approach to simulate an NFA in $O(nm/w)$ worst-case time.

A disadvantage of the bit-parallel Shift-Or pattern matching approach in comparison to simple string matching algorithms is an inability to skip input characters. For example, the Boyer-Moore family of algorithms [5] skip input characters to achieve sublinear times in the average case. Backward Dawg Matching (BDM) string matching algorithms [8] based on suffix automata are able to skip characters. The Backward Nondeterministic Dawg Matching (BNDM) pattern matching algorithm [26] combines the bit-parallel advantages of Shift-Or and with the character skipping advantages of the BDM algorithm. The nrgrep pattern matching tool is built over the BNDM algorithm, and hence the name nrgrep [14].

There has been considerable interest in using parallelization techniques to improve the performance of regular expression matching on parallel hardware such as multi-core processors (CPUs), graphics processing units (GPUs), field-programmable gate arrays (FPGAs), and specialized architectures such as the Cell Broadband Engine (Cell BE). Scarpazza and Braudaway [22] demonstrated that text processing algorithms that exhibit irregular memory access patterns can be efficiently executed on multi-core hardware. In related work, Pasetto et al. presented a flexible tool that performs small-ruleset regular expression matching at a rate of 2.88 Gbps per chip on Intel Xeon E5472 hardware [16]. Naghmouchi et al. demonstrated that the Aho-Corasick (AC) string matching algorithm [1] is well suited for parallel implementation on multi-core CPUs, GPUs and the Cell BE [13, 19]. On each hardware, both thread-level parallelism (additional cores) and data-level parallelism (wide SIMD units) are leveraged for performance. Salapura et al., advocated the use of vector-style processing for regular expressions in business analytics applications and leveraged the SIMD hardware available on multi-core processors to achieve a speedup of better than 1.8 over a range of data sizes of interest [17]. In [20], Scarpazza and Russell presented a SIMD tokenizer that delivered 1.001.78 Gbps on a single Cell BE chip and extended this approach for emulation on the Intel Larrabee instruction set [18]. On the Cell BE, Scarpazza [21] described a pattern matching implementation that delivered a throughput of 40 Gbps for a small dictionary of approximately 100 patterns, and a throughput of 3.3-3.4 Gbps for a larger dictionary of thousands of patterns. Iorio and van Lunteren [9] presented a string matching implementation for automata that achieves 4 Gbps on the Cell BE. In more recent work, Tumeo et al. [24] presented a chunk-based implementation of the AC algorithm for accelerating string matching on GPUs. Lin et al., proposed the Parallel Failureless Aho-Corasick (PFAC) algorithm to accelerate pattern matching on GPU hardware and achieved 143 Gbps throughput, 14.74 times faster than the AC algorithm performed on a four core multi-core processor using OpenMP [11].

[Copyright notice will appear here once 'preprint' option is removed.]

Whereas the existing approaches to parallelization have been based on adapting traditional sequential algorithms to emergent parallel architectures, we introduce both a new algorithmic approach and its implementation on SIMD and GPU architectures. This approach relies on a bitwise data parallel view of text streams as well as a surprising use of addition to match runs of characters in a single step. The closest previous work is that underlying bit-parallel XML parsing using 128-bit SSE2 SIMD technology together with a parallel scanning primitive also based on addition [7]. However, in contrast to the deterministic, longest-match scanning associated with the ScanThru primitive of that work, we introduce here a new primitive MatchStar that can be used in full generality for nondeterministic regular expression matching. We also introduce a long-stream addition technique involving a further application of MatchStar that enables us to scale the technique to n -bit addition in $\lceil \log_{64} n \rceil$ steps. We ultimately apply this technique, for example, to perform synchronized 4096-bit addition on GPU warps of 64 threads.

There is also a strong keyword match between the bit-parallel data streams used in our approach and the bit-parallelism used for NFA state transitions in the classical algorithms of Wu and Manber [25], Baez-Yates and Gonnet [2] and Navarro and Raffinot [15]. However those algorithms use bit-parallelism in a fundamentally different way: representing all possible current NFA states as a bit vector and performing parallel transitions to a new set of states using table lookups and bitwise logic. Whereas our approach can match multiple characters per step, bit-parallel NFA algorithms proceed through the input one byte at a time. Nevertheless, the `agrep` [25] and `ngrep` [14] programs implemented using these techniques remain among the strongest competitors in regular expression matching performance, so we include them in our comparative evaluation.

The remainder of this paper is organized as follows. Section 2 briefly describes regular expression notation and the `grep` problem. Section 3 presents our basic algorithm and MatchStar using a model of arbitrary-length bit-parallel data streams. Section 4 discusses the block-by-block implementation of our techniques including the long stream addition techniques for 256-bit addition with AVX2 and 4096-bit additions with GPGPU SIMT. Section 5 describes our overall SSE2 implementation and carries out a performance study in comparison with existing `grep` implementations. Given the dramatic variation in `grep` performance across different implementation techniques, expressions and data sets, Section 6 considers a comparison between the bit-stream and NFA approaches from a theoretical perspective. Section 7 then examines and demonstrates the scalability of our bitwise data-parallel approach in moving from 128-bit to 256-bit SIMD on Intel Haswell architecture. To further investigate scalability, Section 8 addresses the implementation of our matcher using groups of 64 threads working together SIMT-style on a GPGPU system. Section 10 concludes the paper with a discussion of results and areas for future work.

2. Regular Expression Notation and Grep

We follow common Posix notation for regular expressions. A regular expression specifies a set of strings through a pattern notation. Individual characters normally stand for themselves, unless they are one of the special characters `*+?[\(\|~$.` that serve as metacharacters of the notation system. Thus the regular expression `cat` is a pattern for the set consisting of the single 3-character string “`cat`”. The special characters must be escaped with a backslash to prevent interpretation as metacharacter, thus `\$` represents the dollar-sign and `\\` represent the string consisting of two backslash characters. Character class bracket expressions are pattern elements that allow any character in a given class to be used in a particular con-

text. For example, `[@#%]` is a regular expression that stands for any of the three given symbols. Contiguous ranges of characters may be specified using hyphens; for example `[0-9]` for digits and `[A-Za-z0-9_]` for any alphanumeric character or underscore. If the caret character immediately follows the opening bracket, the class is negated, thus `[^0-9]` stands for any character except a digit. The period metacharacter `.` stands for the class of all characters.

Consecutive pattern elements stand for strings formed by concatenation, thus `[cd][ao][tg]` stands for the set of 8 three-letter strings “`cat`” through “`dog`”. The alternation operator `|` allows a pattern to be defined to have to alternative forms, thus `cat|dog` matches either “`cat`” or “`dog`”. Concatenation takes precedence over alternation, but parenthesis may be used to change this, thus `(ab|cd)[0-9]` stands for any digit following one of the two prefixes “`ab`” or “`cd`”.

Repetition operators may be appended to a pattern to specify a variable number of occurrences of that pattern. The Kleene `*` specifies zero-or-more occurrences matching the previous pattern, while `+` specifies one-or more occurrences. Thus `[a-z][a-z]*` and `[a-z]+` both specify the same set: strings of at least one lower-case letter. The postfix operator `?` specifies an optional component, i.e., zero-or-one occurrence of strings matching the element. Specific bounds may be given within braces: `(ab){3}` specifies the string “`ababab`”, `[0-9A-Fa-f]{2,4}` specifies strings of two, three or four hexadecimal digits, and `[A-Z]{4,}` specifies strings of at least 4 consecutive capital letters.

The `grep` program searches a file for lines containing matches to a regular expression using any of the above notations. In addition, the pattern elements `^` and `$` may be used to match respectively the beginning or the end of a line. In line-based tools such as `grep`, `.` matches any character except newlines; matches cannot extend over lines. Normally, `grep` prints all matching lines to its output. However, `grep` programs typically allow a command line flag such as `-c` to specify that only a count of matching lines be produced; we use this option in our experimental evaluation to focus our comparisons on the performance of the underlying matching algorithms.

3. Matching with Bit-Parallel Data Streams

Whereas the traditional approaches to regular expression matching using NFAs, DFAs or backtracking all rely on a byte-at-a-time processing model, the approach we introduce in this paper is based on quite a different concept: a data-parallel approach to simultaneous processing of data stream elements. Indeed, our most abstract model is that of unbounded data parallelism: processing all elements of the input data stream simultaneously. In essence, data streams are viewed as (very large) integers. The fundamental operations are bitwise logic, stream shifting and long-stream addition.

Depending on the available parallel processing resources, an actual implementation may divide an input stream into blocks and process the blocks sequentially. Within each block all elements of the input stream are processed together, relying the availability of bitwise logic and addition scaled to the block size. On commodity Intel and AMD processors with 128-bit SIMD capabilities (SSE2), we typically process input streams 128 bytes at a time. In this case, we rely on the Parabix tool chain [12] to handle the details of compilation to block-by-block processing. For our GPGPU implementation, we have developed a long-stream addition technique that allows us to perform 4096-bit additions using 64 threads working in lock-step SIMT fashion. Using scripts to modify the output of the Parabix tools, we effectively divide the input into blocks of 4K bytes processed in a fully data-parallel manner.

A key concept in this streaming approach is the derivation of bit streams that are parallel to the input data stream, i.e., in one-

current input marker stream as input to each of them. The final marker streams of the compiled forms of each subexpression are then just combined using a bitwise-or to produce the overall final marker stream of the alternation. That is, a match occurs at any position that can be reached by matching any one of the alternatives.

- If we have a marker stream formed as a concatenation of subexpressions, then we compile each of these in turn, providing the output marker stream of each compilation as the input marker stream for the compilation of the next pattern element.
- If a regular expression is a character class expression or a single character, then we form the bitwise-and of the current marker stream and the character class stream to filter out current marker positions that do not have a match to the class. The result is then shifted forward one position to identify the successful matches.
- If a regular expression is an optional expression of the form $R?$ for some subexpression R , then the output marker stream is simply formed as the bitwise-or of the input marker stream (zero occurrences of R matched) and the output stream produced by compiling R in the context of the current input marker stream (one occurrence matched).
- If a regular expression is a repetition of a character class of the form C^* , then the compiled form uses the MatchStar operation to produce the output marker stream from the input stream and the compiled stream for character class C .
- If a regular expression is a repetition of a non character class of the form R^* , then a Pablo while loop is created conditioned on a control marker stream still having bits marking match positions. The body of the while consists of the compiled form of the expression R , taking as input the marker at the beginning of the iteration and producing output that becomes the input for the next iteration, if any. The final output is the bitwise-or of matches determined in each loop iteration.
- If a regular expression is a bounded repetition of the form $R\{m,n\}$, then it is compiled according to the equivalent form consisting of m concatenations of R followed by $n - m$ concatenations of $R?$.
- If a regular expression is a bounded repetition of the form $R\{m,\}$, then it is compiled according to the equivalent form consisting of m concatenations of R followed by R^* .

The output of the regular expression compiler is then fed as input to the Pablo compiler of the Parabix tool chain. The result is then compiled with a C++ compiler linked with the Parabix runtime libraries.

4. Block-at-a-Time Processing

The unbounded stream model of the previous section must of course be translated an implementation that proceeds block-at-a-time for realistic application. In this, we primarily rely on the Pablo compiler of the Parabix toolchain [12]. Given input statements expressed as arbitrary-length bitstream equations, Pablo produces block-at-a-time C++ code that initializes and maintains all the necessary carry bits for each of the additions and shifts involved in the bitstream calculations.

In the present work, our principal contribution to the block-at-a-time model is the technique of long-stream addition described below. Otherwise, we were able to use Pablo directly in compiling our SSE2 and AVX2 implementations. Our GPU implementation required some scripting to modify the output of the Pablo compiler for our purpose.

Long-Stream Addition. The maximum word size for addition on commodity processors is typically 64 bits. In order to implement long-stream addition for block sizes of 256 or larger, a method for propagating carries through the individual stages of 64-bit addition is required. However, the normal technique of sequential addition using add-with-carry instructions, for example, is far from ideal.

We have developed a technique using SIMD or SIMT methods for constant-time long-stream addition up to 4096 bits. We assume the availability of the following SIMD/SIMT operations operating on vectors of f 64-bit fields.

- `simd<64>::add(X, Y)`: vertical SIMD addition of corresponding 64-bit fields in two vectors to produce a result vector of f 64-bit fields.
- `simd<64>::eq(X, -1)`: comparison of the 64-bit fields of x each with the constant value -1 (all bits 1), producing an f -bit mask value.
- `hsimd<64>::mask(X)`: gathering the high bit of each 64-bit field into a single compressed f -bit mask value, and
- normal bitwise logic operations on f -bit masks, and
- `simd<64>::spread(x)`: distributing the bits of an f bit mask, one bit each to the f 64-bit fields of a vector.

Given these operations, our method for long stream addition of two $f \times 64$ bit values X and Y is the following.

1. Form the vector of 64-bit sums of x and y .

$$R = \text{simd}\langle 64 \rangle::\text{add}(X, Y)$$

2. Extract the f -bit masks of X , Y and R .

$$x = \text{hsimd}\langle 64 \rangle::\text{mask}(X)$$

$$y = \text{hsimd}\langle 64 \rangle::\text{mask}(Y)$$

$$r = \text{hsimd}\langle 64 \rangle::\text{mask}(R)$$

3. Compute an f -bit mask of carries generated for each of the 64-bit additions of X and Y .

$$c = (x \wedge y) \vee ((x \vee y) \wedge \neg r)$$

4. Compute an f -bit mask of all fields of R that will overflow with an incoming carry bit. This is the *bubble mask*.

$$b = \text{simd}\langle 64 \rangle::\text{eq}(R, -1)$$

5. Determine an f -bit mask identifying the fields of R that need to be incremented to produce the final sum. Here we find a new application of MatchStar!

$$i = \text{MatchStar}(c*2, b)$$

This is the key step. The mask c of outgoing carries must be shifted one position ($c*2$) so that each outgoing carry bit becomes associated with the next digit. At the incoming position, the carry will increment the 64-bit digit. However, if this digit is all ones (as signalled by the corresponding bit of bubble mask b , then the addition will generate another carry. In fact, if there is a sequence of digits that are all ones, then the carry must bubble through each of them. This is just MatchStar!

6. Compute the final result Z .

$$Z = \text{simd}\langle 64 \rangle::\text{add}(R, \text{simd}\langle 64 \rangle::\text{spread}(i))$$

Figure 4 illustrates the process. In the figure, we illustrate the process with 8-bit fields rather than 64-bit fields and show all field values in hexadecimal notation. Note that two of the individual 8-bit additions produce carries, while two others produce FF values

X	19	31	BA	4C	3D	45	21	F1
Y	22	12	45	B3	E2	16	17	36
R	3B	43	FF	FF	1F	5B	38	27
x	0	0	1	0	0	0	0	1
y	0	0	0	1	1	0	0	0
r	0	0	1	1	0	0	0	0
c	0	0	0	0	1	0	0	1
c*2	0	0	0	1	0	0	1	0
b	0	0	1	1	0	0	0	0
i	0	1	1	1	0	0	1	0
Z	3B	44	0	0	1F	5B	39	27

Figure 4. Long Stream Addition

that generate bubble bits. The net result is that four of the original 8-bit sums must be incremented to produce the long stream result.

A slight extension to the process produces a long-stream full adder that can be used in chained addition. In this case, the adder must take an additional carry-in bit p and produce a carry-out bit q . This may be accomplished by incorporating p in calculating the increment mask in the low bit position, and then extracting the carry-out q from the high bit position.

$$i = \text{MatchStar}(c*2+p, b)$$

$$q = i \gg f$$

As described subsequently, we use a two-level long-stream addition technique in both our AVX2 and GPU implementations. In principle, one can extend the technique to additional levels. Using 64-bit adders throughout, $\lceil \log_{64} n \rceil$ steps are needed for n -bit addition. A three-level scheme could coordinate 64 groups each performing 4096-bit long additions in a two-level structure. However, whether there are reasonable architectures that can support fine-grained SIMT style coordinate at this level is an open question.

Using the methods outlined, it is quite conceivable that instruction set extensions to support long-stream addition could be added for future SIMD and GPU processors. Given the fundamental nature of addition as a primitive and its novel application to regular expression matching as shown herein, it seems reasonable to expect such instructions to become available.

5. SSE2 Implementation and Evaluation

5.1 Implementation Notes

5.2 Evaluation Methodology

GREP Implementations: We evaluate our Bitwise Data Parallel implementation against GNU `grep` version 2.10 and `nrgrep` version 1.12. GNU `grep` is a popular open-source `grep` implementation that uses DFAs. NR-`grep` is one of the strongest competitors in regular expression matching performance. It uses an NFA-based approach.

Expressions: Each GREP implementation is tested with the five regular expressions in Table 1. `Xquote` matches any of the representations of a quote in xml. It is run on `roads-2.gml`, a 11,861,751 byte `gml` data file. The other four expressions are taken from Benchmark of Regex Libraries[?] and are all run on a concatenated version of the `linux howto[?]` which is 39,422,105 bytes in length. `@` simply matches the “@” character. It demonstrates the overhead involved in matching the simplest regular expression. `Date`, `Email`, and `URIOrEmail` provide examples of common uses for regular expression matching.

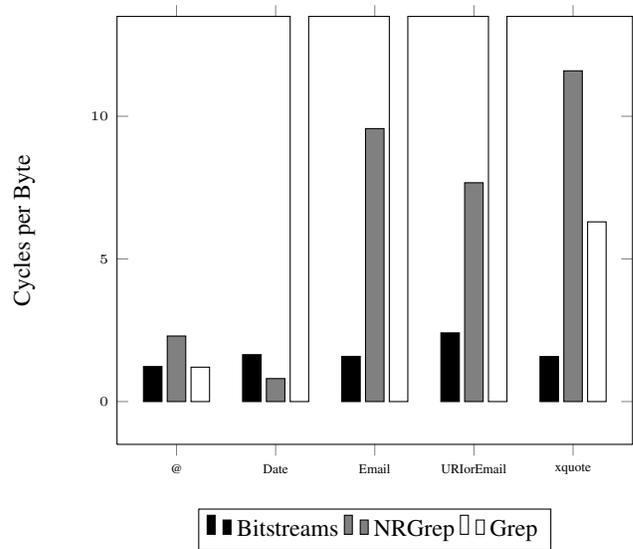


Figure 5. Cycles per Byte

5.3 Comparison

Figure 5 shows cycles per byte. For the three most complicated expressions, the bitstreams implementation had the lowest cycle count.

For the `@` expression, `Grep` had very slightly better performance than the bitstreams implementation. The bitstreams implementation has a fair amount of overhead for transposition that hurts relative performance for very simple expressions.

For the `Date` expression, `NRGrep` is able to skip large portions of the input file since every time it encounters a character that can’t appear in a date, it can skip past six characters. For the more complicated expressions, it loses this advantage.

The bitstreams implementation has fairly consistent performance. As the regular expression complexity increases, the cycle count increases slowly. The largest difference in cycles per byte for the bitstreams implementation is a ratio of 2 to 1. The same cannot be said of `Grep` or `NRGrep`. `NRGrep` uses more than 10 times the cycles per byte for `Xquote` than for `Date`. The number of cycles per byte that `Grep` uses for `URIOrEmail` is almost 900 times as many as it uses for `@`.

Figure 6 shows instructions per byte. The relative values mirror cycles per byte. The bitstreams implementation continues to show consistent performance. This is especially noticeable in Figure 7, which shows instructions per cycle. The bitstreams implementation has almost no variation in the instructions per cycle. Both `Grep` and `NRGrep` have considerably more variation based on the input regular expression.

Figure 8 shows the branch misses per kilobyte. The bitstreams implementation remains consistent here. `NRGrep` and `Grep` have branch miss rates that vary significantly with different regular expressions.

Overall, our performance is considerably better than both `NRGrep` and `Grep` for the more complicated expressions that were tested. Also, our performance scales smoothly with regular expression complexity so it can be expected to remain better for complicated expressions in general.

Name	Expression
@	@
Date	([0-9] [0-9]?)/([0-9] [0-9]?)/([0-9] [0-9] ([0-9] [0-9])?)
Email	([^\s@]+)@([^\s@]+)
URIorEmail	([a-zA-Z][a-zA-Z0-9]*)://([^\s/]+)(/[^\s*]*)? ([^\s@]+)@([^\s@]+)
Xquote	["'] " ' �*3[49]; �*2[27];

Table 1. Regular Expressions

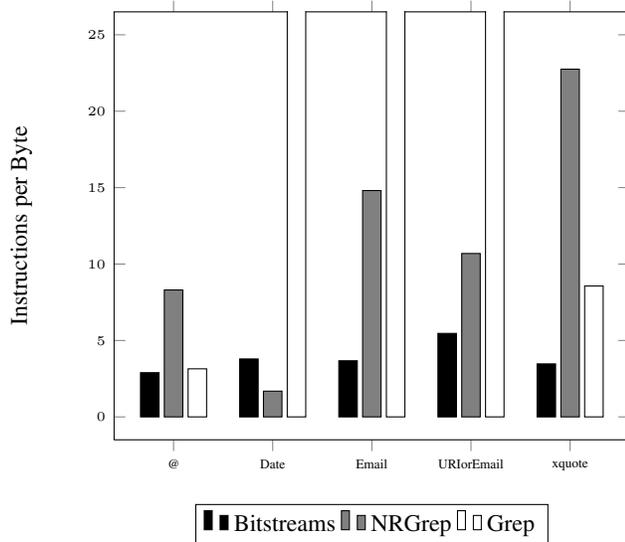


Figure 6. Instructions per Byte

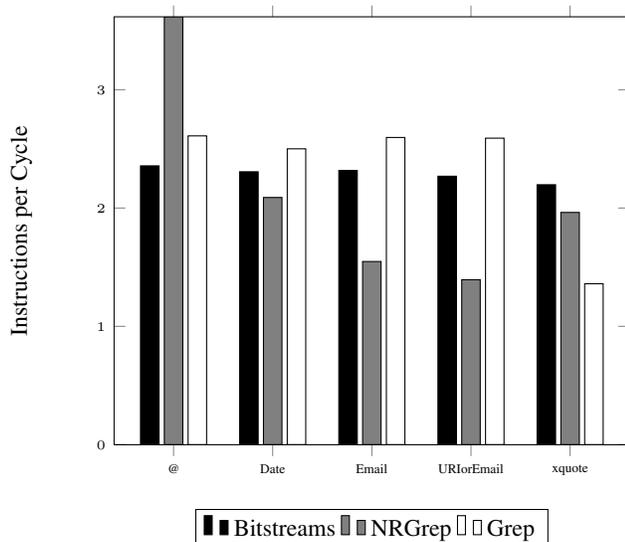


Figure 7. Instructions per Cycle

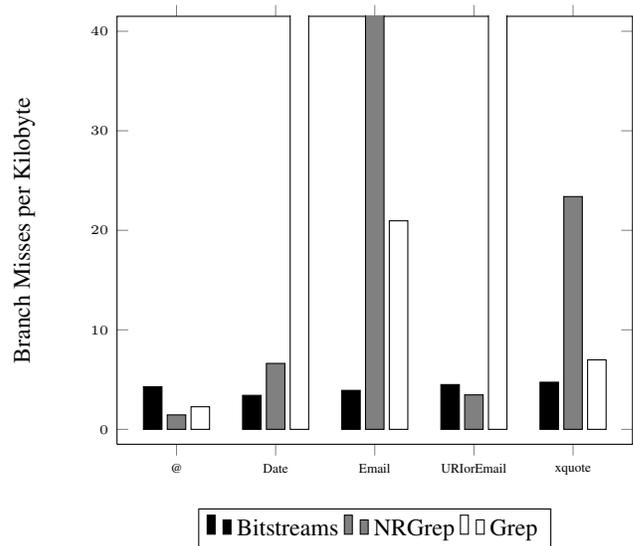


Figure 8. Branch Misses per Kilobyte

6. Running-time Comparison with DFA and NFA Implementations

Our experimental results indicate that regular expression matching using bitstreams can outperform current implementations of NFA- and DFA-based matching. It is worth exploring why this is so, and under what conditions one might expect bitstreams to perform better than NFA- or DFA-based matchers, and vice-versa.

The bitstream method starts with a preprocessing step: the compilation of the regular expression using the `parabix` toolchain. Compilation is an offline process whose time is not counted in our performance measures, as `parabix` is an experimental research compiler that is not optimized. This leads to a bias in our results, as our timings for `nrgrep` and `grep` include the time taken for preprocessing. We have attempted to minimize the bias by performing our tests with large inputs, so that the text-scanning costs dominate the preprocessing costs. We furthermore believe that, if a special-purpose optimized compiler for regular expressions were built, that its inclusion in bitstream `grep` would not substantially increase the running time, particularly for large input texts—the compilation involved is straightforward.

For simplicity, we will first assume that the input regular expressions are restricted to having Kleene closures only of single characters or alternations of single characters. This is a broad class of regular expressions, covering the majority of common uses of `grep`.

The bitstream method compiles a regular expression of size m into bitstream code that is $O(m)$ statements long (with one operation per statement; it is essentially three-address code). This is

translated to machine code and placed inside a loop¹ that executes once per w characters, where w is the width of the processor's word. This gives $O(\frac{nm}{w})$ work. We can further break this down by noting that all of the work in the loop is done by superscalar instructions, with the exception of the additions, which require carry propagation. There will be at most C of these additions in the loop, where C is the number of concatenation and Kleene star operations in the regular expression.

Almost all intermediate bitstreams in the loop body can be kept in registers, requiring no storage in memory. Good register allocation—and limited live ranges for bitstream variables—keeps register spillage to a minimum. For those bitstreams that do require storage in memory, long buffers are allocated, allowing the successive iterations of the loop to access successive memory locations. That is, for the few streams requiring it, memory is accessed in a sequential fashion. As this is the best case for hardware prefetching, we expect few cache misses with bitstream method.

Compare this with NFA methods. In the base NFA method, a state set of approximately m states is kept as a bit set in $\frac{m}{w}$ machine words (or $\frac{m}{8}$ bytes). For each character c of the input, a precomputed transition table, indexed by the c and the current state set, is accessed. Since there are $2^{\Theta(m)}$ state sets, the transition table can have $\sigma 2^{\Theta(m)}$ entries, where σ is the size of the input alphabet. Each entry is a new state set, which requires $\frac{m}{8}$ bytes. Thus, the transition table is of size $\sigma m 2^{\Theta(m)}$, which is quite large: it can become expensive to precompute, and it consumes a lot of memory. For even fairly small m a table of this size will probably not fit in cache memory. Thus, we would expect many cache misses with this base method.

To improve the table size, several authors have separated the transition table into several tables, each indexed by a subset of the bits in the bit set representing the current state. Suppose one uses k bits of the state set to index each table. Ignoring ceilings, this requires $\frac{m}{k}$ tables, each with $\sigma 2^k$ entries of $\frac{m}{8}$ bytes apiece. Each table therefore takes up $2^{k-3} m \sigma$ bytes, and so the collection of them takes up $\frac{m 2^{k-3} \sigma}{k}$ bytes. At each character, the NFA algorithm does one lookup in each table, combining the results with $\frac{m}{k} - 1$ boolean OR operations.

The original NFA method of Thompson uses $k = 1$, which gives a m tables of $\frac{m \sigma}{4}$ bytes each, along with m lookups and $m - 1$ boolean OR operations to combine the lookups, per character.

Navarro and Raffinot use $k = \frac{m}{2}$, giving 2 tables of $2^{\frac{m}{2}-3} m \sigma$ bytes each, two lookups per character, and 1 boolean OR operation per character to combine the lookups.

7. SIMD Scalability

Although commodity processors have provided 128-bit SIMD operations for more than a decade, the extension to 256-bit integer SIMD operations has just recently taken place with the availability of AVX2 instructions in Intel Haswell architecture chips as of mid 2013. This provides an excellent opportunity to assess the scalability of the bitwise data-parallel approach to regular expression matching.

For the most part, adapting the Parabix tool chain to the new AVX2 instructions was straightforward. This mostly involved regenerating library functions using the new AVX2 intrinsics. There were minor issues in the core transposition algorithm because the doublebyte-to-byte pack instructions are confined to independent operations within two 128-bit lanes.

¹Technically, it is inside two loops: an inner one that executes once per w characters in a large buffer, and an outer one that successively fetches buffers until the input is exhausted.

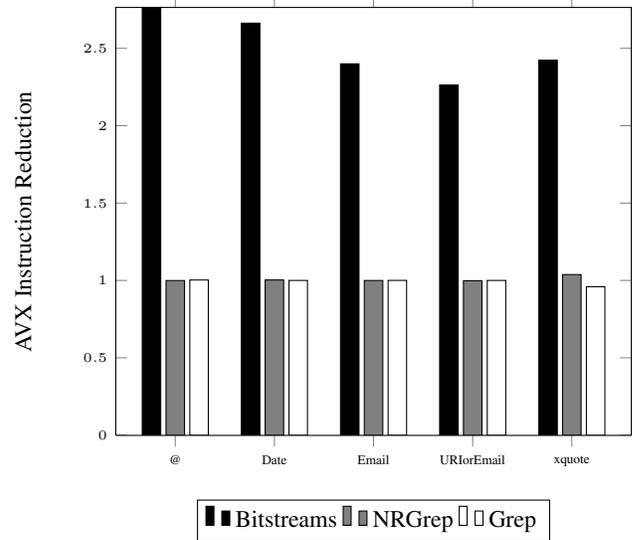


Figure 10. Instruction Reduction

7.1 AVX Stream Addition

Bitstream addition at the 256-bit block size was implemented using the long-stream addition technique. Figure 9 shows our implementation. Spreading bits from the calculated increments mask was achieved somewhat awkwardly with a 64-bit multiply to spread into 16-bit fields followed by SIMD zero extend of the 16-bit fields to 64-bits each.

We also compiled new versions of the `grep` and `ngrep` programs using the `-march=core-avx2` flag in case the compiler is able to vectorize some of the code.

Figure 10 shows the reduction in instruction count achieved for each of the applications. Working at a block size of 256 bytes at a time rather than 128 bytes at a time, the bitstreams implementation scaled dramatically well with reductions in instruction count over a factor of two in each case. Although a factor of two would seem an outside limit, we attribute the change to greater instruction efficiency. AVX2 instructions use a non destructive three-operand form instead of the destructive two-operand form of SSE2. In the two-operand form, binary instructions must always use one of the source registers as a destination register. As a result the SSE2 object code generates many data movement operations that are unnecessary with the AVX2 set.

As expected, there was no observable reduction in instruction count with the recompiled `grep` and `ngrep` applications.

As shown in Figure 11 the reduction in instruction count was reflected in a considerable speed-up in the bitstreams implementation. However, the speed-up was considerably less than expected. As shown in the AVX2 version has lost some of the superscalar efficiency of the SSE2 code. This is a performance debugging issue that we have yet to resolve.

Overall, the results on our AVX2 machine were quite good, demonstrating very good scalability of the bitwise data-parallel approach.

8. GPU Implementation

To further assess the scalability of our regular expression matching using bit-parallel data streams, we implemented a GPGPU version in OpenCL. We arranged for 64 work groups each having 64 threads. Input files are divided in data parallel fashion among the

```

void add_ci_co(bitblock_t x, bitblock_t y, carry_t carry_in, carry_t & carry_out, bitblock_t & sum) {
    bitblock_t all_ones = simd256<1>::constant<1>();
    bitblock_t gen = simd_and(x, y);
    bitblock_t prop = simd_xor(x, y);
    bitblock_t partial_sum = simd256<64>::add(x, y);
    bitblock_t carry = simd_or(gen, simd_andc(prop, partial_sum));
    bitblock_t bubble = simd256<64>::eq(partial_sum, all_ones);
    uint64_t carry_mask = hsimd256<64>::signmask(carry) * 2 + convert(carry_in);
    uint64_t bubble_mask = hsimd256<64>::signmask(bubble);
    uint64_t carry_scan_thru_bubbles = (carry_mask + bubble_mask) &~ bubble_mask;
    uint64_t increments = carry_scan_thru_bubbles | (carry_scan_thru_bubbles - carry_mask);
    carry_out = convert(increments >> 4);
    uint64_t spread = 0x0000200040008001 * increments & 0x0001000100010001;
    sum = simd256<64>::add(partial_sum, _mm256_cvtepu16_epi64(avx_select_lo128(convert(spread))));
}

```

Figure 9. AVX2 256-bit Addition

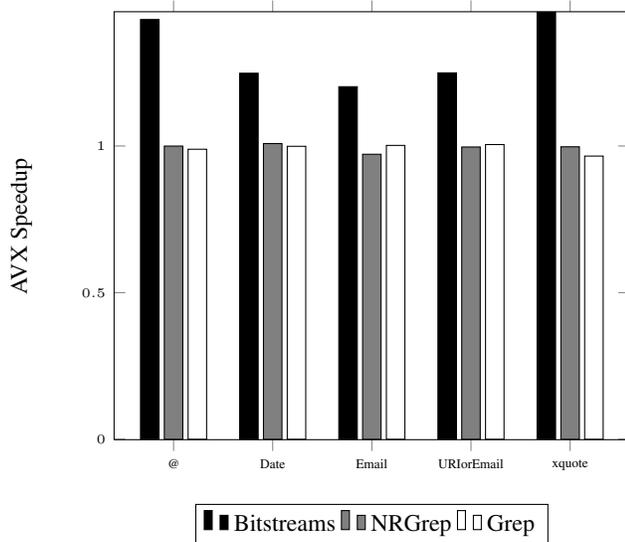


Figure 11. AVX Speedup

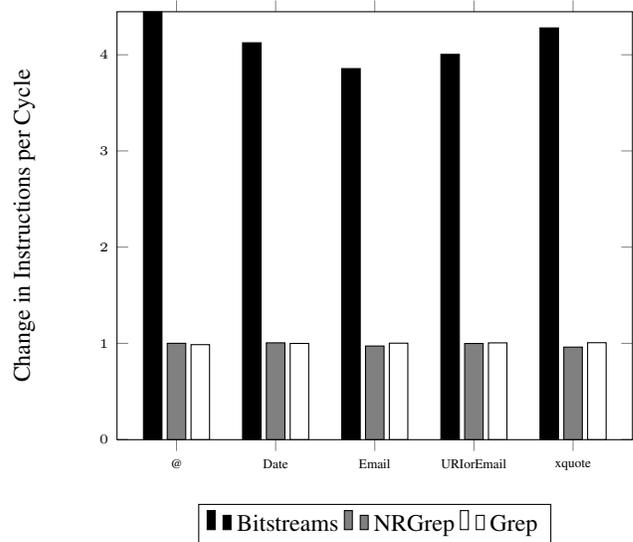


Figure 12. Change in Instructions Per Cycle With AVX

64 work groups. Each work group carries out the regular expression matching operations 4096 bytes at a time using SIMT processing. Figure 13 shows our implementation of long-stream addition on the GPU. Each thread maintains its own carry and bubble values and performs synchronized updates with the other threads using a six-step parallel-prefix style process.

Our GPU test machine was an AMD A10-5800K APU with Radeon(tm) HD Graphics having a processor speed of 4.30 GHz and 32.0GB of memory.

9. Miscellaneous

9.1 Skipping

9.2 Unicode

The introduction of Unicode as a common encoding system including the characters of all the world's written languages and notation systems has introduced some complexity for regular expression matching engines. Nevertheless, most modern tools and libraries do include some form of Unicode support, with varying degrees of performance loss.

UTF-8, UTF-16 and UTF-32 are the three common transformation formats of Unicode depending on whether character code points are encoded using 8-bit, 16-bit or 32-bit code units. In the case of the 32-bit code units of UTF-32, each Unicode character is encoded as a single 32-bit unit, with the high 11 bit positions all zero. A stream of UTF-32 encoded text can then be processed using a straightforward application of the techniques described above by first transforming to a set of 21 parallel bit streams for each of the significant bit positions within UTF-32.

For most practical purposes, UTF-16 can also be processed similarly, considering that each 16-bit code unit represents a single character. For the rarely used characters of the Unicode supplementary plane, two 16-bit code units are required. Such a two code unit sequence is known as a surrogate pair. Following the common practice of treating each member of a surrogate pair as pseudo-character, UTF-16 can also be processed by the straightforward transposition to 16 parallel bit streams and application of the techniques above.

UTF-8 is probably the most widely used of the Unicode formats, primarily because of its compatibility with widely deployed

```

inline BitBlock adc(int idx, BitBlock a, BitBlock b, __local BitBlock *carry, _
    __local BitBlock *bubble, BitBlock *group_carry, const int carryno){
BitBlock carry_mask;
BitBlock bubble_mask;

BitBlock partial_sum = a+b;
BitBlock gen = a&b;
BitBlock prop = a^b;
carry[idx] = ((gen | (prop & ~partial_sum))&CARRY_BIT_MASK)>>(WORK_GROUP_SIZE-1-idx);
bubble[idx] = (partial_sum + 1)? 0:(((BitBlock)1)<<idx);

barrier(CLK_LOCAL_MEM_FENCE);
for(int offset=WORK_GROUP_SIZE/2; offset>0; offset=offset>>1){
carry[idx] = carry[idx]|carry[idx^offset];
bubble[idx] = bubble[idx]|bubble[idx^offset];
barrier(CLK_LOCAL_MEM_FENCE);
}

carry_mask = (carry[0]<<1)|group_carry[carryno];
bubble_mask = bubble[0];

BitBlock s = (carry_mask + bubble_mask) & ~bubble_mask;
BitBlock inc = s | (s-carry_mask);
BitBlock rslt = partial_sum + ((inc>>idx)&0x1);
group_carry[carryno] = (carry[0]|(bubble_mask & inc))>>63;
return rslt;
}

```

Figure 14 compares the performance of our SSE2, AVX and GPU implementations.

Figure 13. OpenCL 4096-bit Addition

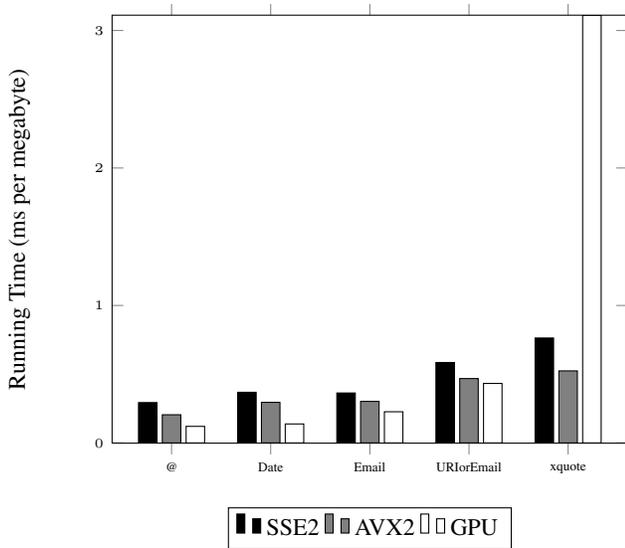


Figure 14. Running Time

networking software based on 8-bit extended-ASCII character representations. UTF-8 is a variable length coding system in which each Unicode character is represented using one to four 8-bit code units.

In order to safely process UTF-8, it is necessary to validate that it is well-formed. Each byte is classified as either an ordinary ASCII byte (high bit clear: range 0x00-0x7F), a UTF-8 prefix byte of a 2-, 3- or 4- byte sequence (in ranges 0xC2-0xDF, 0xE0-0xEF,

and 0xF0-F4, respectively) or as a UTF-8 suffix byte in the range 0x80-0xBF. Parallel bit stream technology achieves this validation easily and efficiently [6].

The UTF-8 byte classification streams produced as a byproduct of validation then enable regular expression on the UTF-8 input streams as follows. For each multibyte character used in the pattern, a character classification bitstream may be formed to identify the occurrence of such characters at each position in the source stream at which the final byte of the character is found. Figure ?? illustrates. (Describe).

Using this approach to multibyte classification, matching of a single k -byte character is straightforward. All current match positions are shifted forward $k - 1$ positions and then combined with the computed character class using bitwise-and. The result is then shifted forward one position to produce the result of the multibyte character match. This method easily extends to character classes comprising multibyte characters all of the same length.

Matching a character class comprising characters of different lengths requires a slightly different strategy. In this case, we take advantage of a bitstream *nonfinal* formed from the UTF-8 byte classification streams to consist of all those positions at which a UTF-8 prefix byte is found, as well as those positions at which the second byte of a 3-byte sequence or the second or third byte of a 4-byte sequence are found. We then apply `ScanThru(current,nonfinal)` to advance all current matches to the final position of the next character in the input. Bitwise-and combination with the character class bitstream produces the result identifying matches with this class, the result is then shifted forward 1 position.

The MatchStar operation for matching arbitrary sequences of a character class can similarly take advantage of the *nonfinal* stream. For this case, we form the bitwise-or of the *nonfinal* stream with the character class stream prior to applying MatchStar.

The result will propagate bits to the first character position of matches and to final positions of nonmatches. We then clear the nonmatches by combining the result stream with the `suffix` byte stream using bitwise-and.

10. Conclusion

Contributions A new class of regular expression matching algorithm has been introduced based on the concept of bit-parallel data streams together with the MatchStar operation. The algorithm is fully general for nondeterministic regular expression matching; however it does not address the nonregular extensions found in Perl-compatible backtracking implementations. Taking advantage of the SIMD features available on commodity processors, its implementation in a `grep` too offers consistently good performance in contrast to available alternatives. While lacking some special optimizations found in other engines to deal with repeated substrings or to perform skipping actions based on fixed substrings, it nevertheless performs competitively in all cases. The algorithm tends to scale very well with regular expression complexity, often with order-of-magnitude performance advantage over even the best of its competitors.

A parallelized algorithm for long-stream addition has also been introduced in the paper making a key contribution to the scalability of the bit-parallel matching technique overall and that of MatchStar in particular. This algorithm has enabled straightforward extension of the matching algorithm to the implementation using 256-bit AVX2 technology as well as 4096-bit SIMT implementation on an AMD GPU.

Future Work An important area of future work is to develop and assess multicore versions of the algorithm to handle regular expression matching problems involving larger rulesets than are typically encountered in the `grep` problem. Such implementations could have useful application in tokenization and network intrusion detection for example.

Another area of interest is to extend the capabilities of the underlying method with addition features for substring capture, zero-width assertions and possibly backreference matching. Extending Unicode support beyond basic Unicode character handling to include full Unicode character class support and normalization forms is also worth investigating.

Acknowledgments

This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada and MITACS, Inc.

References

- [1] A. V. Aho and M. J. Corasick. Efficient string matching: an aid to bibliographic search. *Commun.ACM*, 18(6):333–340, June 1975. URL <http://doi.acm.org/10.1145/360825.360855>.
- [2] R. Baeza-Yates and G. H. Gonnet. A new approach to text searching. *Communications of the ACM*, 35(10):74–82, 1992.
- [3] R. A. Baeza-yates, B. Encalada, and G. H. Gonnet. A new approach to text searching.
- [4] P. Bille and M. Thorup. Faster regular expression matching. In *Automata, Languages and Programming*, pages 171–182. Springer, 2009.
- [5] R. S. Boyer and J. S. Moore. A fast string searching algorithm. *Communications of the ACM*, 20(10):762–772, 1977.
- [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*, pages 91–98. ACM, 2008.
- [7] R. D. Cameron, E. Amiri, K. S. Herdy, D. Lin, T. C. Shermer, and F. P. Popowich. Parallel scanning with bitstream addition: An XML case study. In *Euro-Par 2011 Parallel Processing*, pages 2–13. Springer, 2011.
- [8] M. Crochemore, W. Rytter, and M. Crochemore. *Text algorithms*, volume 698. World Scientific, 1994.
- [9] F. Iorio and J. V. Lunteren. Fast pattern matching on the cell broadband engine. In *2008 Workshop on Cell Systems and Applications (WCSA), affiliated with the*, 2008.
- [10] S. C. Kleene. Representation of events in nerve nets and finite automata. 1951.
- [11] C. Lin, C. Liu, L. Chien, and S. Chang. Accelerating pattern matching using a novel parallel algorithm on gpus. *IEEE Transactions on Computers*, 62(10), 2013.
- [12] D. Lin, N. Medforth, K. S. Herdy, A. Shiraman, and R. Cameron. Parabix: Boosting the efficiency of text processing on commodity processors. In *High Performance Computer Architecture (HPCA), 2012 IEEE 18th International Symposium on*, pages 1–12. IEEE, 2012.
- [13] J. Naghmouchi, D. P. Scarpazza, and M. Berekovic. Small-ruleset regular expression matching on gpgpus: quantitative performance analysis and optimization. In *Proceedings of the 24th ACM International Conference on Supercomputing, ICS '10*, pages 337–348, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0018-6. URL <http://doi.acm.org/10.1145/1810085.1810130>.
- [14] G. Navarro. Nr-grep: A fast and flexible pattern matching tool. *Software Practice and Experience (SPE)*, 31:2001, 2000.
- [15] G. Navarro and M. Raffinot. A bit-parallel approach to suffix automata: Fast extended string matching. In *Combinatorial Pattern Matching*, pages 14–33. Springer, 1998.
- [16] D. Pasetto, F. Petrini, and V. Agarwal. Tools for very fast regular expression matching. *Computer*, 43(3):50–58, 2010.
- [17] V. Salapura, T. Karkhanis, P. Nagpurkar, and J. Moreira. Accelerating business analytics applications. In *High Performance Computer Architecture (HPCA), 2012 IEEE 18th International Symposium on*, pages 1–10. IEEE, 2012.
- [18] D. P. Scarpazza. Is larrabee for the rest of us? *Dr.Dobbs J*, 2009.
- [19] D. P. Scarpazza. Top-performance tokenization and small-ruleset regular expression matching. *International Journal of Parallel Programming*, 39(1):3–32, 2011.
- [20] D. P. Scarpazza and G. F. Russell. High-performance regular expression scanning on the cell/be processor. In *Proceedings of the 23rd international conference on Supercomputing*, pages 14–25. ACM, 2009.
- [21] D. P. Scarpazza and G. F. Russell. High-performance regular expression scanning on the cell/be processor. In *Proceedings of the 23rd international conference on Supercomputing*, pages 14–25. ACM, 2009.
- [22] D. P. Scarpazza, O. Villa, and F. Petrinni. Fast string searches & multicore processors mapping fundamental algorithms on parallel hardware, 2008.
- [23] K. Thompson. Programming techniques: Regular expression search algorithm. *Communications of the ACM*, 11(6):419–422, 1968.
- [24] A. Tumeo, O. Villa, and D. Sciuto. Efficient pattern matching on gpus for intrusion detection systems. In *Proceedings of the 7th ACM international conference on Computing frontiers*, pages 87–88. ACM, 2010.
- [25] S. Wu and U. Manber. agrepa fast approximate pattern-matching tool. *Usenix Winter 1992*, pages 153–162, 1992.
- [26] S. Wu and U. Manber. Fast text searching: allowing errors. *Communications of the ACM*, 35(10):83–91, 1992.