

ParPaRaw: Massively Parallel Parsing of Delimiter-Separated Raw Data

Elias Stehle
Technical University of Munich (TUM)
stehle@in.tum.de

Hans-Arno Jacobsen
Technical University of Munich (TUM)
jacobsen@in.tum.de

ABSTRACT

Parsing is essential for a wide range of use cases, such as stream processing, bulk loading, and in-situ querying of raw data. Yet, the compute-intensive step often constitutes a major bottleneck in the data ingestion pipeline, since parsing of inputs that require more involved parsing rules is challenging to parallelise. This work proposes a massively parallel algorithm for parsing delimiter-separated data formats on GPUs. Other than the state-of-the-art, the proposed approach does not require an initial sequential pass over the input to determine a thread's parsing context. That is, how a thread, beginning somewhere in the middle of the input, should interpret a certain symbol (e.g., whether to interpret a comma as a delimiter or as part of a larger string enclosed in double-quotes). Instead of tailoring the approach to a single format, we are able to perform a massively parallel finite state machine (FSM) simulation, which is more flexible and powerful, supporting more expressive parsing rules with general applicability. Achieving a parsing rate of as much as 14.2 GB/s, our experimental evaluation on a GPU with 3584 cores shows that the presented approach is able to scale to thousands of cores and beyond. With an end-to-end streaming approach, we are able to exploit the full-duplex capabilities of the PCIe bus and hide latency from data transfers. Considering the end-to-end performance, the algorithm parses 4.8 GB in as little as 0.44 seconds, including data transfers.

PVLDB Reference Format:

Elias Stehle, and Hans-Arno Jacobsen. ParPaRaw: Massively Parallel Parsing of Delimiter-Separated Raw Data. *PVLDB*, 13(5): 616-628, 2020.
DOI: <https://doi.org/10.14778/3377369.3377372>

1. INTRODUCTION

Massive amounts of data from a wide range of data sources are made available using delimiter-separated formats, such as comma-separated values (CSV), and various log file formats like the Common Log Format and the Extended Log

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 13, No. 5

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3377369.3377372>

Format [37, 30, 19]. The relevancy of the CSV format, for instance, is highlighted by the plethora of public datasets, some in excess of hundreds of gigabytes in size, that are provided using the CSV format [43, 24]. Log files are another origin of data in a delimiter-separated format that constitute an important source for many analytical workloads. For instance, *Sumo Logic*, a cloud-based log management and analytics service, recently announced that it analyses more than 100 petabytes of data and 500 trillion records daily [42]. With an ever increasing amount of data, there is also a growing need to provide and maintain rapid access to data in delimiter-separated formats. This is also emphasised by ongoing research on in-situ processing of raw data and similar efforts that aim to lower the time to insight [11, 35, 6, 18, 25, 46, 7, 12, 26, 22, 2, 21].

While systems face an ever increasing amount of data that needs to be ingested and analysed, processors are seeing only moderate improvements in sequential processing performance. In order to continue the trend of providing exponentially growing computational throughput, manufacturers have therefore progressively turned towards scaling the number of cores as well as extending single instruction, multiple data (SIMD) capabilities. Graphics Processing Units (GPUs), which have focused on parallelism ever since, now integrate as much as 5120 cores on a single chip [1]. Further, CPUs comprising multiple *chiplets*, as well as research focusing on package-level integration of multiple GPU modules, give an indication that hardware parallelism moves even beyond a single chip, scaling to multiple inherently parallel *chiplets* and GPU modules, respectively, on a package [5].

In order to leverage the current degree of hardware parallelism and benefit from the ongoing trend of an ever growing number of cores, algorithms have to be designed for massive scalability from the ground up [23]. Parsing, as a fundamental and compute-intensive step in the data ingestion pipeline is no exception to this.

Parallel parsing of non-trivial delimiter-separated data formats, however, poses a great challenge, as symbols have to be interpreted differently, depending on the context they appear in. For the CSV format, for instance, RFC 4180 specifies that delimiters (i.e., commas and line breaks), which appear within a field that is enclosed in double-quotes, have to be interpreted as part of the field, instead of being interpreted as actual field or record delimiters [37]. In addition, many formats use a symbol to indicate comments or directives (e.g., '#'), following which, all symbols until the end of line have to be interpreted differently, yet again. Since the context depends on all symbols preceding the symbol

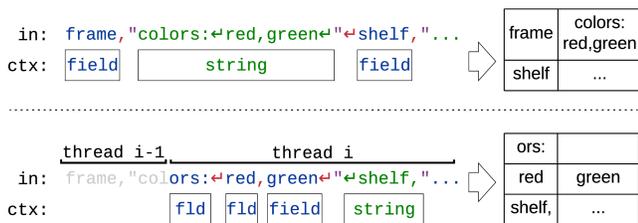


Figure 1: Challenges for parallel parsing: lacking context leads to misinterpretation

currently being interpreted, it is impossible for a thread to simply begin parsing somewhere in the middle of the input. Hence, the input cannot simply be split into multiple chunks that are processed independently.

This is exemplified in Figure 1, where *thread i* begins parsing in the middle of the input. The thread is not aware of the double-quote preceding its chunk that indicates the beginning of a larger string, changing the parsing context. As a result, the thread misinterprets subsequent commas and line breaks as delimiters, while they were actually supposed to be considered as part of the field’s string. A similar challenge arises for determining the records and columns that a chunk of the input belongs to, which, again, depends on all the input preceding the current chunk being interpreted. Finally, threads have to coordinate and possibly collaborate in order to assemble field values that span multiple threads. This may also involve converting symbols to a binary type (e.g., `int`, `float`).

Previous work on parallel loading of delimiter-separated data formats has addressed the challenge of determining a thread’s parsing context by either performing an initial sequential pass over the input or by completely dropping support for inputs with different parsing contexts, such as inputs containing enclosing symbols (e.g., double-quotes), comments, or directives [33, 3]. Another alternative is to tailor the approach to one specific format and exploit the format-specific characteristics, which, however, limits the approach’s flexibility and applicability [29, 28, 39, 36, 16]. One such exploit for a simple CSV format, for instance, is to count the number of double-quotes, inferring the beginning and end of enclosed strings depending on whether the count is odd or even, respectively. More recently, Ge et al. presented an approach for distributed CSV parsing [16]. They aim to circumvent an initial sequential pass by exploiting CSV-specific characteristics to speculate about the parsing context. While such a tailored approach works well with CSV as long as it strictly complies with the format expected by the algorithm, it requires designing a completely new approach from the ground up once the input format deviates. Parsing other delimiter-separated formats, such as log files and their multifaceted formats, poses a challenge for an approach that is tailored to CSVs. Another important characteristic of state-of-the-art approaches is that they are designed for coarse-grained parallelism of distributed and multi-core systems, which renders them infeasible for the fine-grained parallelism required by GPUs [33, 16].

While constraining the input limits the applicability and flexibility, performing a sequential pass over the input contributes a substantial portion of sequential work that limits scalability and, following Amdahl’s law, precludes any speed-ups beyond a certain point. Given the ongoing trend

of increasing hardware parallelism on the one hand and the diversity of data sources that today’s OLAP systems are confronted with on the other hand, addressing these shortcomings is a viable endeavour.

We present **ParPaRaw**, an algorithm for massively parallel parsing of delimiter-separated **raw** data on GPUs that overcomes these scalability issues without compromising applicability or constraining supported input formats. **ParPaRaw** is designed from the ground up to scale linearly with the number of cores, providing robust performance despite the huge diversity of inputs it is confronted with, by employing a data parallel approach with fine-grained parallelism. **ParPaRaw** is designed to leverage the specifics of GPUs. It enables parallelism even beyond the granularity of a single record and ensures load balancing by splitting the input into small chunks of equal size that threads can process independently. Since using a data parallel approach raises the aforementioned challenges, we present an efficient solution for correctly identifying the parsing context of a thread’s chunk, as well as its records and columns. In order to provide a flexible approach that is applicable to a wide range of inputs, we allow specifying the parsing rules in the form of a deterministic finite automaton (DFA). In order to exploit the full-duplex capabilities of the Peripheral Component Interconnect Express (PCIe) bus and lower the end-to-end latency, we present a streaming approach, which parses data on the GPU, while simultaneously transferring raw data to, and parsed data from the GPU.

With a generally applicable approach that does not impose constraints on the input, we are able to parse as much as 14.2 GB/s on the GPU. For end-to-end workloads, including data transfers via the PCIe bus, **ParPaRaw** parses 4.8 GB from the yelp reviews dataset in as little as 0.44 seconds. In summary, the contributions of this paper are four-fold.

1. We present an approach to massively parallel parsing of delimiter-separated data formats that is designed for scalability without sacrificing applicability and flexibility. The approach develops a scalable, data parallel algorithm that addresses three challenges: a) determining a thread’s parsing context without requiring a prior sequential pass, b) determining the records and columns that a thread’s symbols belong to, and c) efficiently coordinating threads to collaboratively generate field values.
2. We address the major challenges that arise when mapping our algorithm to GPUs, which provide only very limited addressable on-chip memory (tens of KB) and, due to their limited register file size, require lightweight threads with only very limited context.
3. We show how to exploit the full-duplex capabilities of the PCIe bus with a streaming extension. This lowers the end-to-end latency and allows parsing data on the GPU, while simultaneously transferring raw data to, and returning parsed data from the GPU.
4. Our experimental evaluation highlights that, given today’s level of hardware parallelism, it is worth to design algorithms for scalability from the ground up, even if it implies a significant increase in the overall work being performed.

This paper is organised as follows. Section 2 gives an overview of related approaches. Section 3 presents the algorithm, its building blocks, and the processing steps. Section 4 introduces optimisations, extensions, and implementation details. Section 5 evaluates the presented approach.

2. RELATED WORK

Even though parsing is fundamental for in-situ processing of raw files and constitutes a major bottleneck in the data ingestion pipeline, there is only limited work on accelerating the process. This is also highlighted by Dziejczak et al., who show that modern Database Management Systems (DBMSs) are unable to saturate available I/O bandwidth [14]. Using a variety of hardware configurations and datasets, Dziejczak et al. provide an extensive analysis of the data loading process for multiple state-of-the-art DBMSs [14]. Their evaluation reveals that data loading is CPU-bound [14].

A notable advancement for parsing delimiter-separated formats is made by Mühlbauer et al. who present improvements along two lines [33]. On the one hand, they introduce optimisations to reduce the number of control flow branches by utilising SIMD instructions for the identification of delimiters. On the other hand, they present an approach for parallel parsing. Their approach splits the input into multiple chunks of equal size that are processed in parallel. Threads start parsing their chunk only from an actual record boundary onward, i.e., the first record delimiter in their chunk. Threads continue parsing beyond their chunk until encountering the end of their last record. This ensures that threads always process complete records, yet makes the approach sensible to the chosen chunk size and the input’s record sizes. For instance, the majority of threads, which work on a record that spans multiple chunks, unsuccessfully search for the beginning of their first record, without performing actual parsing work. Another shortcoming is that threads are not aware of the actual parsing context of their chunk. That is, whether to interpret a field or record delimiter as an actual delimiter or as part of a field’s value. To address this, they introduce a *safe mode* for formats that may contain more involved parsing rules. In *safe mode* a sequential pass over the input is performed, which keeps track of the parsing context, such as quotation scopes, splitting chunks only at actual record delimiters. *Safe mode*, however, introduces a considerable portion of serial work, which, according to Amdahl’s law, precludes any speedup beyond a certain point. By exploiting CSV-specific characteristics, Ge et al., who look at distributed CSV parsing, are able to bypass that initial pass by speculating about the parsing context [16]. Similar to the approach of Mühlbauer et al., Ge et al. require coarse-grained parallelism (distributed parsing), as threads begin parsing only from a chunk’s first record boundary onward. As a result their approach is also sensible to the chosen chunk size and the input’s record sizes. The authors highlight that the performance of their approach degrades with decreasing chunk sizes, once a chunk approaches the size of a record [16]. Moreover, both solutions do not provide parallelism beyond the granularity of an individual record, which makes them susceptible to load-balancing issues, particularly for small chunks and large, varying record sizes. These circumstances render the two approaches infeasible for the fine-grained parallelism required by GPUs [33, 16].

Apart from work addressing delimiter-separated formats, multiple approaches tailored to processing JSON have been proposed. Li et al. present *Mison*, a JSON parser that supports projection and filter pushdown by speculatively predicting logical locations of queried fields based on previously seen patterns [29]. *Mison* deviates from the classic approach of using an FSM while parsing, which allows it to use

SIMD vectorisation to identify structural characters, such as double-quotes, braces, and colons. Whenever a structural character is encountered, its occurrence is recorded in the respective bitmap index (e.g., the double-quotes bitmap index). The beginning and end of a string enclosed in double-quotes can be inferred from looking at the odd and even number of set bits, respectively. While this enables SIMD vectorisation and avoids branch divergence, circumventing the use of an FSM and, hence, tailoring the approach specifically to the JSON format, limits the approach’s applicability to formats with more involved parsing rules. Bonetta et al. introduce *FAD.js*, a runtime system for processing JSON objects that is based on speculative JIT compilation and selective access to data [9]. Palkar et al. propose a technique referred to as *raw filtering*, which is applied on the data’s raw bytestream before parsing [35]. Langdale et al. recently introduced *simdjson*, a standard-compliant, highly-efficient JSON parser that makes use of SIMD instructions. Similar to *Mison*, they focus on the JSON format, which allows them to avoid the use of an FSM while parsing [28].

The parallel prefix scan is a fundamental algorithm of **ParPaRaw** and a frequently recurring building block for data parallel algorithms. Over the years many approaches for a parallel prefix scan have been proposed [31, 44, 17, 32, 13, 8, 20, 10, 27, 40]. For a given binary reduction operator (e.g., addition), it takes an array of input elements and returns an array, where the i -th output element is computed by applying the reduction operator to all input elements up to and including the i -th element [31]: $y_i = \bigoplus_{k=0}^i x_k$

A prefix scan that excludes the i -th input element is called exclusive prefix scan. The prefix scan using addition is called prefix sum. The following table shows an example of the inclusive and exclusive prefix sum, respectively:

x_i	3	5	1	2	9	7	4	2
y_i (incl.)	3	8	9	11	20	27	31	33
y_i (excl.)	0	3	8	9	11	20	27	31

It is worth noting that all efficient parallel approaches require the binary operator to be associative. The prefix scan used in **ParPaRaw** builds on the more recent work from Merrill et al., who propose a single-pass prefix scan [31]. Using the parallel prefix scan has also been considered for parallel parsing of regular languages. In fact, the theory for parallel parsing dates back as long as four decades. In particular, Fischer presents an algorithm that instantiates one FSM for each state defined by the FSM [15]. Hillis et al. illustrate the use of the parallel prefix scan computation by presenting an algorithm that is similar to earlier work from Fischer [20]. Even though the theory dates back several decades, it became feasible only more recently with modern hardware, i.e., GPUs with thousands of cores, that would set off the cost of running multiple FSM instances. We believe that this is the main reason why, to the best of our knowledge, this idea has not been considered for identifying the parsing context when parsing delimiter-separated formats on GPUs. Our approach reconsiders the idea behind the work from Fischer and Hillis et al. to address the sub-problem of identifying the parsing context. For identifying the parsing context, we devise a solution that respects the characteristics of GPUs. We show that, given today’s degree of hardware parallelism, the cost of running multiple FSM instances can be set off with our efficient and scalable solution.

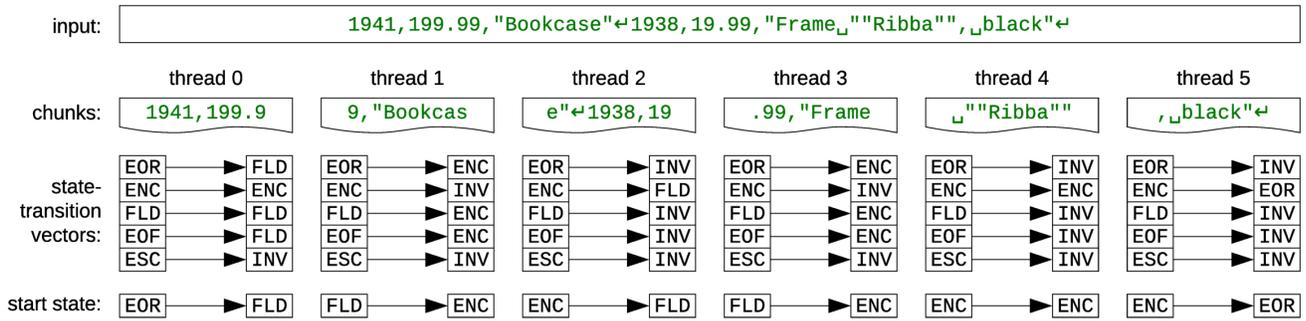


Figure 2: Determining the parsing context

3. MASSIVELY PARALLEL PARSING

In order to achieve scalability, even beyond thousands of cores, we pursue a data parallel approach, which splits the input into multiple chunks (e.g., 32 bytes per chunk) that can be processed independently by the threads. While a data parallel approach allows for massive scalability, there are three key challenges to overcome:

1. Determining the parsing context of a thread's chunk. That is, how a thread is supposed to interpret the symbols within its chunk (Section 3.1).
2. Determining the records and columns that the symbols of a thread's chunk belong to (Section 3.2).
3. Efficient coordination and collaboration between threads to transform a sequence of symbols to the data type of the respective column, e.g., `float` (Section 3.3).

ParPaRaw addresses these challenges in multiple steps. With each step, ParPaRaw gains additional information about each thread's chunk. This information is captured in meta data that subsequent steps can build on.

3.1 Parsing

The first step addresses the challenge of identifying the parsing context of a thread's chunk, allowing a thread to meaningfully interpret its symbols. That is, distinguishing whether a symbol is a control symbol (e.g., delimiting a field or a record) or whether it is part of a field's value.

It is important to note that without constraining the supported input formats and therefore sacrificing the approach's applicability, it is impossible to determine a thread's parsing context without considering all symbols preceding its chunk. However, if a thread is supposed to consider all symbols preceding its chunk, the approach has to either perform an initial sequential pass over the input or wait for all threads working on preceding chunks to finish. Considering all symbols preceding a thread's chunk introduces severe implications on the approach's scalability.

ParPaRaw, however, aims to neither constrain the input nor to introduce sequential work. In order to achieve this, we exploit the fact that there are only few different contexts to consider while parsing. While this increases the overall effort by a constant factor, it enables a fully concurrent approach and allows to scale linearly with the number of cores.

In pursuit of a flexible approach that is generally applicable, ParPaRaw uses a DFA while parsing. The current parsing context is represented by the DFA's state. While a thread iterates over its symbols, it transitions the states of its DFA according to its transition table. One example of a

DFA for parsing a simple CSV format is shown in Figure 3 (for simplicity it omits the invalid state (INV) used to track invalid formats, e.g., reading quotes in FLD state). A sequential approach would simply set the starting state of its DFA and read the symbols of the input beginning to end, always being aware of the current state when reading a symbol. For a data parallel approach, however, a thread, starting to parse somewhere in the middle of the input, cannot simply infer the state it is supposed to start in.

In order to perform meaningful work despite lacking the correct starting state, each thread instantiates one DFA for every state, $s_i \in S$, defined by the DFA, setting the starting state of the i -th DFA-instance to state s_i (see *state-transition vectors* in Figure 2). For reasons of clarity, in the following we assume $s_i = i$, i.e., representing a state by its index, to avoid the intricate differentiation between a state and a state index. An efficient implementation uses the same mechanism during preprocessing to ensure efficient lookups into data structures like the transition table. While the thread is reading the symbols of its chunk, it transitions the states of all its DFA-instances accordingly. Once all the symbols of a chunk have been read, the final state of each DFA-instance is noted in a state-transition vector. We maintain one state-transition vector per thread, with each state-transition vector holding $|S|$ elements. The i -th entry of the state-transition vector represents the final state of the i -th DFA-instance (i.e., the DFA-instance that has originally started in state s_i). Hence, the algorithm can infer that if a thread had started parsing in state s_i , it would end up in the state given by the i -th entry of that thread's state-transition vector after the thread has read all its symbols (see Figure 2).

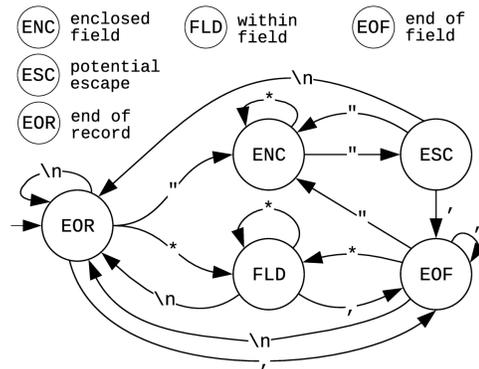


Figure 3: Example for a simple DFA parsing CSVs

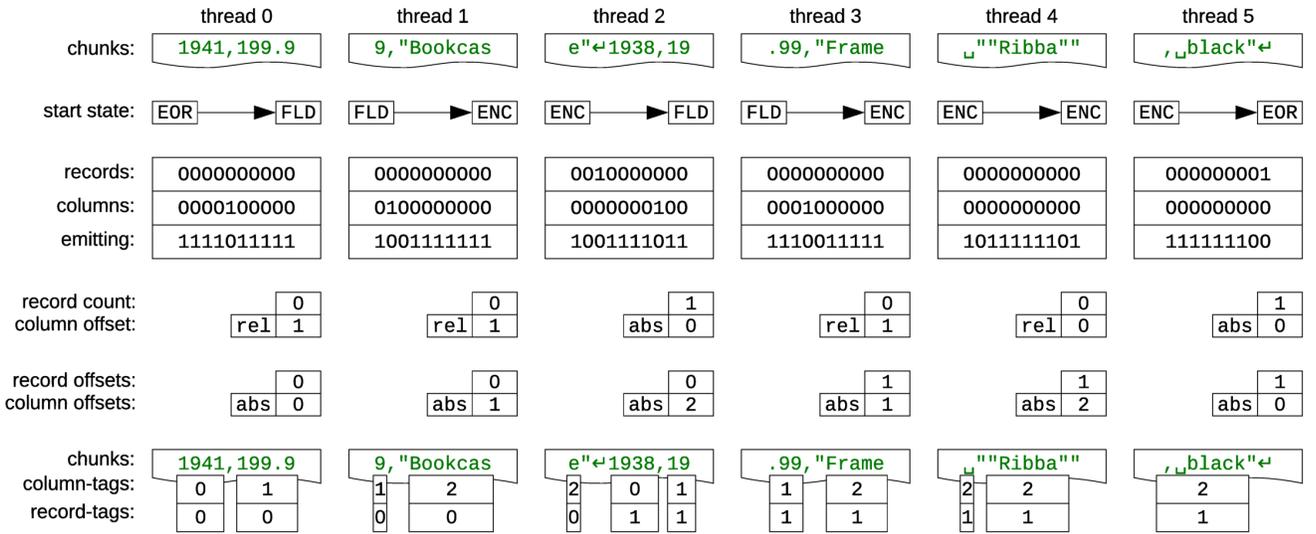


Figure 4: Identifying columns and records

By computing the composite of these state-transition vectors, the algorithm can deduce the starting state for every thread. We define the composite operation $a \circ b$ of two state-transition vectors a and b as:

$$\begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{|S|-1} \end{bmatrix} \circ \begin{bmatrix} b_0 \\ b_1 \\ \vdots \\ b_{|S|-1} \end{bmatrix} = \begin{bmatrix} b_{a_0} \\ b_{a_1} \\ \vdots \\ b_{a_{|S|-1}} \end{bmatrix}$$

Since the composite operation is associative, the algorithm can compute a parallel exclusive scan using the composite operation, which is seeded with the identity vector. After the exclusive scan, the i -th entry of each thread's resulting vector now corresponds to the state that the thread's DFA is supposed to start in, if the sequential DFA's starting state was s_i . For instance, if the sequential DFA's starting state was s_3 , each thread finds its starting state by reading the element at index three from its resulting vector.

Once each thread is aware of its starting state, threads can correctly interpret the symbols from their chunk by simulating a single DFA-instance. While iterating over its symbols, a thread identifies field delimiters, record delimiters, and other control symbols (e.g., an escape sequence) according to the parsing rules specific to the current format being parsed. Since the algorithm addresses delimiter-separated formats, the relevant meta data for each symbol can be represented using three bitmap indexes: one marking symbols that are delimiting a record, one flagging symbols that are delimiting a field, and one indicating whether a symbol is a control symbol (e.g., escape symbol) or whether it is part of the field. Subsequent steps can build on these bitmap indexes without requiring to repeatedly simulate the DFA-instance.

3.2 Identifying Columns and Records

The bitmap indexes from the previous step are used to identify the column and record offset. That is, the record and column that the first few symbols of a thread's chunk belong to, until it encounters the first delimiting symbol. Determining the column and record offsets requires two steps.

First, each thread computes the offset that its chunk adds to the preceding chunk's offset. For the records, the relative offset can easily be computed by counting the records (i.e.,

the number of set bits of a thread's record delimiter bitmap index using POPCNT). For columns, however, this is slightly more involved. If a thread encounters a record delimiter and therefore the beginning of a new record, it can infer the **absolute** column offset for the subsequent chunk (e.g., *thread 2* in Figure 4). Otherwise, all it can infer is that it has seen k field delimiters and, therefore, the next chunk's column offset has an additional offset of k , **relative** to the preceding chunk's column offset. In Figure 4, for instance, *thread 1* encounters one column delimiter but no record delimiter. As *thread 1* is not yet aware of its own column offset, it can only infer that the subsequent thread's column offset increases by one, relative to its own column offset. We distinguish between an absolute and a relative column offset, which are denoted as **abs** and **rel**, respectively, in Figure 4. A column offset is absolute, if there is at least one set bit in the thread's record delimiter bitmap index. The column offset can be computed by zeroing all bits of the column delimiter bitmap index that precede the last set bit in the record delimiter bitmap index, counting the remaining set bits, i.e.: $\text{POPCNT}(\sim\text{BLSMSK}(\text{rec.bidx}) \& \text{col.bidx})$

In a subsequent step, the algorithm computes the exclusive prefix sum over the record counts, which yields each thread's record offset. In order to retrieve the column offsets, we perform an exclusive scan using the following operation, where, for a column offset x , x_t denotes whether a column offset is relative (**rel**) or absolute (**abs**) and x_o denotes the offset value:

$$[a_t] \oplus [b_o] = \begin{cases} [b_t] & \text{if } b_t \text{ is abs} \\ [a_o + b_o] & \text{if } b_t \text{ is rel} \end{cases}$$

Once all absolute column and record offsets have been calculated, threads can correctly identify the column and record that each of its symbols belongs to. In preparation for the next step, which transforms the row-oriented input to a columnar format and, if applicable, converts strings of symbols to the data type of the corresponding column, we tag the symbols with the column and record they belong to, as illustrated at the bottom of Figure 4.

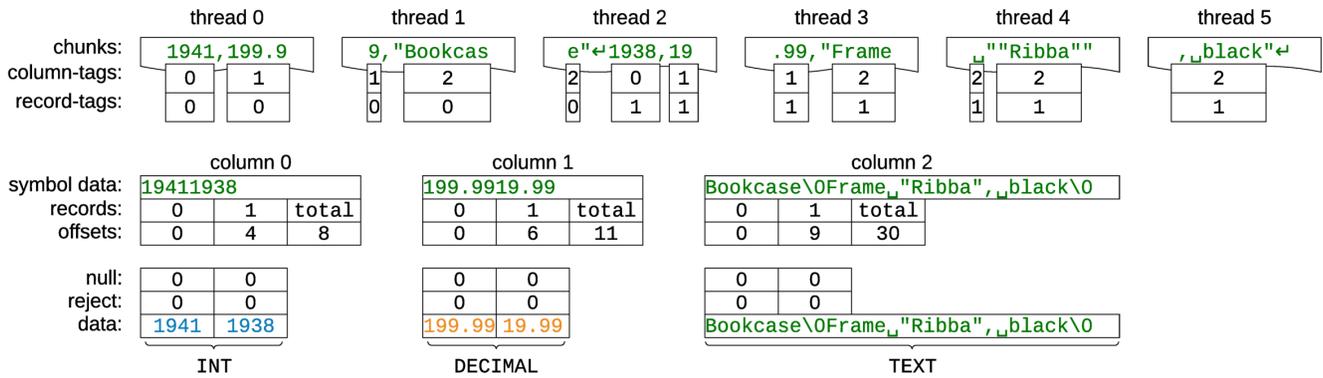


Figure 5: Preparing data in a columnar format and type conversion

3.3 Columnar Format & Type Conversion

Now, that each thread is fully aware of the associated columns and records, threads still need to generate the individual field values in a columnar format. Depending on the column that a string of symbols belongs to, this may require converting to the respective column's type (e.g., `int`, `float`). As shown at the top of Figure 5, symbols belonging to the same field may still span multiple chunks, requiring involved threads to collaborate on generating a single field value. To circumvent the collaboration between threads entirely, one option would be to change the assignment of threads, assigning exactly one thread to exclusively process all symbols required for generating a single field value. This approach, however, may cause considerable load-balancing issues, as the number of symbols per field may be subject to high variance. In particular, values of columns with variable-width types, such as text or Binary Large Objects (BLOBs), may be arbitrarily large.

Another challenge arises due to the fact that symbols are still in a row-oriented format. That is, two threads working on subsequent chunks may be parsing two different columns of two different types. Hence, they may require different rules, executing completely different code segments for parsing the fields' values. For instance, one thread may require generating an integer value, while the next one is extracting a date. The behaviour of different threads executing different code paths is particularly punishing on GPUs, where all threads within a warp (e.g., a group of 32 threads) are executing the same instruction in lockstep.

`ParPaRaw` addresses these challenges by first partitioning all symbols by the column they are associated with. During partitioning, `ParPaRaw` ensures that symbols within a column maintain their order by using a stable radix sort that uses the symbols' column-tags as the sort-key. While sorting, the symbols and the record-tags are moved along with the associated sort-key. The radix sort iterates over the bits of the column-tags, performing a stable partitioning pass on the sequence of bits considered with a given pass. A single partitioning pass involves (1) computing the histogram over the number of items that belong to each partition, (2) computing the exclusive prefix sum over the histogram's counts, and (3) scattering the items to the respective partition.

After partitioning, all symbols belonging to the same column lie cohesively in memory. We refer to all symbols be-

longing to the same column as the concatenated symbol string (CSS) of a column. The histogram that is maintained while sorting is used to identify the offsets of the columns' CSSs. Similar to the symbols, all the symbols' record-tags lie cohesively in memory, indicating which record a symbol belongs to.

Having all symbols in a columnar format allows the algorithm to efficiently process each of the columns. This may include type inference, validation, identifying NULLs, and converting symbol strings to the column's type. First, `ParPaRaw` uses the record-tags to generate an index into the CSS. The index is used to identify the offsets and lengths of the fields' symbol strings. To generate the index, the algorithm performs a run-length encoding on the symbols' record-tags, which yields each field's record and its number of symbols. Computing the exclusive prefix sum over the fields' symbol counts yields the offsets into the CSS, as shown in Figure 5. The symbol count of a field can be inferred using the difference of the successive field's offset and the field's own offset.

Building on the index, `ParPaRaw` can now start generating the fields' values by interpreting the strings of symbols, if that is required for a given column (e.g., numerical or temporal types). In order to address possible load-balancing issues due to having high variance in the number of symbols per field, we use three different collaboration levels: thread-exclusive, block-level, and device-level collaboration. By default, a thread tries to exclusively generate a field value, looking up the offset and number of its symbols in the index. Once the thread has identified the symbols, it starts converting the symbol string to the column's type (e.g., `int`, `float`). If, during lookup, a thread detects that its string of symbols exceeds a certain threshold, it will defer generating that field value for the block- or device-level collaboration. The threshold depends on the on-chip memory of a GPU's streaming multiprocessor and its number of cores. If there are fields left for the block-level collaboration, all threads of a thread-block (e.g., 64 threads) collaborate on generating a field value. Fields that exceed the on-chip memory available to a thread-block (typically in the order of tens of kilobytes) are addressed by the device-level collaboration. Block- and device-level collaboration use the same data parallel approach as the overall approach presented for parsing delimiter-separated inputs. Hence, the same technique for determining a thread's parsing context is employed.

4. EXTENSIONS & IMPLEMENTATION

Having presented the fundamental processing steps for a robust approach to massively parallel parsing in Section 3, this section focuses on optimisations, extensions, and implementation details. We develop two optimised specialisations that can be applied if a given input meets certain conditions (see Section 4.1). Section 4.2 addresses symbols crossing chunk boundaries, such as being encountered when dealing with variable-length encodings. To highlight that not only efficiency but also the approach’s applicability was of great importance to this work, we present a few more capabilities in Section 4.3. With an end-to-end streaming extension, we aim to hide the latency of data transfers via the PCIe bus (see Section 4.4). Finally, Section 4.5 presents how we address the major challenges of mapping the algorithm to the GPU.

4.1 Alternative Tagging Modes

ParPaRaw, as presented in Section 3, focuses on robustness. It is even resilient to inputs that contain records with a varying number of field delimiters per record (e.g., `"1,Apples\n2\n"`). This section focuses on presenting two optimised specialisations that are chosen, if the input provides a constant number of columns per record or if the user prefers to reject records that have an inconsistent number of field delimiters.

Since many of the presented processing steps work at peak memory bandwidth, reading and writing record-tags of four bytes increases the amount of memory transfers and degrades performance. Hence, we aim to lower the amount of memory transfers by reducing the memory footprint of the record-tags. As illustrated in Figure 6, we provide two alternatives to record-tags.

The *inline-terminated CSS* replaces delimiters with a terminator during the tagging phase. Just like the null character for null-terminated strings, the terminator is a unique character that indicates the end of a field’s symbols. Good candidates for terminators are various separators specified by the ASCII standard, such as the *record separator* (0x1E) or the *unit separator* (0x1F). To generate the CSS’s index, the algorithm simply writes the offsets of all occurrences of the terminator symbols to the index. The *inline-terminated CSS* requires that the terminator is not part of a column’s CSS, as those symbols would otherwise get confused for a terminator.

The *vector-delimited CSS* can address this scenario by devoting its own auxiliary boolean vector that delimits the fields within a column. The CSS’s index is generated the same way as for the *inline-terminated CSS* with the minor difference that the algorithm identifies non-zero values in the auxiliary vector instead of terminators from the CSS.

4.2 Variable-Length Symbols

So far, we have not addressed the challenge of symbols crossing chunk boundaries. While this can be easily prevented for fixed-size symbols spanning multiple bytes by adjusting the chunk size to be an integer multiple of the symbol size, it is more involved for variable-length symbols. For instance, if inputs are encoded using a variable-length Unicode Transformation Format (UTF), such as UTF-8 or UTF-16, symbol boundaries become unpredictable and some symbols might be crossing chunks. If a symbol crosses chunk boundaries, the thread working on the chunk at which the symbol

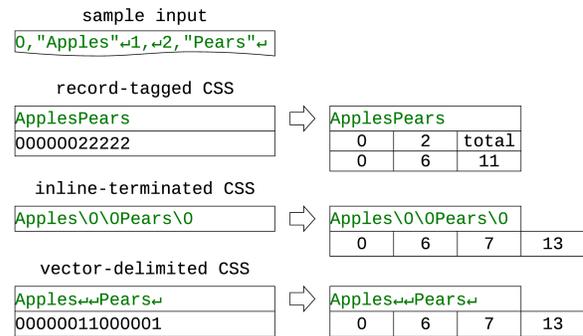


Figure 6: Alternative tagging modes

begins (i.e., the symbol’s leading bytes) is in charge of reading that symbol and transitioning the state of its DFA accordingly. Threads working on subsequent chunks that only read trailing bytes of a symbol skip those bytes. For the variable-length encodings UTF-8 and UTF-16, threads can identify whether the first bytes of a chunk are only trailing bytes of an encoded code point (a code point is a numerical value and most code points are assigned a character). UTF-8 encodes code points using one, two, three, or four bytes. Unless a single byte is used, all trailing bytes have the common binary prefix of `0b10XX XXXX`. Hence, for UTF-8 encoded inputs, threads simply ignore a chunk’s first few bytes with that binary prefix. UTF-16 uses either two bytes to encode code points ranging from `0x0000` to `0xD7FF` and from `0xE000` to `0xFFFF`, and four bytes for code points beyond `0x10000`. If four bytes are used, the two high order bytes, referred to as high surrogate, are in the range of `0xD800` to `0xD8FF`, and the low order bytes, referred to as low surrogate, are in the range of `0xDC00` to `0xDFFF`. Since unicode does not assign any characters in the range of `0xD800` to `0xDFFF`, there is no two-byte combination in that range. Hence, similar to UTF-8, a thread ignores a chunk’s first two bytes if their value is in the range of `0xDC00` to `0xDFFF`.

4.3 Capabilities

This section focuses on pointing out a few more capabilities to highlight **ParPaRaw**’s applicability to real-world requirements.

Validating format — One notable strength of **ParPaRaw** is its ability to simulate an FSM while parsing, which makes it widely applicable and enables more expressive parsing rules. With the presented massively parallel approach for simulating a DFA, **ParPaRaw** is always aware of the DFA’s current state when reading a symbol. Hence, invalid state transitions as well as a non-accepting end state can easily be detected.

Skipping records and selecting columns — **ParPaRaw** is able to ignore a user-specified set of records and columns. While tagging symbols with their associated column and record, all symbols that belong to records or columns that are supposed to be ignored are identified and marked as irrelevant. Irrelevant symbols can be ignored following the partitioning step.

Skipping rows – It is worth noting that rows are different from records, as some records may span multiple rows. Since ignoring rows may interfere with the assignment of symbols to columns and records, **ParPaRaw** has to ensure that rows are ignored early on. Hence, **ParPaRaw** ignores a set of rows

by performing an initial parallel pass over the input, pruning symbols of ignored rows (i.e., parallel stream compaction).

Inferring or validating number of columns — If no schema is provided and therefore the number of columns is not known a priori, **ParPaRaw** can infer the input’s number of columns. Similarly, if **ParPaRaw** is supposed to reject records that do not conform to the expected number of columns the same technique is applied. In either case, during DFA simulation threads need to track three values in addition to the relative or absolute column offset handed over to the subsequent chunk. Firstly, every thread keeps track of the number of field delimiters encountered before reading its very first record delimiter, which subsequently is referred to as *relative min/max*. Further, every thread maintains the minimum and maximum number of columns it counted per record for all records following the chunk’s first record delimiter. We use an extra bit to denote if no minimum and maximum was determined, i.e., the chunk does not contain any record delimiter. After the prefix scan of the column offsets, **ParPaRaw** can resolve the *relative min/max*, turning it into an absolute column offset. The absolute column offset is then incorporated in the respective chunk’s minimum and maximum column count. A subsequent reduction over the maximum is then used to infer the number of columns. Comparing the identified minimum and maximum column counts indicates whether a given chunk conforms to the expected number of columns per record.

Default values for empty strings — If the input has a consistent number of field-delimiters per record, the default value for empty strings is set during type conversion. That is, when field values are parsed, the empty string is parsed as the column’s default value. If the input does not have a consistent number of field-delimiters per record, the column’s data is pre-initialised with the user-specified default value and later overwritten for non-empty fields.

Type inference — **ParPaRaw** is comparably efficient when identifying a column’s type, as, prior to type conversion, all of a column’s symbols lie cohesively in memory. During an initial pass over the column’s symbols, threads identify the minimum numerical type being required to back their field value. A subsequent parallel reduction over the minimum type yields the inferred type of a column. **ParPaRaw** currently only considers type inference for numerical types, but can be extended to cover temporal types.

4.4 End-to-End Streaming

This section provides an extension to **ParPaRaw**’s on-GPU parsing algorithm presented in Section 3 to address inputs that do not reside on the GPU or exceed its available device memory. In order for the GPU to be able to process the input, the input first needs to be transferred via the comparably slow PCIe bus and, once processed, the parsed data has to be returned. It is worth noting that the PCIe bus allows for full-duplex communication, enabling simultaneous data transfers in either direction at peak bandwidth. While the PCIe bus does not necessarily limit the throughput, waiting for the data transfer to complete before and after parsing, respectively, adds a considerable amount of latency to the end-to-end processing time. Hence, rather than waiting for the input to arrive on the GPU, before the GPU begins processing it and, once finished, starts returning the parsed data, we make use of a streaming approach. The streaming approach splits the input into multiple par-

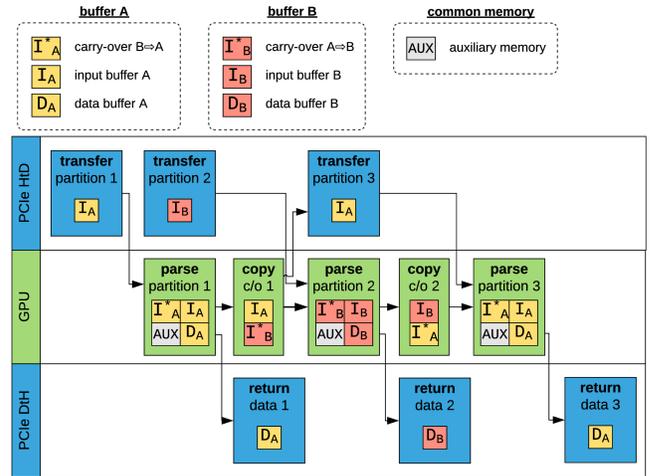


Figure 7: End-to-end streaming

titions. Each partition, at some point, is transferred to the GPU, processed, and its data is returned. Having multiple partitions allows to overlap these stages for subsequent partitions, similar to the pipelined approach in [38, 41]. That is, transferring a partition, while processing its predecessor on the GPU and simultaneously returning parsed data via the interconnect.

For the end-to-end streaming approach, we allocate a double-buffer and some auxiliary memory on the GPU (see top of Figure 7). Each buffer comprises memory for the raw input and the parsed data. One buffer’s raw input allocation is used as input for parsing on the GPU, while the opposing buffer’s raw input allocation is receiving data of the next partition. Similarly, one buffer’s data allocation is used to output parsed data, while data is being returned via the interconnect from the opposing buffer’s data allocation. In addition, we prepend additional memory for a carry-over to the memory allocated for the input of each buffer. The carry-over is used for prepending the last, incomplete record at the end of one buffer’s input to the opposing buffer’s input.

Figure 7 exemplifies the processing steps of the streaming parsing approach. The stages of a partition are (1) *transfer*: transferring the raw input of a partition from the host to the GPU, (2) *parse*: parsing the input of a given partition, including the prepended carry-over and writing the parsed data to the data buffer, and (3) *return*: returning the parsed data from the data buffer to the host. The resources required by each processing step are illustrated by the rectangular symbols within a step (e.g., I_A representing the memory allocated for the input of buffer A). A processing step’s dependency on a preceding processing step is depicted by an incoming edge. An important sequence depicted in Figure 7 is when the GPU switches work from one double-buffer to the other. For instance, after the GPU has finished parsing the input of the first partition (raw input provided by *input buffer A*), the last incomplete record is prepended to the second partition by copying it to the memory of the *carry-over* of *buffer B*. Since copying the carry-over is reading from *input buffer A*, the algorithm ensures that the transfer of the third partition to *input buffer A* does not take place before the carry-over has been copied, as the carry-over would otherwise get corrupted.

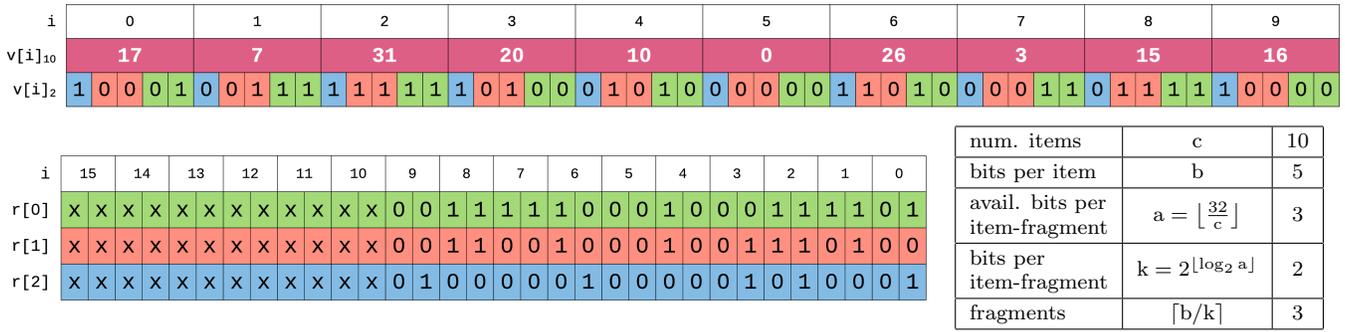


Figure 8: Logical and physical view of the multi-fragment in-register array

4.5 Implementation Details

This section addresses the main challenges faced when mapping ParPaRaw to the GPU. Specifically, we introduce a new data structure, referred to as multi-fragment in-register array (MFIRA), which provides a workaround for the constraint that threads cannot dynamically address into the register file. Since the register file is extremely fast and provides the most on-chip memory, addressing this shortcoming is a viable endeavour. The presented data structure allows to dynamically index into and access elements of a bounded array. MFIRA is particularly efficient for low-cardinality arrays of small integers. This is a recurring pattern in GPU programming, since the GPU’s threads need to be very lightweight, allowing for only very limited context (i.e., using only few registers). Hence, even though MFIRA was designed as an efficient data structure backing various objects when parsing, MFIRA likely would be useful for other use cases as well. Further, we present a branchless algorithm that builds on SIMD within a register (SWAR) to identify the index of a read symbol in the transition table. With that approach, we are able to keep the symbols that the algorithm compares against in the very fast register file. At the same time, it avoids that threads within a warp are executing along different branches.

Multi-fragment in-register array — The idea behind the data structure is that even though thread registers themselves cannot be addressed dynamically, individual bits within a register can be. Specifically, we use the intrinsic functions bit-field insert (BFI) and bit-field extract (BFE), which require only two clock cycles on recent microarchitectures, to efficiently access an arbitrary sequence of bits from a register. We use these two functions to decompose an item that is written to the data structure and distribute the item’s fragments (i.e., partitions of its bits) amongst one or more registers. Similarly, when an item is accessed, it is reassembled from its fragments. Figure 8 illustrates this principle, depicting an array containing up to ten items, each five bits wide. For such an array, the data structure could use up to three bits per fragment. To efficiently compute bit-offsets into a register, however, the number of bits actually being used by the data structure is chosen to be a power of two. This allows replacing the expensive integer multiplication with a bit-shift operation. In the example depicted in Figure 8, the data structure would therefore devote two bits per fragment, using a total of three fragments. The individual fragments of the items are colour-coded in Figure 8 to highlight how the logical view (top of the figure) maps to the physical view (bottom of the figure).

Table 1: Transition table example

symbols	groups	states					
		EOR	ENC	FLD	EOF	ESC	INV
$\backslash n$	0	EOR	ENC	EOR	EOR	EOR	INV
"	1	ENC	ESC	INV	ENC	ENC	INV
,	2	EOF	ENC	EOF	EOF	EOF	INV
*	3	FLD	ENC	FLD	FLD	INV	INV

Symbol matching using SWAR — During DFA simulation, the algorithm uses a transition table to identify the state transition from the DFA’s current state and a read symbol to the DFA’s new state. The transition table is two-dimensional, with states along one and symbols along the other dimension. To compress the transition table’s size, we collapse all the transition table’s symbols that have identical state transitions into symbol groups. As illustrated in Table 1, we have one symbol group per row instead of having symbol groups as columns, which allows coalesced access to all state transitions of a read symbol. This is particularly useful when computing the state-transition vectors. A thread reads a symbol from its chunk, identifies its symbol group, and fetches the row of state transitions for the matched symbol group. For each of its DFA instances, it can then efficiently determine the new state from that row.

Having introduced symbol groups, mapping a symbol to its symbol group is an elementary step. To ensure an efficient mapping, we exploit the fact that delimiter-separated formats typically have only a few symbols to distinguish amongst, such as an escaping symbol, field and record delimiters, and enclosing symbols like quotes or brackets (see Table 1). Hence, for the symbols we use a comparison-based approach, rather than devoting a full lookup-table that maps each character value to its group. Since symbols are often only eight bits wide (e.g., ASCII and UTF-8-encoded ASCII characters), while GPUs implement 32-bit wide arithmetic instructions, we use a branchless SWAR algorithm to perform multiple comparisons at a time (see Table 2). On the one hand, this avoids inefficiencies due to threads executing divergent branches. On the other hand, with the following approach, we are more space-efficient and are able to keep the symbols in the very fast register file. As illustrated in Table 2, we place each of the symbols that we try to match against in the individual bytes of four-byte registers. We refer to these registers as lookup-registers (LU-registers). For later comparison against the LU-registers, whenever a symbol is read, we replicate that symbol in every byte of

Table 2: Identifying a symbol’s index using SWAR

byte	7	6	5	4	3	2	1	0
symbol group			3	2	2	2	1	0
lookup (LU)				\t	—	,	”	\n
read symbol (s)	,	,	,	,	,	,	,	,
$c = LU \text{ XOR } s$	--	--	--	25	50	00	0E	26
$swar = H(c)$	--	--	--	00	00	80	00	00
$\mathbf{bfind}(swar) \gg 3$	1F	FF	FF	FF	00	00	00	02
$idx = \min_{v_{r_i}}(x)$								0x00000002
$\min(idx, 5)$								0x00000002
$H(x) = ((x - 0x01010101) \& (\sim x) \& 0x80808080)$								

a separate register (i.e., the s -register). Computing the exclusive or for each of the LU-registers with the s -register yields a null-byte if the two bytes match. Subsequently applying the bit-twiddling hack to determine a null-byte, as suggested by Mycroft in 1987 [34], sets the most-significant bit for that byte (see definition of $H(x)$ in Table 2). Using the intrinsic function `bfind`, we retrieve the position of the most-significant set bit. If no bit was set, i.e., the read symbol does not match any byte from the LU-registers, `bfind` will return `0xFFFFFFFF`. To retrieve the matching index, we divide the value returned by `bfind` by eight, using bit-shift for efficiency reasons (i.e., shifting it three bits to the right). For LU-registers that contain no match, the matching index is `0xFFFFFFFF`, while for the ones that contain a match, it yields a value between zero and three. To ensure that we consider a match, if present, we compute the minimum over all matching indexes. Finally, in case there was no match, we map the matching index of `0xFFFFFFFF` to the `catch-all` symbol group by using the minimum function. The minimum is computed very efficiently, requiring only one or two cycles on recent microarchitectures and is therefore generally preferable to a conditional expression.

5. EXPERIMENTAL EVALUATION

For the experimental evaluation we use two systems, one to evaluate CPU-only implementations, referred to as *CPU system*, and one system equipped with a GPU (*GPU system*) used for evaluating GPU-based approaches. Both systems are running Ubuntu 18.04. The *CPU system* has four sock-

ets, each equipped with a Xeon E5-4650 clocked at 2.70 GHz. It has a total of 512 GB of DRAM (DDR3-1600). The *GPU system* is equipped with 128 GB of DRAM (DDR4-2400) and a Xeon E5-1650 v4 processor with six physical cores, clocked at 3.60 GHz. The source code was compiled with the `-O3` flag. We used release 10.1.105 of the CUDA toolkit. The *GPU system* hosts an NVIDIA Titan X (Pascal) with 12 GB device memory, 3 584 cores, and a base clock of 1 417 MHz (driver version is 418.40.04).

The output of `ParPaRaw` is configured to comply with the format specified by *Apache Arrow*. *Apache Arrow* specifies a columnar memory format for efficient analytic operations [4]. It is used by a multitude of well-known in-memory analytics projects, such as *OmniSci*, *pandas*, and *Apache Spark*. For `ParPaRaw` we use a DFA that is capable of parsing any RFC4180 compliant input [37]. The DFA defines six states, including one state to track invalid state transitions.

To evaluate the systems, we choose the two dissimilar real-world datasets *yelp reviews* and *NYC taxi trips*. The *yelp reviews* dataset comprises 6.69 million reviews from yelp’s dataset as CSV, with all fields enclosed in double-quotes [45]. The dataset is 4.823 GB large with an average record size of 721.4 bytes per record. Each record is made up of nine columns, covering text-based, numerical, and temporal types. The dataset is of particular interest due to the text-based reviews that may include field and record delimiters, which poses a challenge for many parallel parsers.

The *NYC taxi trips* dataset is 9.073 GB large and comprises 102.8 million yellow taxi trips taken in the year 2018 provided by the *NYC Taxi & Limousine Commission* [43]. The dataset’s 17 columns cover numerical and temporal datatypes. With an average of only 88.3 bytes per record and 5.2 bytes per field, the majority of the fields are very short and of a numerical type, putting the emphasis on data type conversion.

5.1 On-GPU Parsing

This section provides a detailed evaluation of the presented algorithm using on-GPU workloads. Our on-GPU evaluation focuses on identifying efficient configurations and analysing the algorithm’s sensibility to input parameters. Time measurements for the on-GPU parsing experiments represent the GPU wall-clock time, measured using CUDA

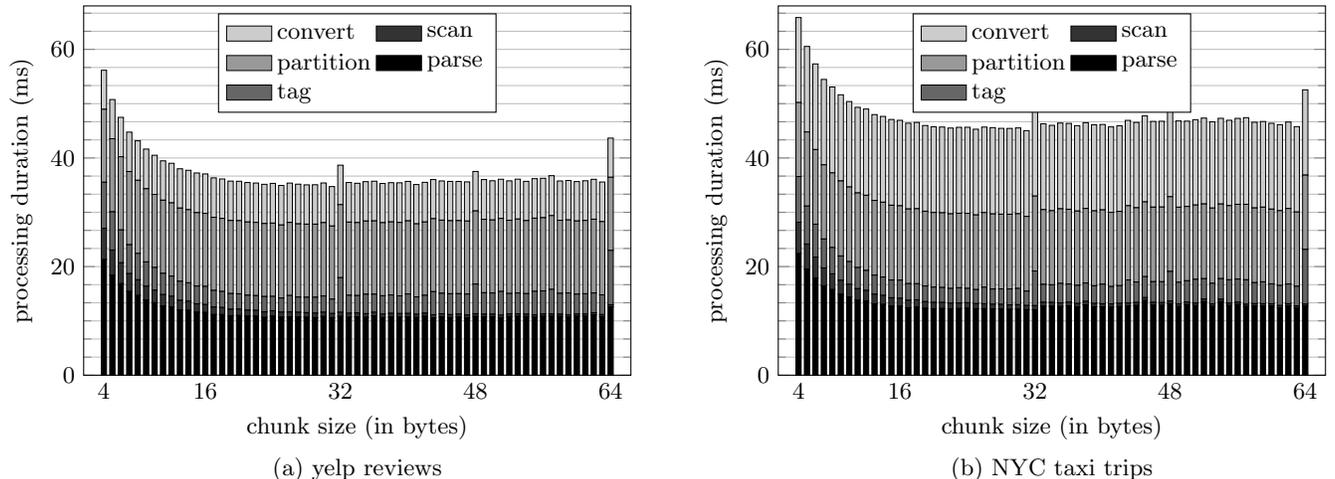


Figure 9: Time spent on individual processing steps depending on the chunk size configuration

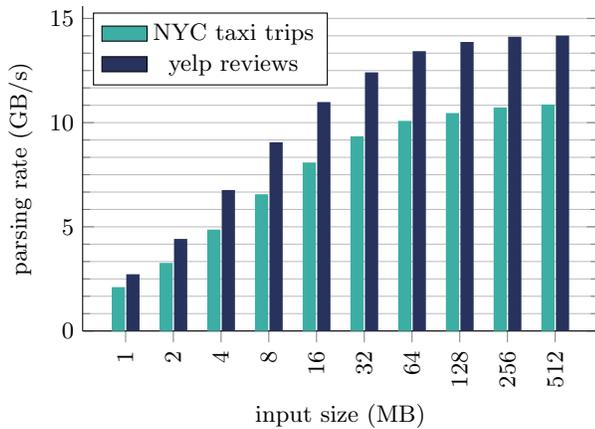


Figure 10: Parsing rate as a function of input size

events. Other than the end-to-end parsing experiments, on-GPU experiments do not include data transfers between host and device. Unless noted otherwise, we use the first 512 MB of each dataset for this evaluation, to be able to evaluate all tagging modes before running out of device memory.

We provide a breakdown of the time spent on the individual processing steps as a function of chunk size in Figure 9. Comparing the breakdown of the two datasets highlights the complexity of converting the many numerical and temporal types of the *NYC taxi trips* dataset, which, on average, make up only 5.2 bytes per value. The type conversion of the *NYC taxi trips* dataset accounts for roughly one third of the total processing time. Type conversion of the *yelp reviews* dataset, in contrast, only contributes approximately 20% to the total processing time, as the text-based reviews make up the majority of the raw record size. The analysis shows that the approach is mostly agnostic to choice of the chunk size, as long as it is reasonably large. Only for tiny chunk sizes of 15 bytes and less, the overhead of initialising and scheduling tens of millions of threads becomes noticeable. For a tiny chunk size, the ratio of actual work being done in relation to the time spent on initialising threads and the amount of meta data being written becomes unfavourable. Choosing a small chunk size is disadvantageous to parsing, tagging, and the prefix scan. As the prefix scan’s complexity is linear in the total number of chunks, its share of the processing time becomes noticeable when using very small chunks. The prefix scan takes less than two percent of the total processing time for most choices of the chunk size. The small spikes for parsing and tagging when using 32, 48, and 64 bytes per chunk, respectively, are due to shared-memory bank conflicts and bad occupancy. The best performance is achieved for 31 bytes per chunk, which will be used as default for the remaining evaluations.

Figure 10 shows **ParPaRaw**’s performance for various different input sizes. Parsing ten megabytes of the *yelp reviews* dataset in as little as one millisecond, **ParPaRaw** shows impressive performance even for small inputs, achieving a parsing rate of 9.75 GB/s. For even smaller inputs, **ParPaRaw** is able to process a single megabyte from either dataset in less than 500 μ s, corresponding to a parsing rate of more than 2.1 GB/s and 2.7 GB/s for the *NYC taxi trips* and the *yelp reviews* dataset, respectively. Even though the absolute performance is impressive, in particular when compared to available parsers (see Section 5.2), **ParPaRaw**’s efficiency

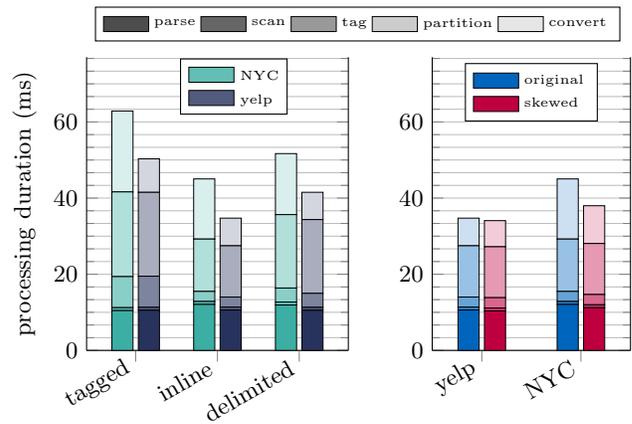


Figure 11: Time breakdown for different tagging modes (left) and skewed input (right)

degrades as the input size decreases. When parsing only five megabytes of either of both datasets, **ParPaRaw**’s performance achieves roughly 50% of its peak performance. A major reason for this, especially for inputs that are parsed in less than a millisecond, is the overhead due to the many kernel invocations during the type conversion step. During type conversion, there are multiple kernel invocations per column, required for the CSS-index generation as well as the type conversion itself. Hence, considering the many columns of the two datasets, kernel invocations, each with an estimated overhead in the order of roughly 5 - 10 μ s, account for a reasonable share of the few hundred microseconds that are required for parsing those tiny inputs.

We also analyse the performance of the different tagging modes (see Figure 11). Compared to the original inputs, the skewed inputs in Figure 11 (right) contain a single record that is 200 MB in size, while the remaining records remain the same. As expected, the use of record-tags (*tagged*) is noticeably slower than the two other tagging modes. In particular the tagging, partitioning, and type conversion steps take more time, as they depend on the choice of the tagging mode. The analysis also highlights the approach’s robustness, providing stable performance for the two dissimilar datasets, even if they are skewed (see Figure 11). On the one hand, the time breakdown shows that, except for the type conversion, all steps take roughly the same time for both datasets. Only type conversion, which involves generating data for more than an order of magnitude more fields in case of the *NYC taxi trips* dataset, shows perceivable performance differences. On the other hand, the approach shows robust performance even for highly skewed inputs.

5.2 End-to-End Parsing

For the end-to-end parsing experiments, we measured the CPU wall-clock time. Measurements include the time for reading the input from RAM and writing the parsed data back to system memory. For **ParPaRaw**, this includes data transfers between the host and the device. The end-to-end parsing approach was compared against *MonetDB*, *Apache Spark*, *pandas*, and the approach presented by Mühlbauer et al. [33] (*Inst. Loading*). In addition, we evaluated the GPU-based parser that is part of NVIDIA’s recently introduced open GPU data science project called *RAPIDS*. For *RAPIDS* we provide two evaluations. Firstly, simply read-

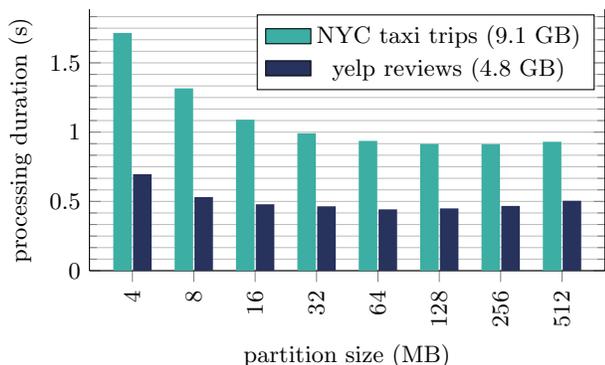


Figure 12: End-to-end parsing performance

ing the input into a GPU-based *DataFrame* called *cuDF* from where the data may be queried and processed with GPU support (*cuDF**). Secondly, exporting the parsed data to the host in the *Apache Arrow* columnar memory format using *cuDF*'s `to_arrow()` method (*cuDF*).

We analyse *ParPaRaw*'s performance depending on the chosen partition size (see Figure 12). Our evaluation shows that *ParPaRaw*'s performance increases with the partition size. Once the partition size grows beyond 128 MB for the *yelp reviews* and 256 MB for the *NYC taxi trips* dataset, however, the end-to-end processing duration starts growing again. This is due to the increased time for copying the very first partition and returning the parsed data of the very last partition (see Figure 7). Larger datasets compensate the effect of larger partitions. It is worth noting that this remains the only noticeable effect for larger inputs, since, with increasing input size, the number of partitions increases linearly, while the time per partition remains the same.

Figure 13 shows the time taken for parsing the respective input end-to-end. The performance numbers reported for parsing the 4.8 GB from the *yelp reviews* dataset highlight the strength of *ParPaRaw*, which takes only 0.44 seconds for the more challenging dataset. Only *cuDF*, which is still roughly 16 times slower than *ParPaRaw*, provides comparable performance. All CPU-based approaches, i.e., *MonetDB*, *Spark*, and *pandas*, are more than two orders of magnitude slower. Unfortunately, the implementation of *Inst. Loading* provided to us by the authors could not handle the *yelp* dataset due to its incomplete handling of quoted strings in parallel loads. Compared to *yelp reviews*, parsing of the *NYC taxi trips* dataset is easier to parallelise, as all line breaks correspond to record delimiters, making it trivial to identify the parsing context. Hence, even though parsing of the *NYC taxi trips* is computationally more expensive due to its many numerical and temporal fields, all CPU-based approaches benefit from the simpler format and see great improvements in the parsing rate. In particular, *Inst. Loading*, the approach proposed by Mühlbauer et al. [33], is about an order of magnitude faster than any other CPU-based implementation. Even though *Inst. Loading* is able to exploit the parallelism of the 32 physical cores for the *NYC taxi trips* dataset, *ParPaRaw* running on a single GPU is still roughly four times faster, despite the fact that *ParPaRaw* performs a full DFA simulation to keep track of the parsing context. Compared to the remaining approaches, *ParPaRaw* is more than ten times faster than *RAPIDS* loading the data into *cuDF* and over 40 times faster than the next best CPU-based approach.

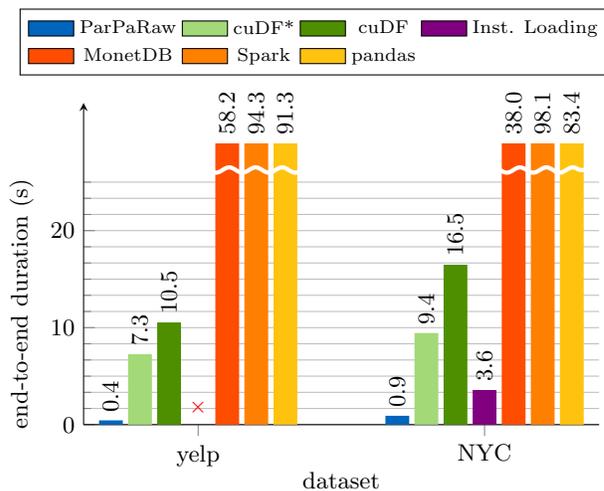


Figure 13: End-to-end performance comparison

6. CONCLUSIONS

This work presents *ParPaRaw*, a novel, massively parallel approach to parsing delimiter-separated formats. Other than the state-of-the-art that targets multicore and distributed systems with a coarse-grained approach, *ParPaRaw* is designed for fine-grained parallelism. Supporting parallelism even beyond the granularity of individual records makes it suitable for GPUs and ensures load balancing despite small chunks or large and varying record sizes. Being designed for scalability from the ground up with a data parallel approach that does not require any serial work, *ParPaRaw* is future-proof and can continue to gain speed-ups, as more cores are being added with future processors. Our approach identifies the parsing context (quotation scopes, comments, directives, etc.) without requiring a prior sequential pass. *ParPaRaw* is flexible and generally applicable. It supports even complex formats with involved parsing rules, as *ParPaRaw* is able to perform a massively parallel FSM simulation. State-of-the-art JSON parsers and the speculative approach by Ge et al., in contrast, have to deviate from the classic approach of using an FSM in order to be able to use SIMD vectorisation and speculation, respectively. This limits their applicability to other formats and requires designing completely different algorithms when confronted with another format (e.g., log files).

We show that *ParPaRaw* provides scalability without sacrificing applicability and flexibility. Achieving a parsing rate of as much as 14.2 GB/s, our experimental evaluation shows that *ParPaRaw* is able to scale to thousands of cores and beyond. With *ParPaRaw*'s end-to-end streaming approach, we are able to exploit the full-duplex capabilities of the PCIe bus while hiding latency from data transfers. For end-to-end workloads, *ParPaRaw* parses 4.8 GB of *yelp reviews* in as little as 0.44 seconds, including data transfers.

7. ACKNOWLEDGMENTS

This research has been supported in part by the Alexander von Humboldt Foundation. We would like to thank the authors of "*Instant loading for main memory databases*", in particular Thomas Neumann, for providing their implementation [33].

8. REFERENCES

- [1] NVIDIA Tesla V100 GPU Architecture. Whitepaper. Technical report, NVIDIA, 2017.
- [2] I. Alagiannis, R. Borovica, M. Branco, S. Idreos, and A. Ailamaki. Nodb: Efficient query execution on raw data files. *SIGMOD*, 2012.
- [3] Alenka. Alenka - a gpu database engine. <https://github.com/antonmks/Alenka>, 2012.
- [4] Apache Software Foundation. Apache Arrow. <https://arrow.apache.org>, 2019.
- [5] A. Arunkumar, E. Bolotin, B. Cho, U. Milic, E. Ebrahimi, O. Villa, A. Jaleel, C.-J. Wu, and D. Nellans. Mcm-gpu: Multi-chip-module gpus for continued performance scalability. *SIGARCH*, 2017.
- [6] T. Azim, M. Karpathiotakis, and A. Ailamaki. Recache: Reactive caching for fast analytics over heterogeneous data. *PVLDB*, 11(3):324–337, 2017.
- [7] S. Blanas, K. Wu, S. Byna, B. Dong, and A. Shoshani. Parallel data analysis directly on scientific file formats. *SIGMOD*, 2014.
- [8] G. E. Blelloch. Scans as primitive parallel operations. *IEEETransComp*, 1989.
- [9] D. Bonetta and M. Brantner. Fad.js: Fast json data access using jit-based speculative optimizations. *PVLDB*, 10(12):1778–1789, 2017.
- [10] R. P. Brent and H. T. Kung. A regular layout for parallel adders. *IEEETransComp*, 1982.
- [11] B. Chandramouli, D. Xie, Y. Li, and D. Kossmann. Fishstore: Fast ingestion and indexing of raw data. *PVLDB*, 12(12):1922–1925, 2019.
- [12] Y. Cheng and F. Rusu. Parallel in-situ data processing with speculative loading. *SIGMOD*, 2014.
- [13] Y. Dotsenko, N. K. Govindaraju, P.-P. Sloan, C. Boyd, and J. Manferdelli. Fast scan algorithms on graphics processors. *ICS*, 2008.
- [14] A. Dziejdzic, M. Karpathiotakis, I. Alagiannis, R. Appuswamy, and A. Ailamaki. Dbms data loading: An analysis on modern hardware. *DaMoN*, 2016.
- [15] C. N. Fischer. On parsing context free languages in parallel environments. Technical report, 1975.
- [16] C. Ge, Y. Li, E. Eilebrecht, B. Chandramouli, and D. Kossmann. Speculative distributed csv data parsing for big data analytics. *SIGMOD*, 2019.
- [17] S. Ha and T. Han. A scalable work-efficient and depth-optimal parallel scan for the gpgpu environment. *TPDS*, 2013.
- [18] R. Hai, S. Geisler, and C. Quix. Constance: An intelligent data lake system. *SIGMOD*, 2016.
- [19] P. M. Hallam-Baker and B. Behlendorf. Extended Log File Format. <https://w3.org/TR/WD-logfile>, 1996.
- [20] W. D. Hillis and G. L. Steele Jr. Data parallel algorithms. *CACM*, 1986.
- [21] S. Idreos, I. Alagiannis, R. Johnson, and A. Ailamaki. Here are my data files. here are my queries. where are my results? *CIDR*, 2011.
- [22] M. Ivanova, Y. Kargin, M. Kersten, S. Manegold, Y. Zhang, M. Datcu, and D. E. Molina. Data vaults: A database welcome to scientific file repositories. *SSDBM*, 2013.
- [23] R. Johnson and I. Pandis. The bionic dbms is coming, but what will it look like? *CIDR*, 2013.
- [24] Kaggle. Kaggle datasets. <https://www.kaggle.com/datasets>, 2018.
- [25] M. Karpathiotakis, I. Alagiannis, and A. Ailamaki. Fast queries over heterogeneous data through engine customization. *PVLDB*, 9(12):972–983, 2016.
- [26] M. Karpathiotakis, M. Branco, I. Alagiannis, and A. Ailamaki. Adaptive query processing on raw data. *PVLDB*, 7(12):1119–1130, 2014.
- [27] P. M. Kogge and H. S. Stone. A parallel algorithm for the efficient solution of a general class of recurrence equations. *IEEETransComp*, 1973.
- [28] G. Langdale and D. Lemire. Parsing gigabytes of JSON per second. *CoRR*, 2019.
- [29] Y. Li, N. R. Katsipoulakis, B. Chandramouli, J. Goldstein, and D. Kossmann. Mison: A fast json parser for data analytics. *PVLDB*, 10(10):1118–1129, 2017.
- [30] A. Luotonen. Logging Control In W3C httpd. <https://www.w3.org/Daemon/User/Config/Logging.html>, 1995.
- [31] D. Merrill and M. Garland. Single-pass parallel prefix scan with decoupled look-back. *NVIDIA, Tech. Rep. NVR-2016-002*, 2016.
- [32] D. Merrill and A. Grimshaw. Parallel scan for stream architectures. Technical report, 2009.
- [33] T. Mühlbauer, W. Rödiger, R. Seilbeck, A. Reiser, A. Kemper, and T. Neumann. Instant loading for main memory databases. *PVLDB*, 6(14):1702–1713, 2013.
- [34] A. Mycroft. String Processing Instruction. <https://groups.google.com/forum/embed#!topic/comp.lang.c/2HtQXvg7iKc>, 1987. *comp.lang.c*.
- [35] S. Palkar, F. Abuzaid, P. Bailis, and M. Zaharia. Filter before you parse: Faster analytics on raw data with sparser. *PVLDB*, 11(11):1576–1589, 2018.
- [36] RAPIDS. Rapids - open gpu data science. <https://rapids.ai>, 2012.
- [37] Y. Shafranovich. RFC4180 - Common format and MIME type for comma-separated values (CSV) files. <https://tools.ietf.org/html/rfc4180>, 2005.
- [38] A. Shahvarani and H.-A. Jacobsen. A hybrid b+-tree as solution for in-memory indexing on cpu-gpu heterogeneous computing platforms. *SIGMOD*, 2016.
- [39] Simantex. Simantex - csvimporter. <https://github.com/Simantex/CSVImporter>, 2012.
- [40] J. Sklansky. Conditional-sum addition logic. *IRE Transactions on Electronic Computers*, 1960.
- [41] E. Stehle and H.-A. Jacobsen. A memory bandwidth-efficient hybrid radix sort on gpus. *SIGMOD*, 2017.
- [42] Sumo Logic. Press release. <https://www.sumologic.com/press/2018-02-27/growth-milestones>, 2018.
- [43] Taxi and Limousine Commission. Tlc trip record data. http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml, 2016.
- [44] S. Yan, G. Long, and Y. Zhang. Streamscan: Fast scan algorithms for gpus without global barrier synchronization. *PPoPP*, 2013.
- [45] Yelp Inc. Yelp Dataset Challenge. www.yelp.com/dataset/challenge, 2019.
- [46] W. Zhao, Y. Cheng, and F. Rusu. Vertical partitioning for query processing over raw data. *SSDBM*, 2015.