



HAL
open science

Supporting Descendants in SIMD-Accelerated JSONPath

Mateusz Gienieczko Envelope, Filip Murlak Envelope, Charles Paperman
Envelope

► **To cite this version:**

Mateusz Gienieczko Envelope, Filip Murlak Envelope, Charles Paperman Envelope. Supporting Descendants in SIMD-Accelerated JSONPath. International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2024), 2024, San Diego (California), United States. 10.4230/LIPIcs . hal-04398350

HAL Id: hal-04398350

<https://hal.science/hal-04398350>

Submitted on 16 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Supporting Descendants in SIMD-Accelerated JSONPath

Mateusz Gienieczko  

Univeristy of Warsaw, Poland

Filip Murlak  

Univeristy of Warsaw, Poland

Charles Paperman  

CRIStAL, Université de Lille, INRIA

Abstract

Harnessing the power of SIMD can bring tremendous performance gains in data processing. In querying streamed JSON data, the state of the art leverages SIMD to fast forward significant portions of the document. However, it does not provide support for descendant, which excludes many real-life queries and makes formulating many others hard. In this work, we aim to change this: we consider the fragment of JSONPath that supports child, descendant, wildcard, and labels. We propose a modular approach based on novel depth-stack automata that process a stream of events produced by a state-driven classifier, allowing fast forwarding parts of the input document irrelevant at the current stage of the computation. We implement our solution in Rust and compare it with the state of the art, confirming that our approach allows supporting descendants without sacrificing performance, and that reformulating natural queries using descendants brings impressive performance gains in many cases.

2012 ACM Subject Classification Author: Please fill in 1 or more `\ccsdsc macro`

Keywords and phrases json, jsonpath, simd, query language, data management

Digital Object Identifier 10.4230/LIPIcs...

Acknowledgements This work was supported by Poland's National Science Centre grant 2018/30/E/ST6/00042. We salute the reviewers, whose insightful comments helped us improve both our work and its presentation.

1 Introduction

JSON is the format of choice for both modern web communication and large datasets. Due to its prevalence, all modern programming frameworks provide some facility for JSON processing. While the technology is relatively mature, substantial performance gains can be still achieved by exploiting the Single Instruction, Multiple Data (SIMD) capabilities of modern commodity processors [20, 23, 24]. Langdale and Lemire harnessed SIMD instructions to parse JSON data, validate it and produce its tree representation (the DOM, Document Object Model), achieving impressive speed-ups over conventional parsers [22]. While DOM allows us to query the document efficiently, it is prohibitively costly for large datasets. Not only does it take up massive amounts of RAM, but constructing it would also take most of the query execution time – parsing can amount to up to 90% of time spent in such an application, while the actual query only touches a small portion of the input data [25]. When faced with terabytes of data to query, the only feasible solution is a streaming algorithm with minimal memory footprint.



© M. Gienieczko and F. Murlak and C. Paperman;
licensed under Creative Commons License CC-BY 4.0
Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1.1 State of the Art

General tools for querying streamed JSON data, such as JsonSurfer [32] and jq [12], are slow: the throughput of JsonSurfer oscillates around 200MB/s and jq is an order of magnitude slower. Meanwhile, Langdale and Lemire’s `simdjson` can parse gigabytes of data per second, providing access via a SAX-like on-demand API [30]. Jiang and Zhao’s recent JSONSki [21] shows that certain queries can be evaluated even faster, surpassing the throughput of `simdjson` thanks to clever SIMD-accelerated fast-forwarding through irrelevant fragments of the stream.

JSONSki supports a useful but limited subset of the popular JSONPath query language [17]. It has no support for descendant selectors, and their wildcard selector implements only a part of the JSONPath specification, stepping into every entry of an array, but not into every field of an object. Importantly, JSONSki relies on knowing whether a selector acts on objects or lists, which the authors identify as an obstacle for supporting descendant, and which is also incompatible with idiomatic wildcard.

1.2 Motivation

Full support for wildcards and descendants is a highly desirable feature. Idiomatic wildcards allow accessing elements without specifying the full path. Descendants allow reaching deep down the document and accessing elements at multiple depths with a single query. When they are supported, tasks like fetching all values associated with a given property name in the document become very easy, as they can be succinctly represented with a descendant query. For example, one could scrape all `url` property values from a document without knowing anything about the paths leading to them, whereas without descendants the user would need to know the depth of the property, and without full wildcard support also specify all labels on the path, leading to an explosion of possibilities. A more general example is the code-as-data application scenario (for instance, harvesting code for AI [13]), fueled by the proliferation of large code repositories and archives such as Software Heritage [28]. Tools such as `cLang` output abstract syntax trees of programs as JSON documents, deep and highly irregular. Exploring such documents without wildcard and descendant is infeasible.

1.3 Our Contribution

We present an engine supporting JSONPath queries with labels, `child`, wildcard, and descendant selectors that does not rely on the advance knowledge of types of processed elements. We propose a novel modular approach based on *depth-stack automata* that process a stream of events produced by a *state-driven classifier* capable of fast-forwarding through currently irrelevant fragments of the stream using SIMD acceleration.

1.3.0.1 Abstract automaton execution.

The classifier consumes the raw JSON stream and produces events associated with symbols meaningful for the query, such as structural symbols and relevant labels. Every step of the automaton is costly, but the classifier generates only the necessary events, allowing the automaton to skip over most of the raw input stream. This modular architecture allows one to compare different classifiers and different automata models. It also provides a general abstraction separating fast branchless SIMD-enhanced stream processing from the heavily branching sequential code implementing the query logic, potentially of use elsewhere.

1.3.0.2 Sparse stack representation.

In the described model, queries combining child and descendant segments can be easily executed by a push-down automaton, but using the stack is potentially costly. Depth-register automata [7] are stackless, but they cannot handle all queries mixing child and descendant. Aiming to get the best of both worlds, we propose *depth-stack automata* which are equivalent to pushdown automata but offer a flexible sparse representation of the stack, allowing to keep its depth to a bare minimum.

1.3.0.3 State-driven classifier.

Because the automaton cares about different events in different states, additional savings can be made by adjusting the set of recognized symbols depending on the current needs of the automaton. This way we reduce the workload of the automaton (fewer events to process), and allow the classifier to pick an algorithm optimized for a particular character set. The idea is implemented at two levels. Internally, our main classifier is capable of *tooggling* specific symbols on the fly. Externally, our *multi-classifier pipeline* allows switching dynamically between the main classifier and additional specialized classifiers. In principle, there is an optimal classifier for each state of the automaton, but the cost of switching often exceeds the gain. This is why we do not switch whenever a state change occurs, but only when the expected benefits justify it. In the remaining cases we rely on the flexibility of the main classifier.

1.3.0.4 Implementation and Experiments.

We implemented our solution in Rust. The resulting tool, dubbed *rsonpath*, is available online [15]. In a series of experiments we compare *rsonpath* with the state of the art, and confirm that our approach allows supporting descendant and idiomatic wildcard without sacrificing performance (actually, with 10-20% performance boost) and that by using descendants some natural queries can be sped up impressively (up to an order of magnitude).

1.4 Outline

In Section 2 provide the necessary background on JSON and JSONPath. Then we explain how the query logic is implemented using depth-stack automata (Section 3). Next, we describe the classifiers and how they are orchestrated to provide the automaton with suitable access to the stream (Section 4). Section 5 is devoted to the experimental evaluation. We conclude in Section 6.

2 Background

JSON is a serialization format for JavaScript objects. A JSON document J is one of the following:

- an *atomic value*: a string, a number, or one of the special values `true`, `false`, `null`;
- an *array* of the form $[J_1, J_2, \dots, J_n]$ where J_1, J_2, \dots, J_n are JSON documents;
- an *object* of the form $\{\ell_1 : J_1, \ell_2 : J_2, \dots, \ell_n : J_n\}$ where here J_1, J_2, \dots, J_n are JSON documents and $\ell_1, \ell_2, \dots, \ell_n$ are strings (called *property names* or *labels*).

We call J_1, J_2, \dots, J_n *direct subdocuments* (or *children*) of J . A string is a sequence of Unicode characters, wrapped in double quotes, using backslash escapes. For instance, `{"a": "\b":2022}"` is an object with a single property `a` whose value is the string

XX:4 Supporting Descendants in SIMD-Accelerated JSONPath

`{"b":2022}`. The escaping mechanism poses additional challenges for rapid processing of JSON documents.

While JSONPointer [33] can point to a specific subdocument, using property names and array indices, a JSONPath query selects all subdocuments accessible by a path specified using more flexible selectors. We work with a fragment of JSONPath that uses selectors `$`, `.l`, `.*`, and `..l`; that is, we allow expressions given by the grammar

$$e ::= \$ \mid e.l \mid e.* \mid e..l$$

where l is a property name. These selectors allow navigating to, respectively, the document root, the property l of the current element, any direct subdocument (object property or array entry) of the current element, the property l of the current element or any of its subdocuments.

More precisely, we apply the *node semantics*, defined formally as follows. When evaluated over a JSON document J , an expression e returns a (possibly empty) set $e(J)$ of subdocuments of J , in the order in which they occur in J , distinguishing different occurrences of syntactically identical subdocuments. The expression $\$$ returns J : $\$(J) = \{J\}$. Suppose

$$e(J) = \{J_1, J_2, \dots, J_n\}$$

with $n \geq 0$. Then, for $s = .l$ or $s = .*$ or $s = ..l$,

$$es(J) = s(J_1) \cup s(J_2) \cup \dots \cup s(J_n),$$

where $s(J_i)$ is the set of subdocuments selected by the selector s in J_i : $.l(J_i)$ is the singleton of the value of property l in J_i or the empty set if J_i is an array or does not have property l ; $.*(J_i)$ is the set of all direct subdocuments of J_i ; and $..l(J_i)$ is the set of values of property l in all objects in J_i , including J_i itself; that is,

$$..l(J_i) = .l(J_i) \cup ..l(J_i^1) \cup ..l(J_i^2) \cup \dots \cup ..l(J_i^{n_i}),$$

where $J_i^1, J_i^2, \dots, J_i^{n_i}$ are the direct subdocuments of J_i . For example, in the document `{a: [{b: {c: 1}}, {b: [2]}]}`, the query `$.a..b.*` returns 1 and 2.

There exists an alternative *path semantics*, in which it matters how a given subdocument is reached. For example, in the document `{a: {a: {a: {b: 42}}}}` the query `$.a..b` can select 42 by matching the selector `..a` either to the child or to the grandchild or to the great-grandchild of the root. Under the node semantics it does not matter. Under the path semantics this is reflected by returning 42 three times: once for each way it can be reached. Path semantics can be defined formally by replacing sets with multisets in the definition above.

Using the JSONPath comparison project [9] we found that most existing implementations of JSONPath use path semantics: out of 44 tested implementations, 34 use path semantics (including the original implementation by Gössner [17]), while only 6 use node semantics (4 were errors); see Appendix E for details. The upcoming RFC has also settled on path semantics [18].

We posit that, despite its prevalence, path semantics is undesirable. The reasons are three-fold.

First, cluttering results with duplicated answers is usually not what the user wants and deduplication *a posteriori* is provably impossible without additional information identifying returned answers with subdocuments of the input document. For instance, under path semantics, on the document `{a: {a: {b: 42}}}` the query `$.a..b` returns 42 twice, each time

referring to the same subdocument. On $\{a: [\{b:42\}, \{b:42\}]\}$ the same query also returns 42 twice, but this time referring to different subdocuments. Under node semantics, 42 is returned once on the first document, and twice on the second document. As the results on the two documents under path semantics are indistinguishable, there is no way to recover node semantics from path semantics.

Second, the multiplicity blow-up can be not only inconvenient, but even harmful. Indeed, by generalising the above example we can see that the result set might be exponentially large in the length of the query. As the implementation of JSONPath in PostgreSQL also adheres to path semantics, one can exploit this blow-up in an attack against a PostgreSQL database by inserting into the database compact antagonistic documents targeting specific queries, expected to be run there.

Third, path semantics is actually harder to compute. One might think it is easier, as it does not require eliminating duplicates, but this intuition is correct only for a naive recursive implementation. An efficient implementation will traverse the document only once, naturally producing no duplicates. As path semantics obviously cannot be recovered from node semantics *a posteriori*, reproducing the multiplicities of path semantics would incur additional cost throughout the computation.

From now on, we consider only node semantics, as not only more useful and safer, but also better suited for optimization in our streaming model, which inherently demands a linear pass over the document.

3 Query execution

Our query execution algorithm has two phases. In the first phase, detailed in Section 3.1 below, the query is compiled into a deterministic *query automaton* recognizing access paths leading to subdocuments selected by the query. In the second phase, the query automaton is simulated over the streamed document. The simulation uses a concise stack representation (Section 3.2) as well as four skipping techniques (Section 3.3) allowing it to fast-forward through irrelevant fragments of the stream. The implementation (Section 3.4) abstracts away access to the stream as an iterator, which can be seen as a SIMD-enhanced lexer, capable of filtering out irrelevant tokens on demand. The iterator is the workhorse of rsonpath; we discuss it in detail in Section 4.

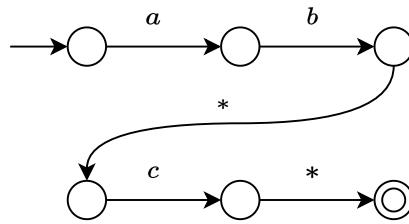
3.1 Constructing query automata

A JSONPath query can be naturally represented as a finite automaton that runs on a word formed by the sequence of labels on a path from the root to a node in the document tree (array entries are given an artificial label, different from property names).

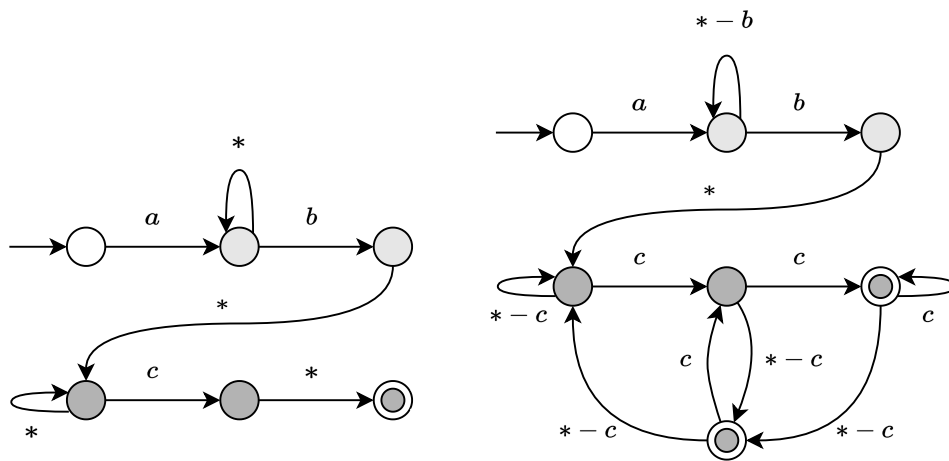
For descendant-free queries the situation is simple. Consider a query of the form $$.l_1.l_2 \cdots .l_n$, where each l_i is either a label or a wildcard. An automaton for such a query has a very regular form, illustrated in Figure 1. It consists of a chain of states corresponding to the selectors of the query. From selector $.l_i$, the automaton moves to selector $.l_{i+1}$ when it sees a matching label. That is, if l_i is a label it must see l_i itself, and if l_i is wildcard every letter is fine. From the last selector, the automaton moves to the accepting state (when it sees a matching label). Importantly, the automaton is a *deterministic finite automaton* (DFA), and as such it is ready for simulation.

For queries with descendants, things get more complicated. One can still construct an automaton with states corresponding to selectors, but in states corresponding to descendants, the automaton can not only pass to the next state when it sees a matching label, but it can

XX:6 Supporting Descendants in SIMD-Accelerated JSONPath



■ **Figure 1** Minimal DFA recognising paths matching query `$.a.b.*.c.*` (trash state omitted).



■ **Figure 2** NFA (top) and minimal DFA (bottom) recognising paths matching query `$.a..b.*..c.*`. Different shades of grey represent segments in the NFA and corresponding components in the DFA.

also loop regardless of the seen label; see Figure 2 (top). We call such states *recursive*. States corresponding to child selectors are called *direct*. As in recursive states the automaton has sometimes two choices for a single label, it is a *nondeterministic finite automaton* (NFA). To ease the simulation, we need to turn the NFA into a minimal DFA. One can, of course, use a general method to determinize and minimize an NFA, but it is useful to understand certain common structural properties of minimal DFAs obtained from JSONPath queries.

Descendant selectors split the query into segments. Each segment begins with a descendant selector (or `$`) and ends just before the next descendant selector (or at the end of the query). For example, the query `$.a..b.*..c.*` consists of three segments: segment `$.a`, segment `..b.*`, and segment `..c.*`. The following *greedy match property* is key: when searching for query matches in the document, as soon as we find a match for a segment on some branch, we can ignore any further matches of the segment on this branch, and start looking for a match for the next segment, etc. Note that we crucially rely on node semantics here. Under the path semantics, we would have to consider each subsequent match too, as they would lead to different matches of the whole query even if the final selected node is the same. For example, for the query `$.a..b.*..c.*` and document `{a:{b:{b:{b:{c:[42]}}}}}`, we would have to consider matches of segment `..b.*` at all `b` nodes. Ultimately we would discover that the first two matches can be extended to the whole query, but the third one cannot, and we would report two matches in total (even though there is only one way to match the last segment). Under node semantics, we can assume that `..b.*` is matched at the first `b` and disregard subsequent matches.

The greedy match property implies that once we reach a given recursive state in the NFA, we can forget about all previous states. This breaks the automaton down into segments such that only one segment needs to be simulated at a time. Consequently, each segment can be determinized separately, usually—but not always—yielding a single strongly connected component of the DFA; see Figure 2 (bottom). Each state has some transitions over concrete labels (typically advancing the matching), and a single *fallback transition* over the remaining labels (typically leading to the *initial state of the component*, corresponding to a recursive state of the NFA). The segment of the NFA corresponding to selectors before the first descendant is always deterministic and is simply copied to the DFA. In this component, fallback transitions from non-wildcard states lead to an all-rejecting trash state (not shown in Figure 2). Other segments may easily yield exponentially large components. For example, the query $\$.a.*.*.\dots.*$, selecting nodes whose ancestor n levels up has label a , reconstructs the classical example of exponential blowup when passing from an NFA to a DFA.

3.2 Simulating query automata on streamed trees

With the determinized query automaton at hand, executing a query boils down to simulating this automaton on all paths in the tree. The tree, however, is streamed. For now, let us take a high-level view and treat the streamed document as a sequence of tokens: structural characters ‘{’, ‘}’, ‘[’, ‘]’, ‘:’, ‘,’ , labels, and atomic values. In a pass over the streamed document we can simulate a DFA easily, as long as we use a stack:

- whenever an opening character (‘[’ or ‘{’) is encountered, the state of the DFA is pushed to the stack;
- each encountered label (including the “non-label” of array entries) triggers a transition of the DFA and, if the new state is accepting, an answer is reported;
- a closing character (‘]’ or ‘}’) pops the stack and restores the state to what it was before visiting the subtree.

In this work, however, we aim at minimising the use of the stack, the hypothesis being that such code will have fewer branches and thus should be more performant when paired with SIMD processing.

Querying streamed trees in a stackless manner was investigated in [7], where the authors characterise the kinds of queries that can be effectively executed on *depth-register automata*, which are finite automata with a constant number of depth registers and access to the current depth in the tree. The only operations allowed on registers are storing the current depth and comparing whether the stored value is less than, equal to, or greater than the current depth. It is easy to see that such automata can be effectively implemented, as the current depth can be tracked in a single integer variable that is incremented on every occurrence of an opening character and decremented on every occurrence of a closing character.

Automata resulting from descendant-only queries can be simulated stacklessly in the depth-register model. A stackless algorithm for the query $\$.l_1.l_2 \dots .l_n$ uses depth registers $\delta_1, \dots, \delta_{n-1}$ and states $1, 2, \dots, n+1$. We start in state 1 and report an answer whenever in state $n+1$. When in state i , there are two kinds of events that trigger a transition:

- if the current depth falls to the value in register δ_{i-1} , move to state $i-1$ (not applicable when $i=1$);
- if label l_i is found, set δ_i to the current depth and move to state $i+1$ (not applicable when $i=n+1$).

For automata resulting from queries mixing descendant and child selectors, the depth-register model is too weak [7]. Intuitively, this is due to the non-local nature of such queries:

two children of the same node can be arbitrarily far away from each other in the input JSON string. For instance, for the query $\$.l_1.l_2$ the automaton would have to remember every occurrence of label l_1 on the current path from the root in order to be able to check all its children, because the children of shallower l_1 nodes might occur both before and after the children of deeper l_1 nodes. While all DFAs can be simulated using a stack, this becomes costly when the stack gets large. As a remedy, we generalize the depth-register model by employing a *depth-stack*. Rather than just symbols from a fixed finite alphabet, a depth-stack stores *stack frames*, which consist of a symbol and a depth. By replacing the registers in a depth-register automaton with a depth-stack we obtain a *depth-stack automaton*. The automaton can pop a frame, push a frame with the current depth, and compare the depth in the top frame with the current depth.

The advantage of the depth-stack over the classical stack is conciseness. In the ordinary stack-based simulation the height of the stack is tied to the depth of the tree. In the depth-stack model we keep track of the depth using the counter and the stack is only used to record when the state of the simulated DFA changes:

- whenever a label is about to trigger a state change, the current state and depth are pushed to the depth-stack;
- whenever the current depth drops to the value in the topmost frame, the frame is popped and the current state is reverted to the one in the frame.

For a child-free query with n selectors this results in $\mathcal{O}(n)$ memory usage, and the at most n frames on the stack correspond directly to the n registers from the stackless algorithm. For a query with child selectors the stack can grow up to the depth of the JSON tree, but for most real-life data this is rare: it requires documents where nodes with the same label are nested in itself, and the query asks for a child of a node with such a label (see query A1 in Section 5). In the implementation we represent the depth-stack using a special Rust structure `SmallVec`. This puts our depth-stack on the actual stack of the executing thread as long as it is relatively shallow (less than 128 elements, bounded by 512 bytes). In the rare cases when it grows larger than that, it is moved to the heap.

Apart from the depth-stack, the simulation algorithm stores the query automaton, whose size is negligible for practically useful queries. The overall memory footprint is linear in the depth of the document. Time complexity is linear in the size of the input data.

3.3 Skipping

While the algorithm described in Section 3.2 dutifully steps through every element of the document, the key insight from [21] is that one can *skip* fragments of the document that are known not to contain query matches. Of course, we cannot really jump over fragments of the stream, but we can fast-forward through them using SIMD processing. Below we discuss in the abstract the four types of skipping used in our query engine; their SIMD implementation is explained Section 4. In order to ensure correct semantics of wildcards, we skip less aggressively than JSONSKI, which assumes that wildcards match only array elements and so skips over objects if the next selector is a wildcard. Also, because automata for queries with wildcard and descendant are structurally non-trivial (as we have seen in Section 3.1), the criteria for skipping become more subtle. Finally, the last kind of skipping is entirely new: it is impossible for queries without descendant selectors.

Skipping leaves.

We call a state *internal* if it has no transitions to accepting states. When the simulated DFA is in an internal state, it makes no sense to visit leaves (i.e. atomic values) in the tree because the automaton needs to descend at least two levels to accept. That is, when going through the children of some node v we would like to skip all leaves ahead of us and jump straight to the next child that has some descendants. For example, consider the query `$.a..b.*..c.*` from Figure 2 and the document

```
{b:"Long string with no matches for sure", c:[1,2,...,1000], a:{...}, ...
}
```

After reading the initial `{` we are in the initial state of the DFA in Figure 2 (bottom), which is an internal state. We would now like to skip the `b` child, because the DFA has nowhere to go below it. This can be done by jumping to the next opening character. However, we should also be mindful of closing characters because, if all remaining siblings are leaves, the next opening character will be already outside of the scope of the current subtree, which affects the simulation. Overall, when skipping leaves, we will be tracking structural characters `{`, `}`, `[`, `]`, fast-forwarding between their successive occurrences. In order to simulate the DFA, we will also need the label before each `{`, but we can get it by backtracking. When not skipping leaves, we will have to pay attention also to the remaining structural characters: commas and colons.

Skipping children.

When reading the label of the current node v moves the simulated DFA to the trash state, it makes no sense to visit the children of v . Instead, we would like to jump straight to the closing character marking the end of the whole subtree. Continuing the example above, after jumping to the first `[`, we see a label c before it which would lead us to the trash state (not depicted in Figure 2). We would like to skip all the children of this node (i.e., all the entries in the array) and jump straight to the closing `]`. Importantly, since we have already seen the opening character, we know what the matching closing character will be, so we need to track only two characters: either `{`, `}` or `[`, `]`.

Skipping siblings.

We call a state *unitary* if it has a single transition over a concrete label and its fallback transition leads to the trash state. (Such states correspond to non-wildcard selectors in the query before the first descendant selector.) Suppose that upon reading the label of the current node v the simulated DFA enters a unitary state with a transition over label ℓ . As soon as we discover a child v' of v with label ℓ , it makes no sense to visit the remaining siblings of v' , because labels do not repeat among siblings. That is, after processing the subtree rooted at v , just like in the previous case, we would like to jump straight to the closing character marking the end of the subtree rooted at the parent of v , and we know that this character will be `}`. Moving on with our example, after reading the character `}` marking the end of the subdocument `a:{...}`, no matter if we have found some query matches there or not, we know that there can be no more properties with label `a` in the root object, so we can jump directly to the closing brace.

Skipping to a label.

We call a state *waiting* if it has exactly one transition over a concrete label ℓ and its fallback transition is looping. Such a state corresponds to descendant selectors $..l$. When the simulated DFA is in such a state, it would make sense to jump directly to the next descendant of the current node that has label ℓ . This is hard in general, as it requires monitoring the depth and the label, but if the waiting state is the initial state of the automaton (this is the case for queries that begin with $$.l$) then the current node is the root and we can jump to the next occurrence of label ℓ .

3.4 Main algorithm

The pseudocode of the main algorithm, shown in Listing 1 in a Python-like syntax for brevity, combines the depth-stack based algorithm with skipping leaves, children, and siblings. It uses an iterator that abstracts away all access to the stream. The iterator allows advancing to the next relevant structural character with method `next()` and peeking it without advancing with method `peek()`. Both these methods yield items of an algebraic sum type with four variants: **Closing** for `{` and `[`, **Colon** for `:`, **Comma** for `,`, and **Opening** for `}` and `]`. Whenever an opening character is found, the method `get_label()` is used to get the label corresponding to the subdocument. The iterator does that simply by backtracking through whitespace characters (and possibly a colon) to the label and returning it. If instead it finds a comma or an opening character, it means that the encompassing element is an array and there is no label. In that case, an artificial label is returned, falling under the fallback transition of the simulated automaton.

By default, the iterator returns only opening and closing characters, which amounts to skipping leaves. The method `toggle()` checks if the automaton can accept in a single step from the current state in the current type of element (object or list). If so, it extends the set of structural characters to be iterated over: it adds commas if the current element is a list and colons if it is an object. (It would be possible to use commas in objects as well, but we found the current solution more efficient.) We make sure that only leaves are processed in cases **Colon** and **Comma** by checking if the next structural character is not an **Opening** character. If it is, the current node is not a leaf and it is handled in the main two cases. An additional corner case is the first item of an array, which will not be caught in the **Comma** case nor in the **Opening** case if it is leaf. We handle it using `try_match_first_item()` in the `Open('[')` case for the encompassing array, which reports an answer if: the array is non-empty, the next structural character is not opening, and the target of the fallback transition is accepting.

Skipping children and siblings is handled by calling the method `skip()` for the opening character that needs to be closed. Skipping to label is implemented outside of the above algorithm: if the query starts from a descendant selector $..l$, the engine finds the first occurrence of ℓ in the stream using the highly optimized SIMD-based `memmem` function from the `memchr` crate [14], and runs the described algorithm from there. When the whole subdocument associated with the first occurrence of ℓ is processed, the external loop identifies the next occurrence of ℓ and runs the algorithm again, and so on.

4 Vectorised classification

The core of our engine is a state-driven SIMD pipeline that quickly locates relevant characters and fast-forwards through irrelevant fragments of the stream, allowing to implement the iterator used in the main algorithm from Section 3.4. The pipeline consists of: a *structural*

```
1 state = automaton.init_state()
2 stack = init_stack()
3 while event = iterator.next():
4     match event:
5         case Opening(c):
6             label = iterator.get_label()
7             target = state.transition(label)
8             if target.is_rejecting():
9                 iterator.skip(c)
10                continue
11            if target != state:
12                stack.push(state, depth, c)
13                state = target
14            depth = depth + 1
15            if state.is_accepting():
16                report_match()
17            iterator.toggle(state, c)
18            if c == '[':
19                try_match_first_item()
20        case Closing(c):
21            depth = depth - 1
22            prev = stack.top()
23            if depth == prev.depth:
24                state = prev.state
25                stack.pop()
26                if state.is_unitary():
27                    iterator.skip('{')
28                    continue
29                iterator.toggle(state, prev.c)
30            else:
31                iterator.toggle(state, '{')
32        case Colon:
33            if iterator.peek() == Opening(_):
34                continue
35            label = iterator.get_label()
36            target = state.transition(label)
37            if target.is_accepting():
38                report_match()
39            if state.is_unitary():
40                iterator.skip('{')
41        case Comma:
42            if iterator.peek() != Opening(_):
43                report_match()
```

■ Listing 1 Main algorithm

XX:12 Supporting Descendants in SIMD-Accelerated JSONPath

■ **Table 1** JSON Structural Characters

Character	{	}	[]	:	,
UTF	0x7b	0x7d	0x5b	0x5d	0x3a	0x2c

■ **Table 2** Classification on Intel Skylake (S) and Ice Lake (IL). Cycles per iteration (cpi) computed with uops.info [1].

Values	Naive										
	1	2	3	4	5	6	7	8	16	32	
cpi (S)	1.33	2.00	2.25	2.82	3.50	4.20	5.20	5.92	11.40	21.89	
cpi (IL)	1.10	1.86	2.11	2.81	3.48	4.16	5.23	5.92	11.35	21.60	

	Non-overlapping	Few groups	General
cpi (S)	2.25	2.50	4.00
cpi (IL)	2.03	2.31	3.67

classifier that recognizes structural characters shown in Table 1 (with the possibility to toggle commas and colons) and fast-forwards through whitespaces, labels, and atomic values; a *depth classifier* used to fast-forward through entire subdocuments; and a *quote classifier* that detects characters placed between quotes and handles escaping.

The structural classifier is obtained as a special case of a general method for *raw classification* (Section 4.1), which can be of use for fast parsing of other document formats. Along with the quote classifier (Section 4.2), it underlies the main functionality of the iterator (Section 4.3), whereas the depth classifier powers skipping (Section 4.4). Together, the classifiers build up a flexible and extensible multi-classifier pipeline (Section 4.5).

4.1 Raw classification

Our algorithm has to solve a special case of a more general *binary classification problem*, asking to classify an input vector of n bytes into 2 buckets.

► **Problem 1** (Binary Classification). *Fix a classification set*

$$A \subseteq \{0x00, 0x01, \dots, 0xff\}.$$

Given a vector $v = [v_0, v_1, \dots, v_{n-1}]$ of n bytes, compute the n -bit classification mask whose i -th bit is set iff $v_i \in A$.

In the engine, we run the classification algorithm in a loop, processing consecutive blocks of the input. We seek to minimize the *reciprocal throughput* of this loop, expressed in *cycles per iteration (cpi)*, measured as the average number of CPU cycles it takes to perform a single iteration of the loop's body, assuming it iterates forever. To estimate the reciprocal throughput on specific microarchitectures we use the uops.info Code Analyzer powered by nanoBench [1, 2, 3].

A naive classification algorithm compares (vectorially) elements of the input vector with each value from the classification set separately, and aggregates the answers using (vectorial) OR. This method is very fast when the classification set is small, but becomes costly as it grows (see Table 2). For 4 and more values one can do better.

Following [22], we use the `shuffle` instruction, which essentially allows one to compose functions represented as vectors, except that the values of the inner vector are trimmed to the

lower *nibbles* (4-bit parts). For 128-bit vectors \mathbf{a} and \mathbf{b} , the result of `shuffle_epi8(a, b)` is a 128-bit vector with $\mathbf{a}[\mathbf{b}[i] \& 0x0F]$ at position i , for all $0 \leq i < 16$ (assuming that the upper nibbles of \mathbf{b} are zeroed).¹ For instance, if $\mathbf{b} = [15, 14, \dots, 0]$, then `shuffle_epi8(a, b)` is the reverse of \mathbf{a} .

While the original purpose of `shuffle` was to reorder entries of vector \mathbf{a} based on vector \mathbf{b} , as in the example above, one can also think of it as a way to perform parallel lookup in a table \mathbf{a} at the positions indicated by vector \mathbf{b} . We can use it to quickly classify a vector \mathbf{b} of bytes into at most 16 buckets depending on their lower nibble. We can do the same for the upper nibbles (by simply shifting all bytes right by 4). While not every binary classification problem over bytes can be easily expressed as a combination of classifications of upper and lower nibbles, in what follows we describe two cases when it is possible. In both cases we construct two lookup tables such that the results of the lookups can be easily combined to provide the final classification. We also describe an original working solution for the general case, readily applicable wherever fast lexing is needed.

Fix a classification set A . We identify each byte with a pair of an upper and lower nibble; e.g. `0x3a` is identified with $\langle 3, \mathbf{a} \rangle$. We let $\text{Nib} := \{0, 1, \dots, \mathbf{e}, \mathbf{f}\}$.

► **Definition 1.** *The acceptance set $\text{low}(u)$ of a nibble u is the set of all nibbles l such that $\langle u, l \rangle$ is accepted, i.e.*

$$\text{low}(u) = \{l \in \text{Nib} \mid \langle u, l \rangle \in A\}.$$

► **Definition 2.** *An acceptance group is defined by a maximal set of upper nibbles that have the same acceptance sets. For a nibble u let $U_u = \{u' \in \text{Nib} \mid \text{low}(u') = \text{low}(u)\}$. The set G of all acceptance groups is defined as*

$$G = \{\langle U_u, \text{low}(u) \rangle \mid u \in \text{Nib}\}.$$

Note that $|G| \leq |\text{Nib}| = 16$.

► **Definition 3.** *Acceptance groups $\langle U_1, L_1 \rangle, \langle U_2, L_2 \rangle \in G$ are overlapping if $U_1 \neq U_2$ and $L_1 \cap L_2 \neq \emptyset$.*

For instance, if the classification set is

$$A = \{0\mathbf{x}\mathbf{a}\mathbf{1}, 0\mathbf{x}\mathbf{a}\mathbf{2}, 0\mathbf{x}\mathbf{b}\mathbf{1}, 0\mathbf{x}\mathbf{b}\mathbf{2}, 0\mathbf{x}\mathbf{c}\mathbf{2}\}$$

then

$$\text{low}(\mathbf{a}) = \text{low}(\mathbf{b}) = \{1, 2\}, \quad \text{low}(\mathbf{c}) = \{2\},$$

$$G = \{\langle \{\mathbf{a}, \mathbf{b}\}, \{1, 2\} \rangle, \langle \{\mathbf{c}\}, \{2\} \rangle\}.$$

The two groups in G are overlapping, as they share the element 2.

Depending on the properties of the set G of all acceptance groups, we distinguish three cases of increasing complexity.

¹ Zeroing the upper nibbles is important because the exact semantics of `shuffle` make bits there meaningful – if the most significant bit is set, the result is mapped to zero. For all our purposes we want the upper nibbles to be zeroed.

XX:14 Supporting Descendants in SIMD-Accelerated JSONPath

Non-overlapping groups.

If G contains no overlapping groups, the simplest solution is to map lower and upper nibbles from a single group to a unique value, compare the results with `cmpeq`, and convert the resulting vector into a bit mask with `movemask`, as follows. Take an arbitrary enumeration $\langle U_1, L_1 \rangle, \dots, \langle U_{|G|}, L_{|G|} \rangle$ of groups in G . Then construct an upper table as a vector *utab* such that $utab[x] = i$ if $x \in U_i$, and $utab[x] = 254$ if there is no such i . Analogously construct a lower table *ltab* as $ltab[x] = i$ if $x \in L_i$, and $ltab[x] = 255$ if there is no such i . Naturally, $|G| < 254$. The classification mask is then obtained as follows:

```
let usrc = shiftright_epi8(src, 4);
let llookup = shuffle_epi8(ltab, src);
let ulookup = shuffle_epi8(utab, usrc);
let b = cmpeq_epi8(llookup, ulookup);
let mask = movemask_epi8(b);
```

An example is shown in Figure 3. Note that x86 SIMD does not have the `shiftright_epi8` instruction, but it can be simulated with two instructions: first, 16-bit right shift by 4, then zero the upper nibbles with a precomputed mask. Assume that the mask is in register `ymm1`, the block of input is in `ymm3`, while the lower and upper lookup tables are in `ymm0` and `ymm2`, respectively. The compiled body of the classification loop is then:

```
vpsrlw    ymm4, ymm3, 4
vpand    ymm4, ymm4, ymm1
vpshufb  ymm4, ymm2, ymm4
vpshufb  ymm3, ymm0, ymm3
vpcmpeqb ymm3, ymm3, ymm4
vpmovmskb edx, ymm3
```

For Intel Skylake the predicted reciprocal throughput is 2.25 cpi, while for the modern Ice Lake microarchitecture we get 2.03 cpi [6]. This is on par or better than naive already for 3 values (see Table 2).

As it happens, the non-overlapping case is sufficient for our JSON structural classifier. JSON structural characters are shown in Table 1. The groups are

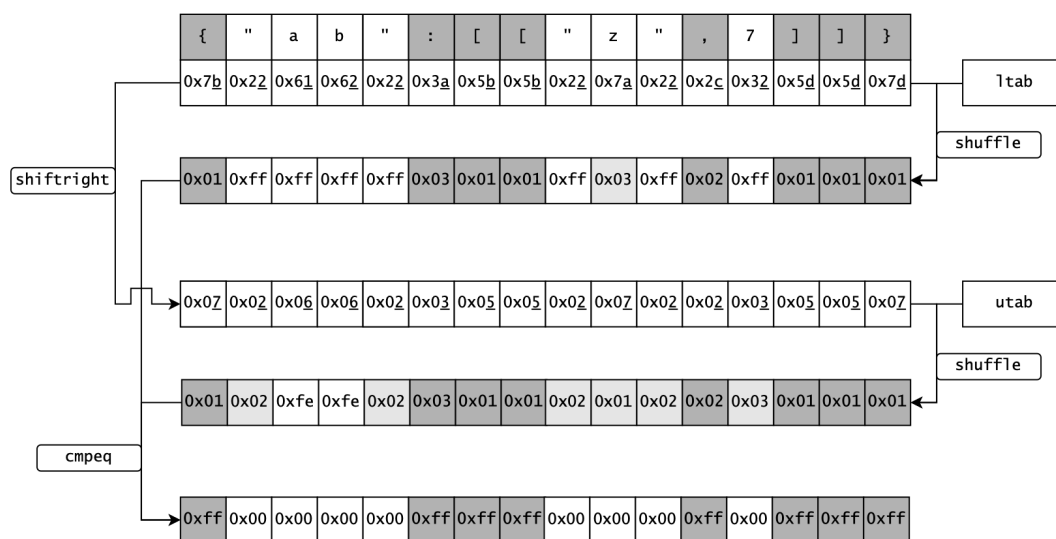
$$\{\langle\{5, 7\}, \{b, d\}\rangle, \langle\{2\}, \{c\}\rangle, \langle\{3\}, \{a\}\rangle\}$$

and they are non-overlapping. Consequently, the lower and upper lookup tables used in our structural classifier, and in the example in Figure 3, are

```
utab = [0xfe, 0xfe, 0x02, 0x03, 0xfe, 0x01, 0xfe, 0x01,
        0xfe, 0xfe, 0xfe, 0xfe, 0xfe, 0xfe, 0xfe, 0xfe],
ltab = [0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff, 0xff,
        0xff, 0xff, 0x03, 0x01, 0x02, 0x01, 0xff, 0xff].
```

Compared with the naive method for 6 values, we save 2 cpi (see Table 2). `Simdjson` [22] uses the same code with the same lookup tables, but does not abstract the general method as we do here.

An additional challenge in `rsonpath` is that skipping leaves requires toggling commas and colons on request. This has an elegant solution. We can simultaneously disable all symbols with upper nibble b by zeroing *utab* at position b . Because *ltab* contains only non-zero



■ **Figure 3** JSON document classified using the structural classifier's `ltab` and `utab` lookup tables.

values, `cmpeq` will return 0 for all symbols with upper nibble `b`. As commas and colons do not share their upper nibble with any other accepted symbol, we can toggle them independently by XOR-ing `utab` with the following bitmasks:

$$\text{toggle_comma} = [0x00, 0x00, 0x02, 0x00, 0x00, \dots, 0x00],$$

$$\text{toggle_colon} = [0x00, 0x00, 0x00, 0x03, 0x00, \dots, 0x00].$$

When both commas and colons are disabled, our classifier has only 4 symbols to accept, but it is still faster than the naive method (see Table 2).

4.1.0.1 Few groups.

Another case that can be efficiently solved is when $|G| \leq 8$. The idea is to assign a unique index from 0 to 7 to each group and then let the upper nibble lookup zero the bit at the index corresponding to the unique group containing the nibble, and the lower nibble lookup set bits at all indices corresponding to groups whose acceptance set contains the nibble. That is, we take an arbitrary enumeration $\langle U_1, L_1 \rangle, \dots, \langle U_{|G|}, L_{|G|} \rangle$ of groups in G . We construct the upper table `utab` such that $\text{utab}[x] = 2^8 - 1 - 2^{i-1}$ if $x \in U_i$, and $\text{utab}[x] = 0$ if no such i exists. The lower table is defined as

$$\text{ltab}[x] = 2^{i_1-1} + 2^{i_2-1} + \dots + 2^{i_c-1},$$

where $x \in L_{i_1}, x \in L_{i_2}, \dots, x \in L_{i_c}$. Then, for every byte $b = \langle u, l \rangle$, $f(b) = 1$ if and only if the bitwise OR of `utab[u]` and `ltab[l]` is $2^8 - 1$. The classification vector b is obtained with just one more operation than in the non-overlapping case:

```
let usrc = shiftright_epi8(src, 4);
let llookup = shuffle_epi8(ltab, src);
let ulookup = shuffle_epi8(utab, usrc);
let lookup = or(llookup, ulookup);
let b = cmpeq_epi8(lookup, ALL_ONES);
let mask = movemask_epi8(b);
```


XX:16 Supporting Descendants in SIMD-Accelerated JSONPath

This is faster than the naive method for 4 or more symbols. For Intel Skylake the predicted reciprocal throughput is 2.50 cpi, while for the modern Ice Lake microarchitecture we get 2.31 cpi [4], which is slightly worse than the non-overlapping case.

4.1.0.2 General case.

We have not found an elegant solution to the general case for $8 < |G| \leq 16$. A working approach is to apply the algorithm for the small case twice. First partition the set G into G_1, G_2 such that $|G_1| \leq 8$ and $|G_2| \leq 8$. Then classify the bytes according to G_1 and G_2 , possibly with local parallelism. In the end, take the OR of both classifications to obtain a classification for G .

```
let usrc = shiftright_epi8(src, 4);
let llookup1 = shuffle_epi8(ltab1, src);
let ulookup1 = shuffle_epi8(utab1, usrc);
let lookup1 = or(llookup1, ulookup1);
let llookup2 = shuffle_epi8(ltab2, src);
let ulookup2 = shuffle_epi8(utab2, usrc);
let lookup2 = or(llookup2, ulookup2);
let lookup = or(lookup1, lookup2);
let b = cmpeq_epi8(lookup, ALL_ONES);
let mask = movemask_epi8(b);
```

For Intel Skylake the reciprocal throughput is 4.00 cpi, and for Ice Lake it is 3.67 cpi [5]. This is significantly worse than we could do in the two simpler cases, and there might be better solutions for other special cases, but it is also significantly faster than the naive method, because the general method is applied only when $|G| > 8$, and the number of accepted symbols is at least $|G| - 1$ (see Table 2).

4.2 Recognising quoted sequences

The next step is to ignore characters recognised as structural by the lookup, but located inside JSON strings, that is, between unescaped double quotes. A double quote is escaped if and only if it is preceded by a sequence of backslashes of odd length: "x\" contains a single escaped double quote, while "x\\" contains a single escaped backslash and none of the double quotes are escaped.

We use the same solution as `simdjson` [22]. First, we mark all backslash and quote characters with a `cmpeq`. Then we mark starts of backslash sequences and use *add-carry propagation* to find their ends. Based on this we mark unescaped double quote characters. Finally, we compute the prefix-xor of the resulting bitmask using *carry-less multiplication* (the `cmul` instruction) by a vector with all bits lit, and the bits set in the result indicate the positions between quotes.

We deal with block boundaries by maintaining a state storing two bits of information: whether the previous block's last character was an unescaped backslash and whether the last block ended while still within quotes.

4.3 Structural iterator

All the classification we have performed up to this point was on a single block of data. To feed information to the main algorithm we need an abstraction on top of a block classifier

that will give us a classifier for the entire input stream. As explained in Section 3.4, the abstraction is provided by an iterator.

The iterator operates on classified blocks of structural characters, which contain a bitmask with all structural characters marked as described in the sections above, along with a reference to the original input block. The iterator begins by classifying the first block of input. Then, when asked for the next structural character, it examines the current classified block and its bitmask. If it is all-zeroes, then there are no structural characters in the current block and we need to classify the next one. If it is not, then we extract the position of the first set bit by calculating trailing zeroes of the mask and check the original input block for the character located at that position. Additionally, we modify the stored bitmask by zeroing the set bit we just processed.

While the SIMD pipeline calculating structural bitmasks is completely branchless, the outer loop of the iterator contains branching when we compare the mask with zero, and then when we branch on the input character to return a proper variant of the algebraic sum type.

The method `toggle()` enables and disables colons and commas by XOR-ing the upper lookup tables with precomputed bit masks, as explained in Section 4.1. When a character is enabled, the current block has to be reclassified. When a character is disabled, its occurrences in the remainder of the block are simply stepped over by the outer loop of the iterator.

4.4 Depth classifier

Skipping children and siblings is implemented by iterating over closing characters (of the suitable kind) and checking the relative depth at each time, until it drops to 0 for the first time. We use a specialised *depth classifier* here, because it has fewer characters to classify but additionally needs to maintain the depth.

We need to keep track of two characters only: either '[' and ']', or '{' and '}'. Hence, instead of using the complex method of the structural classifier (5 SIMD instructions), we can mark all opening and closing characters in the block using two `cmpeq` instructions. This results in two separate bitmasks, rather than a common bitmask for all relevant characters. As it happens, this suits us even better, because we can now find the next closing character by locating the first set bit in the respective bitmask and we can update the relative depth by counting set bits in the relevant chunk of the other bitmask.

We also use an additional heuristic that goes beyond JSONSKI's methods. If the number of closing characters in the whole block is strictly smaller than the current relative depth, nowhere in the block the relative depth can drop to 0. We then jump directly to the end of the block instead of iterating through all closing characters and checking the relative depth for each of them. We update the relative depth based on the total numbers of opening and closing symbols in the block. This makes the algorithm much more efficient, especially on real-life JSON documents

4.5 Multi-classifier pipeline

As skipping children and siblings is performed on demand, we need to be able to trigger depth classification at any moment, pausing the full structural classifier for the duration of the fast-forward. After we are done skipping, we need to resume structural classification. To facilitate this we propose a robust multi-classifier pipeline ready for extension with different classifiers in the future.

Our pipeline consists of the core classifier, the *quote classifier*, which is responsible for creating the bitmasks marking characters that are within quotes and should be ignored.

XX:18 Supporting Descendants in SIMD-Accelerated JSONPath

This classification is always required to maintain correctness. On top of that, we can run either the structural or the depth classifier. To allow quick switching between them, they both expose `stop` and `resume` methods. The `stop` returns an object that encapsulates the state of the underlying quote classifier – all of its internal structures and the last classified block along with the index to which classification was performed. The `resume` counterpart can take such an object, restore the quote classifier, and begin the top-level depth or structural classification. By leveraging Rust’s ownership system we can easily pass the internal quote classifier structures between components without copying (which would be slow) or complicated memory management (which would be error-prone and hard to modify): Rust statically ensures that only one classifier has write access to the shared structures at a time.

While our implementation has only two classifiers in the pipeline, one can easily add further specialised classifiers by simply providing them with appropriate `stop` and `resume` methods. One can envision a progression of more and more refined classifiers. Our depth classifier allows to keep track of the depth to stay within an element. This can be extended to a classifier that allows to fast-forward to the next occurrence of a label within an object. Such a classifier could be leveraged to speed up the execution of nested descendant selectors. Finally, one could additionally demand staying at the same depth, which would be useful in processing child states.

We believe that a flexible pipeline is an important engineering milestone that can enable programmers to construct SIMD solutions that are naturally composable and easily modifiable. This increases the level of abstraction on which we operate when designing SIMD accelerations, while not sacrificing low-level control required to produce very performant code.

5 Experiments

We performed a series of experiments comparing our engine with existing solutions. Our goals were three-fold: estimate the overall overhead incurred by supporting descendants and idiomatic wildcards; showcase the benefits from working directly with descendant queries rather than semantically equivalent descendant-free queries; and identify improvement opportunities. The benchmark code and all datasets are available online [15].

5.1 Competitors

For the baseline we chose `JsonSurfer` [32], which works in the streaming model and supports full JSONPath, but does not apply any SIMD optimizations. We exclude the popular `jq`, which supports a much more expressive variant of JSONPath, because it is much slower than `JsonSurfer` (one order of magnitude slower on the query A1 from Table 6). Our main point of reference is `JSONSki` [21], to the best of our knowledge the only JSONPath engine using SIMD optimization. We do not include `simdjson` [22] in the comparison, because it does not support JSONPath queries directly; instead, queries have to be implemented by hand using an on-demand API. From previous studies we know that `JSONSki` performs better [21].

`JsonSurfer` is implemented in Java. We run it without any modifications via JNI [29] using Java 11 (OpenJDK). The overhead incurred by Rust-Java interop amounts to nanoseconds per call, which is negligible considering that the query execution times are in the range of milliseconds. `JSONSki` is implemented in C++. In order to integrate it with our Rust-based benchmark harness, we introduced a few tweaks. First, we removed gathering results with `std::vector` in favour of a simple match counter. This slightly increases the performance of `JSONSki` compared to the original [21]. Second, we fixed a few memory-management

■ **Table 3** Datasets used in experiments

Name	Size [MB]	Depth	Verbosity
* AST (A)	25.6	102	14.3
BestBuy (B)	1044.6	8	24.5
* Crossref (C)	550.5	9	27.0
GoogleMap (G)	1136.1	10	36.9
NSPL (N)	1210.2	10	13.8
Twitter (T)	842.5	12	29.0
* Twitter small (Ts)	0.7	11	50.6
Walmart (Wa)	995.4	5	96.9
Wikimedia (Wi)	1099.0	13	18.7

issues that manifest when the engine is run many subsequent times within the same process. These were identified by running the executable under Valgrind and are minor changes in constructor and destructor code, with no effect on the actual performance. To compile JSONSki we used gcc version 10.2.1 20210110 (Debian 10.2.1-6) with flags

```
-O3 -std=c++11 -mavx2 -mpclmul -static-libstdc++
```

For rsonpath we used rustc 1.67.1 (d5a82bbd2 2023-02-07) with flags

```
-C target-cpu=native -C opt-level=3 -C lto=thin
```

5.2 Datasets

In our experiments we use the 6 datasets from the JSONSki benchmark as well as 3 additional datasets representing characteristic usecases. The first one is the `twitter.json` file extracted from simdjson’s quick-start tutorial in the project’s repository [31]. It is a typical file obtained by querying an API, small but irregular. The second dataset is an arbitrarily chosen fragment of the datadump [11] of the Crossref service [10] collecting metadata of several hundred millions of scientific publications. The document is highly regular: it contains collections of sub-documents of very similar shape. The third dataset is a deep and highly irregular file: it is the abstract syntax tree of a single large C file (23K lines of code) obtained with `clang`. It falls within the code-as-data application scenario mentioned in Section 1.2. Detailed characteristics of the datasets are shown in Table 3; the asterisk indicates added datasets. By *verbosity* we mean the ratio of the size of the document to the number of nodes in the underlying tree: the lower the verbosity, the harder it is to achieve high throughput. All datasets are available on Zenodo [27].

5.3 Setup

We use Rust Criterion [19] as the benchmarking harness. A benchmark is executed as follows. First, a warm-up is performed, ensuring that the low-level caches are filled for the actual measurements. Measurements are performed on a number of samples, each of which consists of many iterations of the benchmarked routine. The mean execution time of all iterations is taken as a single sample. Finally, collected samples are analysed to find the statistical distribution, outliers are detected, and a mean time is reported. We calculate throughput based on that.

All experiments were run in a stable environment, on machines hosted by Grid’5000, a federated testbed setup providing a high diversity of architectures [16]. We show the results

■ Table 4 JSONSki queries

ID	Query	Matches
B1	<code>\$.products.*.categoryPath.*.id</code>	697440
B2	<code>\$.products.*.videoChapters.*.chapter</code>	8857
B3	<code>\$.products..videoChapters</code>	769
G1	<code>\$.*.routes.*.legs.*.steps.*.distance.text</code>	1716752
G2	<code>\$.*.available_travel_modes</code>	90
N1	<code>\$.meta.view.columns.*.name</code>	44
N2	<code>\$.data.*.*.*</code>	8774410
T1	<code>\$.*.entities.urls.*.url</code>	88881
T2	<code>\$.*.text</code>	150135
W1	<code>\$.items.*.bestMarketplacePrice.price</code>	15892
W2	<code>\$.items.*.name</code>	272499
Wi	<code>\$.*.claims.P150.*.mainsnak.property</code>	15603

for the *Chetemi* nodes which offer *Intel Xeon E5-2630 v4* CPUs (Intel Broadwell) with 256 GiB of (8x16GiB) of 24000 MHz RAM. Plots for further architectures (Intel Skylake, Intel Cascade-Lake, and AMD Zen 2) can be found in Appendix C. Results in tabular form can be found in Appendix D, and raw results can be found on Zenodo [26].

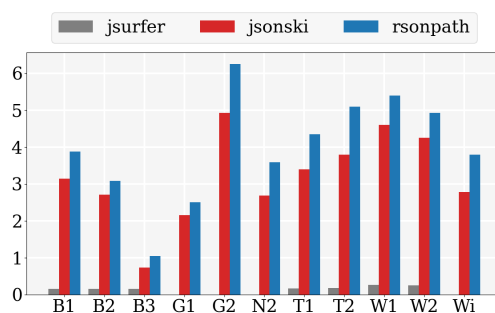
5.4 Experiment I: Overhead

Here we aimed to estimate how the added support for descendants and idiomatic wildcards influences the performance on descendant-free queries. We compared the performance of *jsurfer*, *JSONSki*, and *rsonpath* on the *JSONSki* benchmark (Table 4). We modified queries B1 and N2 by replacing slices (which we do not support) with wildcards, adding more equivalent work for all implementations, and dropped the other Wikipedia query because it coincides with Wi when slice is replaced with wildcard. We also added one query, indicated with *.

The results (Figure 4) show that the multi-classifier pipeline not only incurs no overhead, but actually boosts the performance by 10-20%. With no SIMD acceleration to muster, *jsurfer* consistently performs an order of magnitude worse (missing bar indicates that *jsurfer* was unable to load the dataset into memory due to the limitations of the JVM).

We added query B3 to explore the effect of *JSONSki*'s assumption that wildcard only traverses array elements. B3 was obtained from B2 by removing the last two selectors. It returns much fewer answers but, because of the missing wildcard selector, *JSONSki* cannot conclude that `.videoChapters` must match an array and needs to check objects and atomic values as well. In consequence, *JSONSki* is around three times faster on B2 than on B3. While *JSONSki* exploits this assumption for each wildcard selector, it does not give it an edge one might expect. Interestingly, *rsonpath* is also slower on B3, because it cannot apply leaf skipping any more.

On query N1 (not plotted), after processing the `meta` element contained in a small initial fragment of the document, *JSONSki* exits early instead of skipping the remaining children of the root in the usual way by fast-forwarding to the corresponding closing character. Because the overall size of the document is large, this gives terabytes per second of throughput. We have not implemented this optimization in *rsonpath*, but it is easy to add. When early exit is factored out by computing throughput with respect to the processed part of the document, we see that the raw performance of *rsonpath* and *JSONSki* on N1 follows the trend visible in Figure 4.



■ **Figure 4** Throughput [GB/s] for descendant-free queries.

■ **Table 5** JSONSki queries rewritten with descendants

ID	Query	
B1 ^r	<code>\$.categoryPath..id</code>	697440
B2 ^r	<code>\$.videoChapters..chapter</code>	8857
*B3 ^r	<code>\$.videoChapters</code>	769
G2 ^r	<code>\$.available_travel_modes</code>	90
W1 ^r	<code>\$.bestMarketplacePrice.price</code>	15892
W2 ^r	<code>\$.name</code>	272499
Wi ^r	<code>\$.P150..mainsnak.property</code>	15603

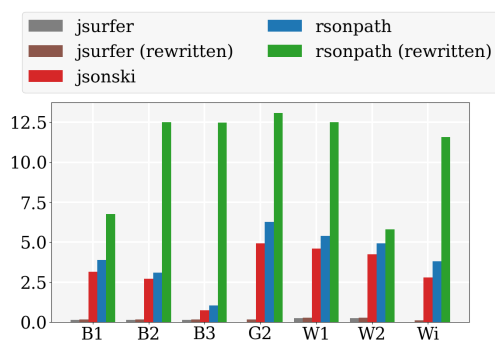
5.5 Experiment II: Benefits from descendants

Here we explore how the performance changes when queries are reformulated using descendants (without altering the semantics). Whenever such rewritings exist, they are simpler than the originals (Table 5). The results, shown in Figure 5, confirm that descendants not only simplify formulating queries, but also significantly boost performance, reaching throughput above 10GB/s for most tested queries. This is largely due to the `memmem`-based skipping to a label: `rsonpath` essentially jumps from one shallowest occurrence of the first label in the query to another. The boost is particularly impressive for the added query B3, whose original version is particularly hard for both JSONSki and `rsonpath`. The two queries with throughput below 10GB/s have an order of magnitude more matches than others, and the sheer cost of label comparison and the branching code to accumulate the results cannot be compensated by SIMD. `Jsurfer` performs orders of magnitude worse, essentially the same on the original and rewritten queries.

5.6 Experiment III: Limitations and Opportunities

In this experiment we aim to learn more about the behaviour of `rsonpath` in diverse scenarios. We use the new datasets and queries shown in Table 6; the results are shown in Figure 6.

Queries A1, A2, and C1 cannot be expressed without descendant because of relevant labels appearing at different levels. We plot them in green to put them visually in the same group with rewritings, as they are all descendant queries. Our target is the `memmem`-based skipping to a label. A1 illustrates how fast one can go for highly selective queries. On A2 the performance is very low (only x4 speedup over `JsonSurfer`). This is because `inner` labels are nested and the query is highly ambiguous, which makes the depth-stack grow deep, as described in Section 3.2. No good implementation for such queries is known, even theoretically. The purpose of C1 is to test the `memmem` acceleration in a very low selectivity



■ **Figure 5** Throughput [GB/s] for descendant-free queries and their rewritings using descendants.

■ **Table 6** Additional queries of interest

ID	Query	Matches
A1	<code>\$.decl.name</code>	35
A2	<code>\$.inner.inner.type.qualType</code>	78129
C1	<code>\$.DOI</code>	1073589
C2	<code>\$.items.*.author.*.affiliation.*.name</code>	64495
C2 ^r	<code>\$.author..affiliation..name</code>	64495
C3	<code>\$.items.*.editor.*.affiliation.*.name</code>	39
C3 ^r	<code>\$.editor..affiliation..name</code>	39
Ts	<code>\$.search_metadata.count</code>	1
Ts ^r	<code>\$.count</code>	1
Ts ^p	<code>\$.search_metadata.count</code>	1

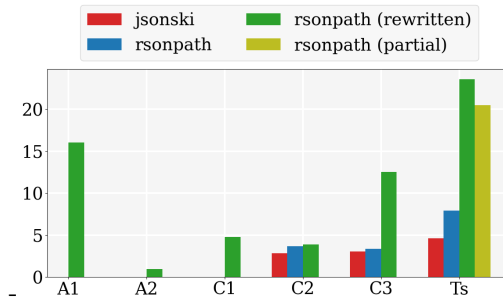
situation. As expected, performance is mediocre because repeated calls to `memmem` result in numerous short fast-forwards, rather a few long ones. Yet, low selectivity (C1) does not slow `rsonpath` down as much as nested labels (A2).

Queries C2 and C3 illustrate how the gain from rewriting with descendants may depend on the structure of the document. The queries have the same structure, yet the gain for C2 is negligible. Indeed, the rewriting of C2 is hard for `rsonpath`. `JSONSki` only examines `author` nodes at depth 3 while `rsonpath` goes through all `author` nodes (12x more), and each time looks through the whole subdocument searching for affiliation names (and finds no additional ones). We would improve our performance on this query if we could fast-forward to a given label within an element (see Section 4.5): this would allow us to quickly discard authors without affiliations. In contrast, the rewriting of C3 is easy for `rsonpath`, because all editors are at the same level, so there are no additional subdocuments to process.

The query Ts^p is a partial rewriting of Ts: it keeps the label `search_metadata` even though it is semantically redundant. Thus, the query Ts, its rewriting (Ts^r), and partial rewriting (Ts^p), all return the same results, but specify the access path with varying precision. They show that in our approach the less specified the path, the better. On Ts `rsonpath` and `JSONSki` display similar performance. If we knew how to find the next label at the same depth (a sibling) we could do better, possibly closing the gap to the rewritten variants.

5.7 Experiment IV: Scalability

To test the scalability of `rsonpath` on various architectures (see Appendix C), we ran the query `$.affiliation..name` on fragments of the Crossref dataset of increasing size (Table 7).



■ **Figure 6** Throughput [GB/s] for additional queries and their rewritings.

As expected, we observed no significant variation.

■ **Table 7** Throughput [GB/s] on the Crossref dataset

Size [GB]	0.3	0.5	1.1	2
Broadwell	7.9	7.8	7.9	7.8
Skylake	8.2	8.2	8.1	8.1
Cascade-Lake	8.3	8.1	8.1	8.1
Zen 2	9.8	10.0	10.1	9.8

6 Looking ahead

We have shown that it is possible to support descendants and idiomatic wildcards in SIMD-accelerated JSONPath processing without paying a penalty and that natural queries can be sped up significantly by using descendants. Supporting array indexing is compatible with our approach and we plan to implement it in the near future. Filters, on the other hand, are highly problematic because multiple potential matches must be stored until all filters are evaluated and it transpires which matches should be returned. This requires unbounded working memory and departs from the streaming paradigm. Another challenge is compositionality; that is, processing queries in succession, with the output of one query fed directly to another.

The key lesson from our work is that it pays off to split complex SIMD optimization tasks into a simple lower layer with high potential for SIMD acceleration, and a stateful sequential upper layer. This facilitates optimizations that would otherwise be impossible to design. One can apply it to any stateful SIMD processing task: regexp matching, DNA processing, XML parsing and querying, etc. Moreover, the separation into a hardware-dependent branchless layer and a higher-level sequential layer is relevant not only for SIMD but also for FPGAs and ASICs, except that acceleration options are richer. Ideally, one would want to use the acceleration to compute matching parentheses in the stream, but even theoretically efficient algorithms for this task are not known.

A Artifact Appendix

A.1 Abstract

This artifact includes the source code of `rsonpath`, the source code of our benchmarking workflow, and instructions on how to build, test, and run the experiments. The instructions

XX:24 Supporting Descendants in SIMD-Accelerated JSONPath

provided here are for reproducing the data behind the three core figures in our paper, 4, 5, and 6 (also given in Table 9 in Appendix D). All datasets are separately archived on Zenodo [27], and the artifact includes code that automatically downloads them and verifies their integrity. The base requirement to run the benches is a standard Rust toolchain with `cargo`. The artifact is also published in the newest version in the `rsonpath` GitHub repository, but for reproducibility we *recommend* using the Zenodo archived version [15].

A.2 Artifact check-list (meta-information)

- **Algorithm:** new JSONPath evaluation algorithm.
- **Program:** custom benchmarks.
- **Compilation:** Standard rustc compiler, lld linker, Java Runtime.
- **Data set:** all published on Zenodo [27], automatically downloaded by benchmark script, approx. size 6GB
- **Run-time environment:** Linux OS; cargo toolchain version at least 1.70.1, glibc in version at least 2.32, including lld; Java runtime in version at least 1.8 for jsurfer benchmarks.
- **Hardware:** x86_64 CPU with AVX2 extensions.
- **Execution:** 1–2 hours.
- **Metrics:** execution time and throughput.
- **Output:** console output with time, throughput, and statistical analysis; CSV file with results; PNG files with charts similar to 4–6.
- **Experiments:** ran by an automated workflow using Criterion.
- **Disk space required (approximately):** 6GB.
- **Time needed to prepare workflow (approximately):** 15 minutes.
- **Time needed to complete experiments (approximately):** 1–2 hours.
- **Publicly available:** <https://github.com/V0ldek/rsonpath>.
- **Code licenses (if publicly available):** MIT.
- **Data licenses (if publicly available):** CC BY v4.
- **Workflow framework used:** <https://github.com/bheisler/criterion.rs>.
- **Archived (provide DOI):** <https://doi.org/10.5281/zenodo.8400854>.

A.3 Description

A.3.1 How to access

Recommended: download the archive from <https://doi.org/10.5281/zenodo.8400854>.
Alternative: get the newest version from GitHub: <https://github.com/V0ldek/rsonpath>.

A.3.2 Hardware dependencies

To achieve similar level of performance a CPU with the AVX2 extension set available is required. Our main benchmark is executed on Intel’s Broadwell architecture. AMD’s Zen2 is also sufficient. The latest GitHub version also accelerates on older platforms. Outside of x86_64 SIMD acceleration is unavailable – the application will run, but our results concern SIMD acceleration specifically.

A.3.3 Software dependencies

Stable Rust toolchain in version at least 1.65.0 (1.71.0 for the GitHub version) and the lld linker in version 2.32 are required. The benchmark script requires the `JAVA_HOME` environment variable to be set to the directory of the JDK (on debian `export JAVA_HOME=/usr/lib/jvm/java-11-openjdk-amd64/`) in version at least 1.8, and a C++ compiler capable of compiling JSONSki available. We recommend installing `just` (<https://github.com/casey/just>), which will allow one to run command aliases we defined. This is easiest to do with `cargo install just`.

A.3.4 Data sets

In the experiments described in the body of the paper we use the 6 datasets from the JSONSki benchmark (BestBuy, GoogleMap, NSPL, Twitter, Walmart, Wikimedia) and 3 additional datasets (AST, Crossref, Twitter). These 9 datasets are archived on Zenodo [27]. The archive also contains the original C file underlying the AST dataset, and an additional OpenFood dataset, used in experiments that did not make it to the body of the paper (see Appendix D). The benchmark script automatically looks into the benchmark's `data` directory and validates the integrity of present datasets against a SHA256 hash. They are automatically downloaded from Zenodo when invalid or not present. One can also download them manually and put them into the correct location (see Section A.4). The `src/framework/dataset.rs` contains all dataset definitions: paths, names, Zenodo links, and checksums.

A.4 Installation

Setup all required dependencies, most notably the Rust toolchain. Unpack the artifact, go to the `crates/rsonpath-benchmarks` and build the benchmarks with

```
just build    or    cargo build -release.
```

If downloading datasets manually, put them into the directory

```
crates/rsonpath-benchmarks/data
```

following the file hierarchy below:

```
ast
  ast.json
crossref
  crossref0.json
  crossref1.json
  crossref2.json
  crossref4.json
pison
  bestbuy_large_record.json
  google_map_large_record.json
  nspl_large_record.json
  twitter_large_record.json
  walmart_large_record.json
  wiki_large_record.json
twitter
  twitter.json
```

A.5 Experiment workflow

There are four separate benchmark group defined. To get all results included in our charts you need to run them all.

- `just bench paper_parity;`
- `just bench paper_rewrite;`
- `just bench paper_other;`
- `just bench paper_jsurfer.`

Criterion produces raw results into `target/criterion`. To produce the charts use our python script:

```
python3 -m chart ./target/criterion -o <OUT>
```

Three PNG charts will be written to the `<OUT>` directory:

- `main` corresponding to Figure 4,
- `rewrite` corresponding to Figure 5,
- `other` corresponding to Figure 6.

Raw data is put into `data.csv`. For details consult help:

```
python3 -m chart -h
```

A.6 Evaluation and expected results

During benchmark execution the Criterion framework will print times and throughputs of all experiments ran, as follows:

```
./data/ast/ast.json_decl_name/rsonpath
time:   [1.7530 ms 1.7580 ms 1.7637 ms]
thrpt:  [14.511 GB/s 14.558 GB/s 14.600 GB/s]
```

The three values are the confidence interval, with the middle value being the mean (for more information on Criterion's statistical properties see the user guide: <https://bheisler.github.io/criterion.rs/book>). These results can be postprocessed with the aforementioned python script.

The absolute numbers will heavily depend on the architecture chosen. The experiments from the body of the paper were ran on a Broadwell-microarchitecture CPU – an older architecture should be expected to perform worse, and a newer one should perform better. The key results that should be reproduced are:

- SIMD solutions are at least an order of magnitude faster than no-SIMD: `rsonpath` and `JSONSki` outclass `jsurfer`;
- `rsonpath` performs on par or better than `JSONSki`;
- `rsonpath` performs much better in most experiments with the rewritten query, especially on Ts; notable experiments where the difference is less pronounced are W2 and C2, as these were specifically chosen to show limitation of rewriting; for details consult Section 5.

A.7 Experiment customization

All our Criterion benchmarks are easily extensible.

A.7.1 Adding a benchmark

To add a custom set of benchmarks it is recommended to create a new file under the `benches` subdirectory. Assume we create `mybench.rs`. Then, a relevant entry at the end of the project's `Cargo.toml` needs to be added:

```
[[bench]]
name = "mybench"
harness = false
```

In `mybench.rs`, first include our workflow prelude:

```
use rsonpath_benchmarks::prelude::*;
```

Then define one or more benchmark functions. These need to be `pub`, take a single parameter of type `&mut Criterion`, and return a `Result<(), BenchmarkError>`. Inside, create a new `Benchset` with a unique ID and referencing the dataset to be benchmarked. After that add the targets to be benchmarked, keeping in mind each can be referenced at most once. For example, if we wished to compare how `rsonpath` fares against `jsonski` when enumerating all topmost symbols in the AST dataset, we could create a function like:

```
pub fn a4(c: &mut Criterion) -> Result<(), BenchmarkError> {
    let bench = Benchset::new("A3_inner", dataset::ast())?
        .add_target(BenchTarget::JsonSki("$.inner[*]"))?
        .add_target(BenchTarget::Rsonpath("$.inner[*]"))?
        .finish();
    bench.run(c);
    Ok(())
}
```

Finally, reference all functions you wish to run at the end of the file with the `benchsets` macro, providing a name of the benchset first and then each function:

```
benchsets!(my_custom_benches, a4);
```

You can then run such a bench with just `bench mybench`.

A.7.2 Adding a dataset

All datasets are defined in `src/framework/datasets.rs`. To include a custom one, you need to provide a path and the SHA 256 checksum of the file. It is easiest if your dataset is available via a URL. For example, to add a gzipped JSON file hosted on Zenodo, declare a new function in the file, such as

```
pub const fn my_dataset() -> Dataset {
    Dataset {
        name: "my_new_dataset",
        path: dataset_path!("custom/directory/my.json"),
        source: DatasetSource::UrlArchive(DatasetArchive {
            url: "https://zenodo.org/record/<ID>/files/my.json.gz",
            checksum: hex!("<CHECKSUM OF my.json.gz>"),
        }),
        checksum: hex!("<CHECKSUM OF UNCOMPRESSED my.json>"),
    }
}
```

This dataset can then be referenced in a `Benchset::new` invocation.

A.7.3 Benchmarking the binary directly

One may wish to benchmark the binary directly. To do so, build it from the root folder of the entire archive with

```
just build release    or    cargo build -release
```

The binary will be located at `target/release/rsonpath`. It takes two arguments, the JSONPath query and a path to the file. For example, the A1 benchmark can be executed directly as:

```
./target/release/rsonpath '\$.decl.name'  
./crates/rsonpath-benchmarks/data/ast/ast.json -rcount
```

This can be then benchmarked using external tools like `perf`. Note that measuring time of this invocation will provide worse results than our workflow, as it will also measure the file loading time, which we exclude in our experiments.

A.8 Notes

If the system is under memory pressure, the benchmarks will degrade. It is important to have plenty of free RAM, and in general it is recommended to have an otherwise “quiet” system – no other CPU or RAM intensive tasks running during benchmarks.

Memory pressure is particularly pronounced with `jsurfer`, where the JVM might run out of heap space and panic. The mitigations are, in the order of severeness:

- make sure the system has more memory;
- allow the JVM to allocate more heap space by increasing the value in the `-Xmx` JVM option (line 88 of `crates/rsonpath-benchmarks/src/rust_jsurfer.rs`);
- do not run `jsurfer` benches – they are not core to the paper.

A.9 Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-badging>
- <http://cTuning.org/ae/submission-20201122.html>
- <http://cTuning.org/ae/reviewing-20201122.html>

References

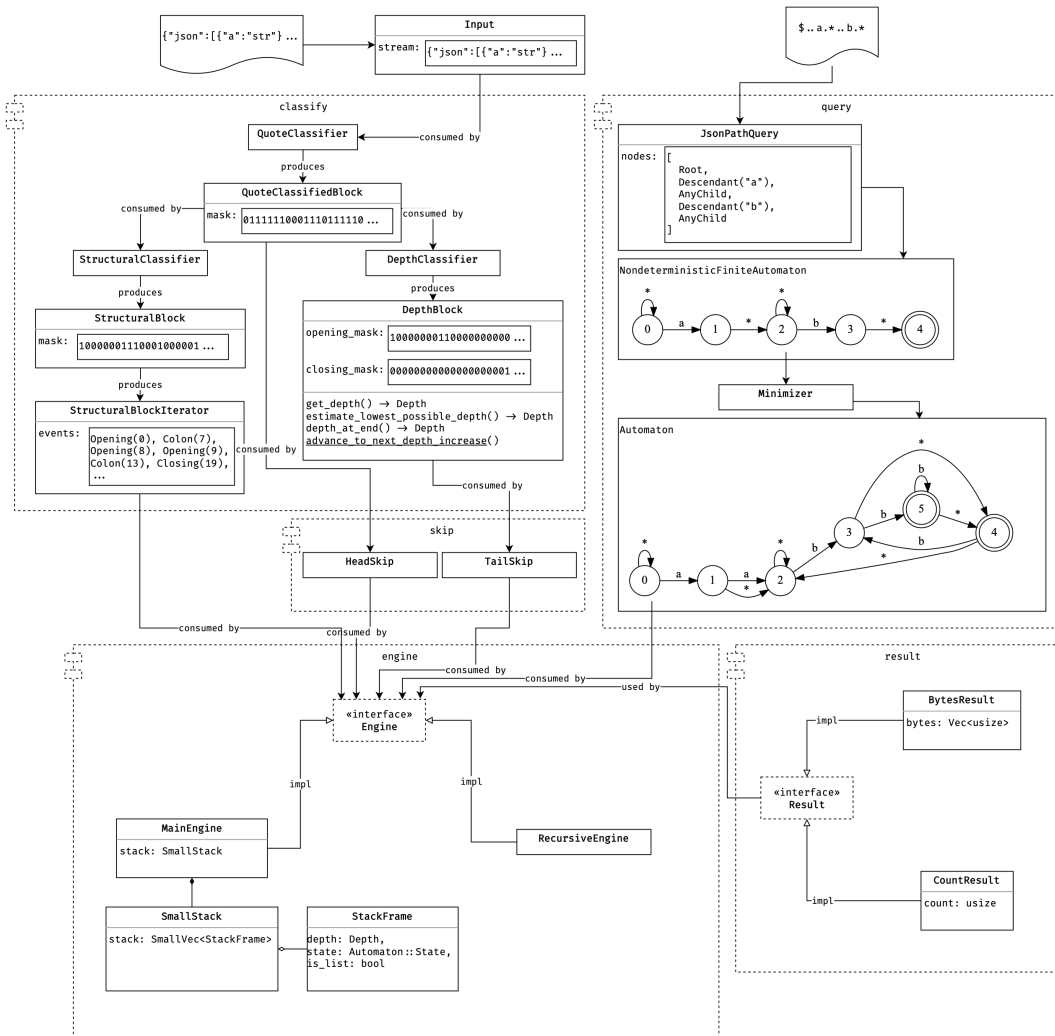
- 1 Andreas Abel and Jan Reineke. `uops.info`: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures. In *ASPLOS*, ASPLOS '19, pages 673–686, New York, NY, USA, 2019. ACM.
- 2 Andreas Abel and Jan Reineke. `nanobench`: A low-overhead tool for running microbenchmarks on x86 systems. In *2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 34–46, 2020.
- 3 Andreas Abel and Jan Reineke. `uiCA`: Accurate throughput prediction of basic blocks on recent Intel microarchitectures. In Lawrence Rauchwerger, Kirk Cameron, Dimitrios S. Nikolopoulos, and Dionisios Pnevmatikatos, editors, *ICS '22: 2022 International Conference on Supercomputing, Virtual Event, USA, June 27-30, 2022*, ICS '22, pages 1–12. ACM, June 2022.
- 4 Andreas Abel and Jan Reineke. The `uops.info` Code Analyzer: analysis of few groups classification loop, September 2023. <https://bit.ly/3sWaSGb>.

- 5 Andreas Abel and Jan Reineke. The uops.info Code Analyzer: analysis of general classification loop, September 2023. <https://tinyurl.com/2s3byyzf>.
- 6 Andreas Abel and Jan Reineke. The uops.info Code Analyzer: analysis of non-overlapping group classification loop, September 2023. <https://bit.ly/3RhtCKH>.
- 7 Corentin Barloy, Filip Murlak, and Charles Paperman. Stackless processing of streamed trees. In *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, PODS'21, page 109125, New York, NY, USA, 2021. Association for Computing Machinery.
- 8 Raphaël Bolze, Franck Cappello, Eddy Caron, Michel Daydé, Frédéric Desprez, Emmanuel Jeannot, Yvon Jégou, Stéphane Lanteri, Julien Leduc, Nouredine Melab, Guillaume Mornet, Raymond Namyst, Pascale Primet, Benjamin Quétier, Olivier Richard, El-Ghazali Talbi, and Iréa Touche. Grid'5000: A Large Scale And Highly Reconfigurable Experimental Grid Testbed. *International Journal of High Performance Computing Applications*, 20(4):481–494, 2006.
- 9 Christoph Burgmer et al. json-path-comparison, 2019. <https://github.com/cburgmer/json-path-comparison>.
- 10 Crossref. Crossref. <https://cd/www.crossref.org/>.
- 11 Crossref. Crossref datadump, 2022. <https://www.crossref.org/blog/2022-public-data-file-of-more-than-134-million-metadata-records-now-available/>.
- 12 Stephen Dolan. jq. <https://stedolan.github.io/jq/>.
- 13 James Finnie-Ansley, Paul Denny, Brett A. Becker, Andrew Luxton-Reilly, and James Prather. The robots are coming: Exploring the implications of OpenAI codex on introductory programming. In *Australasian Computing Education Conference*. ACM, February 2022.
- 14 Andrew Gallant. memchr, 2015. <https://crates.io/crates/memchr>.
- 15 Mateusz Gienieccko, Filip Murlak, and Charles Paperman. rsonpath, 2023. <https://zenodo.org/record/8400854>.
- 16 Groupement d'Intérêt Scientifique. Grid'5000 hardware, 2018.
- 17 Stefan Gössner. JSONPath, 2007. <https://goessner.net/articles/JsonPath/>.
- 18 Stefan Gössner, Glyn Normington, and Carsten Bormann. JSONPath: Query expressions for JSON. Internet-Draft draft-ietf-jsonpath-base-20, Internet Engineering Task Force, August 2023. Work in Progress.
- 19 Brook Heisler. criterion-rs, 2017. <https://crates.io/crates/criterion-rs>.
- 20 Lin Jiang, Junqiao Qiu, and Zhijia Zhao. Scalable structural index construction for JSON analytics. *Proc. VLDB Endow.*, 14(4):694–707, December 2020.
- 21 Lin Jiang and Zhijia Zhao. JSONSki: Streaming semi-structured data with bit-parallel fast-forwarding. In *ASPLOS*, pages 200–211. ACM, 2022.
- 22 Geoff Langdale and Daniel Lemire. Parsing gigabytes of JSON per second. *VLDB J.*, 28(6):941–960, 2019.
- 23 Yinan Li, Nikos R. Katsipoulakis, Badrish Chandramouli, Jonathan Goldstein, and Donald Kossmann. Mison. *Proc. VLDB Endow.*, 10(10):1118–1129, June 2017.
- 24 Tobias Mühlbauer, Wolf Rödiger, Robert Seilbeck, Angelika Reiser, Alfons Kemper, and Thomas Neumann. Instant loading for main memory databases. *Proc. VLDB Endow.*, 6(14):1702–1713, September 2013.
- 25 Shoumik Palkar, Firas Abuzaid, Peter D. Bailis, and Matei A. Zaharia. Filter before you parse: Faster analytics on raw data with sparser. *Proc. VLDB Endow.*, 11:1576–1589, 2018.
- 26 Charles Paperman. Rsonpath benchmark data files, Sep 2023. <https://doi.org/10.5281/zenodo.8386339>.
- 27 Charles Paperman, Mateusz Gienieccko, and Filip Murlak. Dataset for benchmarking rsonpath, Sep 2023. <https://doi.org/10.5281/zenodo.8386388>.
- 28 Antoine Pietri, Diomidis Spinellis, and Stefano Zacchiroli. The software heritage graph dataset: Large-scale analysis of public software development history. In *Proceedings of the 17th International Conference on Mining Software Repositories*, pages 1–5, 2020.
- 29 Prevoty, Inc. and jni-rs contributors. jni. <https://crates.io/crates/jni>.

XX:30 Supporting Descendants in SIMD-Accelerated JSONPath

- 30 simdjson. On-demand API. https://simdjson.org/api/0.6.0/md_doc_ondemand.html.
- 31 simdjson. simdjson. <https://github.com/simdjson/simdjson>.
- 32 Leo Wang. A streaming JSONPath processor in Java. <https://github.com/jsurfer/JsonSurfer/>.
- 33 K Zyp. Rfc 6901: Javascript object notation (json) pointer, 2013.

B Architecture overview



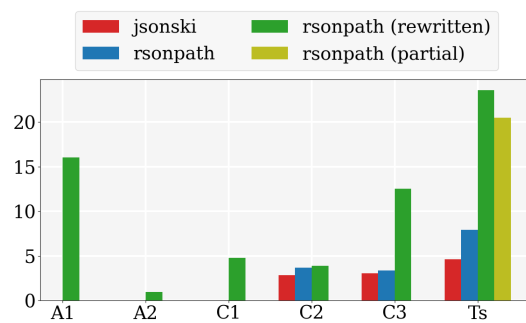
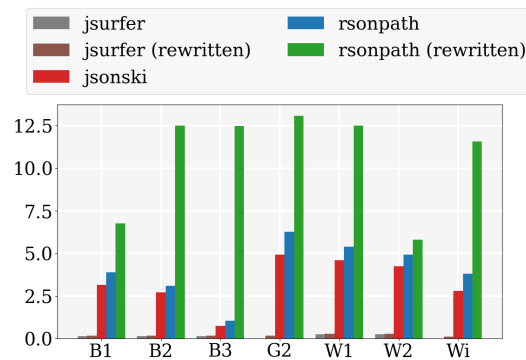
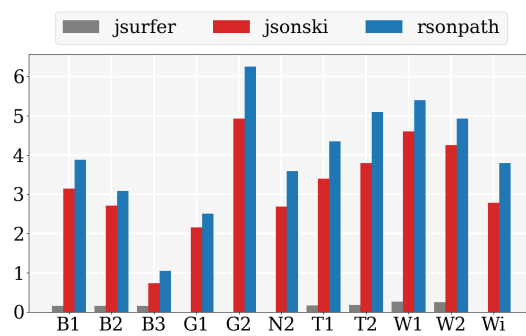
■ **Figure 7** An overview of rsonpath's architecture

C Additional experiments

The CPU information is available from the G5000 website [8]. The RAM information is extracted by the output of the `lshw` command. All value are in GB/s

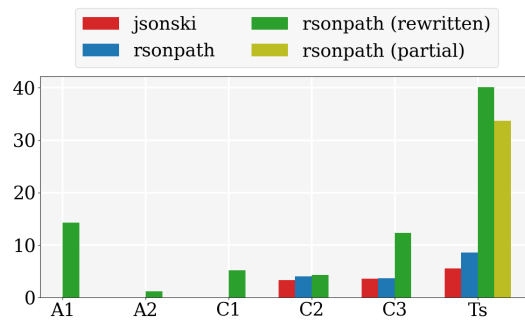
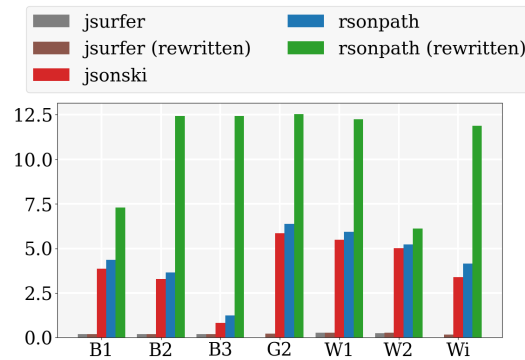
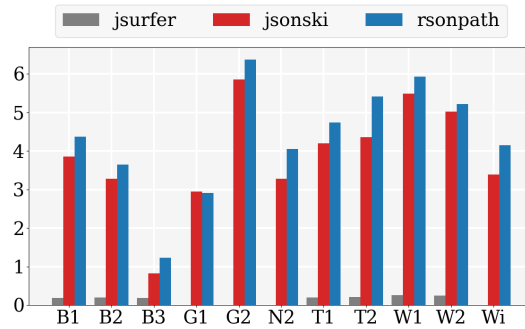
C.1 Broadwell (Chetemi): Throughput comparison

- CPU: Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz (Broadwell)
- RAM: 8x16GiB DIMM DDR4 Synchronous Registered (Buffered) 2400 MHz



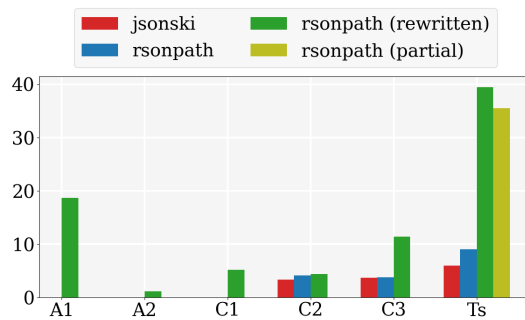
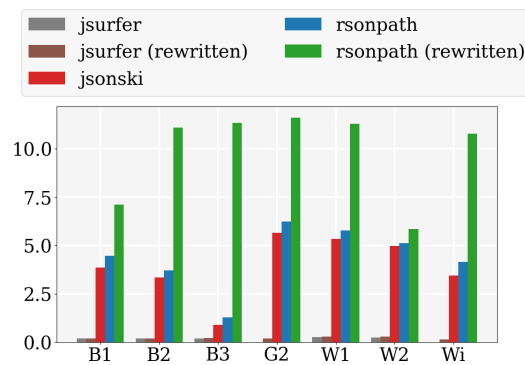
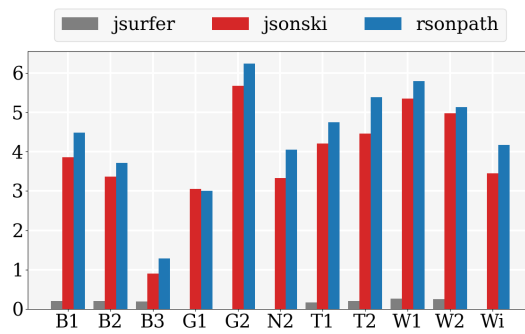
C.2 Skylake (Chiffnot): Throughput comparison

- CPU: Intel(R) Xeon(R) Gold 6126 CPU @ 2.60GHz (Skylake)
- RAM: 12x16GiB DIMM DDR4 Synchronous Registered (Buffered) 2666 MHz



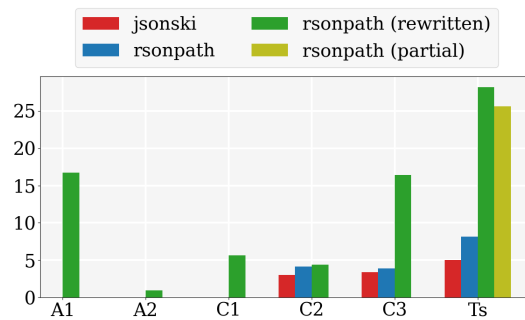
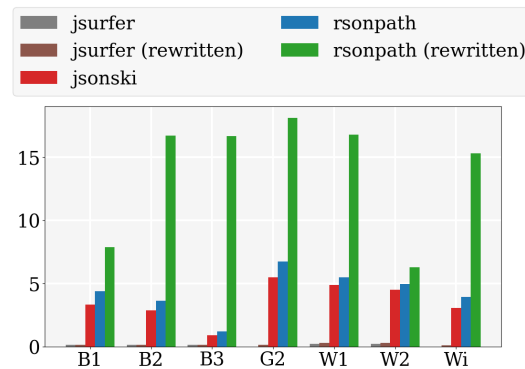
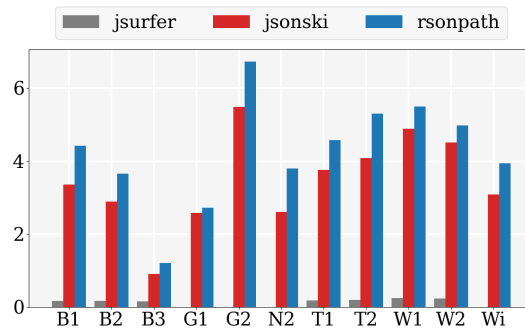
C.3 Cascade-Lake (Troll): Throughput comparison

- CPU: Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz (Cascade-Lake)
- RAM:
 - 12x32GiB DIMM DDR4 Synchronous Registered (Buffered) 2933 MHz
 - 12x128GiB DIMM DDR4 Synchronous Non-volatile LRDIMM 2666 MHz



C.4 Zen 2 (Servan): Throughput comparison

- CPU: AMD EPYC 7352 (Zen 2)
- RAM: 8x16GiB DIMM DDR4 Synchronous Registered (Buffered) 3200 MHz



D Experiment results in tabular format■ **Table 8** All queries used in the experiments

id	rsonpath_id	dataset	query
A1	decl_name	ast	\$.decl.name
A2	nested_inner	ast	\$.inner..inner..type.qualType
A3	included_from	ast	\$.loc.includedFrom.file
B1	BB1_products_category	bestbuy_large_record	\$.products[*].categoryPath[*].id
B1r	BB1'_products_category	bestbuy_large_record	\$.categoryPath..id
B2	BB2_products_video	bestbuy_large_record	\$.products[*].videoChapters[*].chapter
B2r	BB2'_products_video	bestbuy_large_record	\$.videoChapters..chapter
B3	BB3_products_video_only	bestbuy_large_record	\$.products[*].videoChapters
B3r	BB3'_products_video_only	bestbuy_large_record	\$.videoChapters
C1	DOI	crossref2	\$.DOI
C2	author_affiliation	crossref2	\$.items[*].author[*].affiliation[*].name
C2r	author_affiliation_descendant	crossref2	\$.author..affiliation..name
C3	editor	crossref2	\$.items[*].editor[*].affiliation[*].name
C3r	editor_descendant	crossref2	\$.editor..affiliation..name
C4	title	crossref2	\$.items[*].title
C4r	title_descendant	crossref2	\$.title
C5	orcid	crossref2	\$.items[*].author[*].ORCID
C5r	orcid_descendant	crossref2	\$.author..ORCID
G1	GMD1_routes	google_map_large_record	\$.routes[*].legs[*].steps[*].distance.text
G2	GMD2_travel_modes	google_map_large_record	\$.available_travel_modes
G2r	GMD2'_travel_modes	google_map_large_record	\$.available_travel_modes
N1	NSPL1_meta_columns	nspl_large_record	\$.meta.view.columns[*].name
N2	NSPL2_data	nspl_large_record	\$.data[*][*][*]
O1	vitamins_tags	openfood	\$.products[*].vitamins_tags
O1r	vitamins_tags_descendant	openfood	\$.vitamins_tags
O2	added_counties_tags	openfood	\$.products[*].added_countries_tags
O2r	added_countries_tags_descendant	openfood	\$.added_countries_tags
O3	specific_ingredients	openfood	\$.products[*].specific_ingredients[*].ingredient
O3r	specific_ingredients_descendant	openfood	\$.specific_ingredients..ingredient
S0	scalability_affiliation0	crossref0	\$.affiliation..name
S1	scalability_affiliation1	crossref1	\$.affiliation..name
S2	scalability_affiliation2	crossref2	\$.affiliation..name
S4	scalability_affiliation4	crossref4	\$.affiliation..name
T1	TT1_entities_urls	twitter_large_record	\$.entities.urls[*].url
T2	TT2_text	twitter_large_record	\$.text
Ts1	metadata_1	twitter	\$.search_metadata.count
Ts2	metadata_2	twitter	\$.search_metadata.count
Ts3	metadata_3	twitter	\$.count
Ts4	all_hashtags	twitter	\$.hashtags..text
Ts5	hashtags_of_retweets	twitter	\$.retweeted_status..hashtags..text
W1	WM1_items_price	walmart_large_record	\$.items[*].bestMarketplacePrice.price
W1r	WM1'_items_price	walmart_large_record	\$.bestMarketplacePrice.price

XX:36 Supporting Descendants in SIMD-Accelerated JSONPath**Table 8** All queries used in the experiments

id	rsonpath_id	dataset	query
W2	WM2_items_name	walmart_large_record	\$.items[*].name
W2r	WM2'_items_name	walmart_large_record	\$.name
Wi	WP1_claims_p150	wiki_large_record	\$.claims.P150[*].mainsnak.property
Wir	WP1'_claims_p150	wiki_large_record	\$.P150..mainsnak.property

■ **Table 9** Results for all queries over all architectures

id	servan			chetemi			chiffnot			troll	
	rsonpath	jsonski	jsurfer	rsonpath	jsonski	jsurfer	rsonpath	jsonski	jsurfer	rsonpath	jsonski
A1	16.7		0.1	16.0		0.1	14.3		0.1	18.7	
A2	1.0		0.1	1.0		0.1	1.1		0.1	1.2	
A3	12.0		0.1	13.2		0.1	13.0		0.1	14.8	
B1	4.4	3.4	0.2	3.9	3.2	0.2	4.4	3.8	0.2	4.5	3.0
B1r	7.9		0.2	6.8		0.2	7.3		0.2	7.1	
B2	3.7	2.9	0.2	3.1	2.7	0.2	3.6	3.3	0.2	3.7	3.0
B2r	16.7		0.2	12.5		0.2	12.4		0.2	11.1	
B3	1.2	0.9	0.2	1.0	0.7	0.2	1.2	0.8	0.2	1.3	0.0
B3r	16.7		0.2	12.5		0.2	12.4		0.2	11.3	
C1	5.6		0.2	4.8		0.2	5.2		0.2	5.2	
C2	4.2	3.0	0.2	3.7	2.8	0.2	4.0	3.3	0.2	4.2	3.0
C2r	4.4		0.2	3.9		0.2	4.3		0.2	4.4	
C3	3.9	3.4	0.2	3.4	3.1	0.2	3.7	3.6	0.2	3.8	3.0
C3r	16.4		0.2	12.5		0.2	12.3		0.2	11.4	
C4	3.7	3.1	0.2	3.3	2.9	0.2	3.6	3.4	0.2	3.7	3.0
C4r	11.9		0.2	9.2		0.2	9.4		0.2	9.2	
C5	4.0	2.9	0.2	3.5	2.7	0.2	3.8	3.1	0.2	3.9	3.0
C5r	3.9		0.2	3.4		0.2	3.8		0.2	3.9	
G1	2.7	2.6		2.5	2.2		2.9	2.9		3.0	3.0
G2	6.7	5.5		6.3	4.9		6.4	5.9		6.2	5.0
G2r	18.1		0.2	13.1		0.2	12.6		0.2	11.6	
N1	7.6	∞		6.9	∞		6.9	∞		6.7	∞
N2	3.8	2.6		3.6	2.7		4.0	3.3		4.0	3.0
O1	3.0	2.0	0.2	2.7	1.8	0.2	3.1	2.0	0.2	3.3	2.0
O1r	34.6		0.2	20.1		0.1	28.8		0.2	30.9	
O2	8.1	4.3	0.2	7.8	3.9	0.2	8.6	4.5	0.2	9.1	4.0
O2r	27.4		0.2	17.4		0.1	24.5		0.2	25.4	
O3	4.2	2.9	0.2	3.7	2.6	0.2	4.3	3.0	0.2	4.5	3.0
O3r	26.0		0.2	16.5		0.1	22.6		0.2	23.4	
S0	9.8			9.6			8.3			8.3	
S1	10.0			9.5			8.2			8.1	
S2	10.1			7.9			8.2			8.1	
S4	9.8			9.6			8.1			8.1	
T1	4.6	3.8	0.2	4.4	3.4	0.2	4.7	4.2	0.2	4.7	4.0
T2	5.3	4.1	0.2	5.1	3.8	0.2	5.4	4.4	0.2	5.4	4.0
Ts	8.2	5.0	0.3	7.9	4.6	0.3	8.6	5.6	0.3	9.0	6.0
Ts4	18.1		0.3	14.6		0.2	20.4		0.3	21.4	
Ts5	8.5	4.1	0.3	8.2	3.9	0.2	10.3	4.7	0.3	10.6	5.0
Tsp	25.6	5.0	0.3	20.5	4.6	0.2	33.8	5.6	0.3	35.6	5.0
Tsr	28.2	5.0	0.3	23.6	4.6	0.2	40.1	5.6	0.3	39.5	5.0
W1	5.5	4.9	0.2	5.4	4.6	0.3	5.9	5.5	0.3	5.8	5.0
W1r	16.8		0.3	12.5		0.3	12.2		0.3	11.3	
W2	5.0	4.5	0.2	4.9	4.2	0.3	5.2	5.0	0.2	5.1	5.0
W2r	6.3		0.3	5.8		0.3	6.1		0.3	5.8	
Wi	3.9	3.1		3.8	2.8		4.2	3.4		4.2	3.0
Wir	15.3		0.1	11.6		0.1	11.9		0.2	10.8	

E JSONPath implementations – node and path semantics

Using the `json-path-comparison` project [9] we ran a comparison of existing implementations of JSONPath w.r.t. node and path semantics, as described in Section 2. Results are shown in Table 10 below. We used the following JSON document:

```
{
  "person": {
    "name": "A",
    "thesis": {
      "name": "B",
      "advisors": [
        {
          "person": {
            "name": "C"
          }
        },
        {
          "person": {
            "name": "D"
          }
        }
      ]
    }
  }
}
```

The query witnessing the difference between the semantics is `$.person.name`. Ignoring the ordering, the expected results are:

- ["A", "B", "C", "D"] under node semantics; or
- ["A", "B", "C", "D", "C", "D"], under path semantics.

To reproduce the experiment, put the JSON document into a `source.json` file and execute, from the root of `json-path-comparison`:

```
cat source.json | ./src/with_docker.sh ./src/one_off.sh '$.person.name';
```

■ **Table 10** Semantics chosen by known JSONPath implementations. Node semantics is highlighted in dark grey. Light grey indicates errors.

Implementation	Output	Classification
Bash JSONPath.sh	["A", "B", "C", "D"]	node
C json-glib	["A", "B", "C", "D"]	node
Clojure json-path	["A", "B", "C", "D", "C", "D"]	path
C++ jsoncons	["A", "B", "C", "D", "C", "D"]	path
Dart json_path	["A", "B", "C", "D", "C", "D"]	path
Elixir ExJsonPath	["A", "B", "C", "D", "C", "D"]	path
Elixir jaxon	["A"]	error
Elixir warpath	["A", "B", "C", "D", "C", "D"]	path
Erlang ejsonpath	["A", "B", "C", "D", "C", "D"]	path
Go PaesslerAG/jsonpath	["A", "B", "C", "D", "C", "D"]	path
Go jsonslice	[["A", "B", "C", "D"], ["C"], ["D"]]	path ²
Go ojg	["C", "D", "C", "D", "A", "B"]	path
Go oliveagle/jsonpath	not supported	error
Go ajson	["A", "C", "D", "B", "C", "D"]	path
Go yaml-jsonpath	["A", "B", "C", "D", "C", "D"]	path
Haskell jsonpath	["A", "B", "C", "D", "C", "D"]	path
JavaScript Goessner	["A", "B", "C", "D", "C", "D"]	path
JavaScript brunerd	["A", "B", "C", "D", "C", "D"]	path
JavaScript jsonpath	["A", "B", "C", "D", "C", "D"]	path
JavaScript jsonpath-plus	["A", "B", "C", "D", "C", "D"]	path
Java jsurfer	["A", "B", "C", "D"]	node
Java jsonpath	["A", "B", "C", "D", "C", "D"]	path
Kotlin jsonpathkt	["A", "B", "C", "D", "C", "D"]	path
Objective-C SMJJSONPath	["A", "B", "C", "D", "C", "D"]	path
PHP Goessner	["A", "B", "C", "D", "C", "D"]	path
PHP galbar/jsonpath	["A", "B", "C", "D"]	node
PHP remorhaz/jsonpath	["A", "B", "C", "D"]	node
PHP softcreatr/jsonpath	["A", "B", "C", "D", "C", "D"]	path
Perl JSON-Path	["A", "B", "C", "D", "C", "D"]	path
Python jsonpath	["A", "B", "C", "D", "C", "D"]	path
Python jsonpath-ng	["A", "B", "C", "D", "C", "D"]	path
Python jsonpath-rw	["A", "B", "C", "D", "C", "D"]	path
Python jsonpath2	["A", "B", "C", "D", "C", "D"]	path
Raku JSON-Path	["C", "D"]	error
Ruby jsonpath	["A", "B", "C", "D", "C", "D"]	path
Rust jsonpath	not supported	error
Rust jsonpath_lib	["A", "B", "C", "D", "C", "D"]	path
Rust jsonpath_plus	["A", "B", "C", "D", "C", "D"]	path
Scala jsonpath	["A", "B", "C", "D"]	node
Swift Sextant	["A", "B", "C", "D", "C", "D"]	path
.NET Json.NET	["A", "B", "C", "D", "C", "D"]	path
.NET JsonCons.JsonPath	["A", "B", "C", "D", "C", "D"]	path
.NET JsonPath.Net	["A", "B", "C", "D", "C", "D"]	path
.NET JsonPathLib	["A", "B", "C", "D", "C", "D"]	path
.NET Manatee.Json	["A", "B", "C", "D", "C", "D"]	path