Tuplex: Data Science in Python at Native Code Speed

Leonhard Spiegelberg

Rahul Yesantharao[†]
Brown UniversityMalte Schwarzkopf[†]
MIT CSAILTim Kraska[†]

Abstract

Today's data science pipelines often rely on user-defined functions (UDFs) written in Python. But interpreted Python code is slow, and Python UDFs cannot be compiled to machine code easily.

We present Tuplex, a new data analytics framework that just-in-time compiles developers' natural Python UDFs into efficient, end-to-end optimized native code. Tuplex introduces a novel *dual-mode execution model* that compiles an optimized fast path for the common case, and falls back on slower exception code paths for data that fail to match the fast path's assumptions. Dual-mode execution is crucial to making end-to-end optimizing compilation tractable: by focusing on the common case, Tuplex keeps the code simple enough to apply aggressive optimizations. Thanks to dual-mode execution, Tuplex pipelines always complete even if exceptions occur, and Tuplex's post-facto exception handling simplifies debugging.

We evaluate Tuplex with data science pipelines over real-world datasets. Compared to Spark and Dask, Tuplex improves end-to-end pipeline runtime by 5–91× and comes within 1.1–1.7× of a hand-optimized C++ baseline. Tuplex outperforms other Python compilers by 6× and competes with prior, more limited query compilers. Optimizations enabled by dual-mode processing improve runtime by up to 3×, and Tuplex performs well in a distributed setting.

ACM Reference Format:

Leonhard Spiegelberg, Rahul Yesantharao, Malte Schwarzkopf, Tim Kraska. 2021. Tuplex: Data Science in Python at Native Code Speed. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21)*, June 20–25, 2021, Virtual Event, China. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3448016.3457244>

1 Introduction

Data scientists today predominantly write code in Python, as the language is easy to learn and convenient to use. But the features that make Python convenient for programming—dynamic typing, automatic memory management, and a huge module ecosystem—come at the cost of low performance compared to hand-optimized code and an often frustrating debugging experience.

Python code executes in a bytecode interpreter, which interprets instructions, tracks object types, manages memory, and handles exceptions. This infrastructure imposes a heavy overhead, particularly if Python *user-defined functions* (UDFs) are inlined in a larger parallel computation, such as a Spark [71] job. For example, a PySpark job

over flight data [6] might convert a flight's length from kilometers to miles via a UDF after joining with a carrier table:

```
carriers = spark.read.load('carriers.csv')
fun = udf(lambda m: m * 1.609, DoubleType())
spark.read.load('flights.csv')
    .join(carriers, 'code', 'inner')
    .withColumn('distance', fun('distance'))
    .write.csv('output.csv')
```

This code will load data and execute the join using Spark's compiled Scala operators, but must execute the Python UDF passed to the `withColumn` operator in a Python interpreter. This requires passing data between the Python interpreter and the JVM [43], and prevents generating end-to-end optimized code across the UDFs. For example, an optimized pipeline might apply the UDF to distance while loading data from `flights.csv`, which avoids an extra iteration. But the lack of end-to-end code generation prevents this optimization.

Could we instead generate native code (e.g., C++ code or LLVM IR) from the Python UDF and optimize it end-to-end with the rest of the pipeline? Unfortunately, this is not feasible today. Generating, compiling, and optimizing code ahead-of-time that handles all possible code paths through a Python program is not tractable because of the complexity of Python's dynamic typing. Dynamic typing ("duck typing") requires that code always be prepared to handle *any* type: while the above UDF expects a numeric value for `m`, it may actually receive an integer, a float, a string, a null value, or even a list. The interpreter has to handle these possibilities through extra checks and exception handlers, but the sheer number of cases to deal with makes it difficult to compile optimized code even for this simple UDF.

Tuplex is a new analytics framework that generates optimized end-to-end native code for pipelines with Python UDFs. Its key insight is that targeting the common case simplifies code generation. Developers write Tuplex pipelines using a LINQ-style API similar to PySpark's and use Python UDFs without type annotations. Tuplex compiles these pipelines into efficient native code with a new *dual mode execution model*. Dual-mode execution separates the common case, for which code generation offers the greatest benefit, from exceptional cases, which complicate code generation and inhibit optimization but have minimal performance impact. Separating these cases and leveraging the regular structure of LINQ-style pipelines makes Python UDF compilation tractable, as Tuplex faces a simpler and more constrained problem than a general Python compiler.

Making dual-mode processing work required us to overcome several challenges. First, Tuplex must establish what the common case is. Tuplex's key idea is to sample the input, derive the common case from this sample, and infer types and expected cases across the pipeline. Second, the behavior of Tuplex's generated native code must match a semantically-correct Python execution in the interpreter. To guarantee this, Tuplex separates the input data into two row classes: those for which the native code's behavior is identical to Python's, and those for which it isn't and which must be processed

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMOD '21, June 20–25, 2021, Virtual Event, China

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8343-1/21/06...\$15.00

<https://doi.org/10.1145/3448016.3457244>

in the interpreter. Third, Tuplex’s generated code must offer a fast bail-out mechanism if exceptions occur within UDFs (e.g., a division by zero), and resolve these in line with Python semantics. Tuplex achieves this by adding lightweight checks to generated code, and leverages the fact that UDFs are stateless to re-process the offending rows for resolution. Fourth, Tuplex must generate code with high optimization potential but also achieve fast JIT compilation, which it does using tuned LLVM compilation.

Dual mode processing enables compilation, but has another big advantage: it can help developers write more robust pipelines that never fail at runtime due to dirty data or unhandled exceptions. Tuplex detects exception cases, resolves them via slow-path execution if possible, and presents a summary of the unresolved cases to the user. This helps prototype data wrangling pipelines, but also helps make production pipelines more robust to data glitches.

The focus of this paper is primarily on multi-threaded processing on a single server, but Tuplex is a distributed system, and we show results for a preliminary backend based on AWS lambda functions. In summary, we make the following principal contributions:

- (1) We combine ideas from query compilation with speculative compilation techniques in the *dual-mode processing* model for data analytics: an optimized common-case code path processes the bulk of the data, and a slower fallback path handles rare, non-conforming data without inhibiting optimization.
- (2) We observe that data analytics pipelines with Python UDFs—unlike general Python programs—have sufficient structure to make compilation without type annotations feasible.
- (3) We build and evaluate Tuplex, the first data analytics system to embed a Python UDF compiler with a query compiler.

We evaluated our Tuplex prototype over real-world datasets, including Zillow real estate adverts, a decade of U.S. flight data [6], and web server logs from a large university. Tuplex outperforms single-threaded Python and Pandas by 5.8–18.7×, and parallel Spark and Dask by 5.1–91× (§6.1). Tuplex outperforms general-purpose Python compilers by 6–24×, and its generated code comes within 2× of the performance of Weld [52] and Hyper [27] for pure query execution time, while achieving 2–7× faster end-to-end runtime in a realistic data analytics setting (§6.3). Tuplex’s dual-mode processing facilitates end-to-end optimizations that improve runtime by up to 3× over simple UDF compilation (§6.4). Finally, Tuplex performs well on a single server and distributedly across a cluster of AWS Lambda functions (§6.5); and anecdotal evidence suggests that it simplifies the development and debugging of data science pipelines (§7). Tuplex is open-source at <https://tuplex.cs.brown.edu>.

2 Background and Related Work

Many prior attempts to speed up data science via compilation or to compile Python to native code exist, but they fall short of the ideal of compiling end-to-end optimized native code from UDFs written in natural Python. We discuss key approaches and systems in the following; Table 1 summarizes the key points.

Python compilers. Building compilers for arbitrary Python programs, which lack the static types required for optimizing compilation, is challenging. PyPy [57] reimplements the Python interpreter in a compilable subset of Python, which it JIT-compiles via LLVM (i.e., it creates a self-compiling interpreter). GraalPython [50] uses

the Truffle [25] language interpreter to implement a similar approach while generating JVM bytecode for JIT compilation. Numba [32] JIT-compiles Python bytecode for annotated functions on which it can perform type inference; it supports a subset of Python and targets array-structured data from numeric libraries like NumPy [2].

All of these compilers either myopically focus on optimizing hotspots without attention to high-level program structure, or are limited to a small subset of the Python language (e.g., numeric code only, no strings or exceptions). Pyston [41] sought to create a full Python compiler using LLVM, but faced memory management and complexity challenges [40], and offers only a 20% performance gain over the interpreter in practice [42].

Python transpilers. Other approaches seek to cross-compile Python into other languages for which optimizing compilers exist. Cython [4] unrolls the CPython interpreter and a Python module into C code, which interfaces with standard Python code. Nuitka [18] cross-compiles Python to C++ and also unrolls the interpreter when cross-compilation is not possible. The unrolled code represents a specific execution of the interpreter, which the compiler may optimize, but still runs the interpreter code, which compromises performance and inhibits end-to-end optimization.

Data-parallel IRs. Special-purpose native code in libraries like NumPy can speed up some UDFs [24], but such pre-compiled code precludes end-to-end optimization. Data-parallel intermediate representations (IRs) such as Weld [52] and MLIR [33] seek to address this problem. Weld, for example, allows cross-library optimization and generates code that targets a common runtime and data representation, but requires libraries to be rewritten in Weld IR. Rather than requiring library rewrites, Mozart [53] optimizes cross-function data movement for lightly-annotated library code. All of these lack a general Python UDF frontend, assume static types, and lack support for exceptions and data type mismatches.

Query compilers. Query compilers turn SQL into native code [1, 29, 60, 62, 72], and some integrate with frameworks like Spark [14]. The primary concern of these compilers is to iterate efficiently over preorganized data [28, 61], and all lack UDF support, or merely provide interfaces to call precompiled UDFs written e.g. in C/C++.

Simple UDF compilers. UDF compilation differs from traditional query compilation, as SQL queries are declarative expressions. With UDFs, which contain imperative control flow, standard techniques like vectorization cannot apply. While work on peeking inside imperative UDFs for optimization exists [20], these strategies fail on Python code. Tuplware [8] provides a UDF-aware compiler that can apply some optimizations to black-box UDFs, but its Python integration relies on static type inference via PYLLVM [19], and it lacks support for common features like optional (None-valued) inputs, strings, and exceptions in UDFs. Tuplex supports all of these.

Exception handling. Inputs to data analytics pipelines often include “dirty” data that fails to conform to the input schema. This data complicates optimizing compilation because it requires checks to detect anomalies and exception handling logic. Load reject files [10, 39, 56] help remove ill-formed inputs, but they solve only part of the problem, as UDFs might themselves encounter exceptions when processing well-typed inputs (e.g., a division by zero, or None values). Graal speculatively optimizes for exceptions [13] via polymorphic

System Class	Examples	Limitations
Tracing JIT Compilers	PyPy [57], Pyston [41]	Require tracing to detect hotspots, cannot reason about high-level program structure, generated code must cover full Python semantics (slow).
Special Purpose JIT Compilers	Numba [32], XLA [34], Glow [58]	Only compile well-formed, statically typed code, enter interpreter otherwise; use their own semantics, which often deviate from Python's.
Python Transpilers	Cython [4], Nuitka [18]	Unrolled interpreter code is slow and uses expensive Python object representation.
Data-parallel IRs	Weld [52], MLIR [33]	No compilation from Python; static typing and lack exception support.
SQL Query Compilers	Flare [14], Hyper [47]	No Python UDF support.
Simple UDF Compiler	Tupeware [8]	Only supports UDFs for which types can be inferred statically, only numerical types, no exception support, no polymorphic types (e.g., NULL values).

Table 1: Classes of existing systems that compile data analytics pipelines or Python code. All have shortcomings that either prevent full support for Python UDFs, or prevent end-to-end optimization or full native-code performance.

inline caches—an idea also used in the V8 JavaScript engine—but the required checks and guards impose around a 30% overhead [12]. Finally, various dedicated systems track the impact of errors on models [30] or provide techniques to compute queries over dirty data [66, 68], but they do not integrate well with compiled code.

Speculative processing. Programming language research on speculative compilation pioneered native code performance for dynamically-typed languages. Early approaches, like SELF [7], compiled multiple, type-specialized copies of each control flow unit (e.g., procedure) of a program. This requires variable-level speculation on types, and results in a large amount of generated code. State-of-the-art tracing JITs apply a dynamic variant of this speculation and focus on small-scale “hot” parts of the code only (e.g., loops).

A simpler approach than trying to compile general Python is to have Python merely act as a frontend that calls into a more efficient backend. Janus [21, 22] applies this idea to TensorFlow, and Snek [11] takes it one step further by providing a general mechanism to translate imperative Python statements of any framework into calls to a framework’s backend. While these frameworks allow for imperative programming, the execution can only be efficient for Python code that maps to the operators offered by the backend. To account for Python’s dynamic types, such systems may have to speculate on which backend operators to call. In addition, the backend’s APIs may impose in-memory materialization points for temporary data, which reduce performance as they add data copies.

In big data applications, efficient data movement is just as important as generating good code: better data movement can be sufficient to outperform existing JIT compilers [53]. Gerenuk [46] and Skyway [48] therefore focus on improving data movement by specializing serialization code better within the HotSpot JVM.

Tuplex. In Tuplex, UDFs are first-class citizens and are compiled just-in-time when a query executes. Tuplex solves a more *specialized* compilation problem than general Python compilers, as it focuses on UDFs with mostly well-typed, predictable inputs. Tuplex compiles a fast path for the common-case types (determined from the data) and expected control flow, and defers rows not suitable for this fast path to slower processing (e.g., in the interpreter). This simplifies the task sufficiently to make optimizing compilation tractable.

Tuplex supports natural Python code rather than specific libraries (unlike Weld or Numba), and optimizes the full end-to-end pipeline, including UDFs, as a single program. Tuplex generates at most three different code paths to bound the cost of specialization (unlike SELF); and it speculates on a per-row basis, compared to a per-variable basis

in SELF and whole-program speculation in Janus. Tuplex uses the fact that UDFs are embedded in a LINQ-style program to provide high-level context for data movement patterns and to make compilation tractable. Finally, Tuplex makes exceptions explicit, and handles them without compromising the performance of compiled code: it collects exception-triggering rows and batches their processing, rather than calling the interpreter from the fast path.

3 Tuplex Overview

Tuplex is a data analytics framework with a similar user experience to e.g., PySpark, Dask, or DryadLINQ [70]. A data scientist writes a processing pipeline using a sequence of high-level, LINQ-style operators such as `map`, `filter`, or `join`, and passes UDFs as parameters to these operators (e.g., a function over a row to `map`). E.g., the PySpark pipeline shown in §1 corresponds to the Tuplex code:

```
c = tuplex.Context()
carriers = c.csv('carriers.csv')
c.csv('flights.csv')
  .join(carriers, 'code', 'code')
  .mapColumn('distance', lambda m: m * 1.609)
  .tocsv('output.csv')
```

Like other systems, Tuplex partitions the input data (here, the CSV files) and processes the partitions in a data-parallel way across multiple executors. Unlike other frameworks, however, Tuplex compiles the pipeline into end-to-end optimized native code before execution starts. To make this possible, Tuplex relies on a **dual-mode processing** model structured around two distinct execution modes:

- (1) an optimized, *normal-case* execution; and
- (2) an *exception-case* execution.

To establish what constitutes the normal case, Tuplex samples the input data and, based on the sample, determines the expected types and control flow of the normal-case execution. Tuplex then uses these assumptions to generate and optimize code to classify a row into normal or exception cases, and specialized code for the normal-case execution. It lowers both to optimized machine code via LLVM.

Tuplex then executes the pipeline. The generated classifier code performs a single, cheap initial check on each row to determine if it can proceed with normal-case execution. Any rows that fail this check are placed in an exception pool for later processing, while the majority of rows proceed to optimized normal-case execution. If any exceptions occur during normal-case execution, Tuplex moves the offending row to the exception pool and continues with the next row. Finally, after normal-case processing completes, Tuplex attempts to resolve the exception-case rows. Tuplex automatically resolves some

Figure 1: Tuplex uses an input sample to compile specialized code for normal-case execution (blue, left), which processes most rows, while the exception-case (red, right) handles the remaining rows. Compiled parts are shaded in yellow.

exceptions using general, but slower code or using the Python interpreter, while for other exceptions it uses (optional) user-provided resolvers. If resolution succeeds, Tuplex merges the result row with the normal-case results; if resolution fails, it adds the row to a pool of failed rows to report to the user.

In our example UDF, a malformed flight row that has a non-numeric string in the distance column will be rejected and moved to the exception pool by the classifier. By contrast, a row with distance set to None, enters normal-case execution if the sample contained a mix of non-None and None values. However, the normal-case execution encounters an exception when processing the row and moves it to the exception pool. To tell Tuplex how to resolve this particular exception, the pipeline developer can provide a resolver:

```
# ...
.join(carriers, 'code', 'code')
.mapColumn('distance', lambda m: m * 1.609)
.resolve(TypeError, lambda m: 0.0)
# ...
```

Tuplex then merges the resolved rows into the results. If no resolver is provided, Tuplex reports the failed rows separately.

4 Design

Tuplex’s design is derived from two key insights. First, Tuplex can afford slow processing for exception-case rows with negligible impact on overall performance if such rows are rare, which is the case if the sample is representative. Second, specializing the normal-case execution to common-case assumptions simplifies the generated logic by deferring complexity to the exception case, which makes JIT compilation tractable and allows for aggressive optimization.

4.1 Abstraction and Assumptions

Tuplex’s UDFs contain natural Python code, and Tuplex must ensure that their execution behaves exactly as it would have in a Python interpreter. We make only two exceptions to this abstraction. First, Tuplex never crashes due to unhandled top-level exceptions, but instead emulates an implicit catch-all exception handler that records unresolved (“failed”) rows. Second, Tuplex assumes that UDFs are

pure and stateless, meaning that their repeated execution (on the normal and exception paths) has no observable side-effects.

The top-level goal of matching Python semantics influences Tuplex’s design and implementation in several important ways, guiding its code generation, execution strategy, and optimizations.

4.2 Establishing the Normal Case

The most important guidance for Tuplex to decide what code to generate for normal-case execution comes from the observed structure of a sample of the input data. Tuplex takes a sample of configurable size every time a pipeline executes, and records statistics about structure and data types in the sample, as follows.

Row Structure. Input data may be dirty and contain different column counts and column orders. Tuplex counts the columns in each sample row, builds a histogram and picks the prevalent column structure as the normal case.

Type Deduction. For each sample row, Tuplex deduces each column type based on a histogram of types in the sample. If the input consists of typed Python objects, compiling the histogram is simple. If the input is text (e.g., CSV files), Tuplex applies heuristics. For example, numeric strings that contain periods are floats, integers that are always zero or one and the strings “true” and “false” are booleans, strings containing JSON are dictionaries, and empty values or explicit “NULL” strings are None values. If Tuplex cannot deduce a type, it assumes a string. Tuplex then uses the most common type in the histogram as the normal-case type for each column (except for null values, described below).

This *data-driven* type deduction contrasts with classic, static type inference, which seeks to infer types from program code. Tuplex uses data-driven typing because Python UDFs often lack sufficient information for static type inference without ambiguity, and because the actual type in the input data may be different from the developer’s assumptions. In our earlier example (§3), for instance, the common-case type of `m` may be `int` rather than `float`.

For UDFs with control flow that Tuplex cannot annotate with sample-provided input types, Tuplex uses the AST of the UDF to trace the input sample through the UDF and annotates individual nodes with type information. Then, Tuplex determines the common cases within the UDF and prunes rarely visited branches. For example, Python’s power operator (`**`) can yield integer or float results, and Tuplex picks the common case from the sample trace execution.

Option types (None). Optional column values (i.e., “nullable”) are common in real-world data, but induce potentially expensive logic in the normal case. Null-valued data corresponds to Python’s `None` type, and a UDF must be prepared for *any* input variable (or nested data, e.g., in a list-typed row) to potentially be `None`. To avoid having to check for `None` in cases where null values are rare, Tuplex uses the sample to guide specialization of the normal case. If the frequency of null values exceeds a threshold δ , Tuplex assumes that `None` is the normal case; and if the frequency of null values is below $1 - \delta$, Tuplex assumes that null values are an exceptional case. For frequencies in $(1 - \delta, \delta)$, Tuplex uses a polymorphic optional type and generates code for the necessary checks.

4.3 Code Generation

Having established the normal case types and row structure using the sample, Tuplex generates code for compilation. At a high level, this involves parsing the Python UDF code in the pipeline, typing the abstract syntax tree (AST) with the normal-case types, and generating LLVM IR for each UDF. The type annotation step is crucial to making UDF compilation tractable, as it reduces the complexity of the generated code: instead of being prepared to process any type, the generated code can assume a single static type assignment.

In addition, Tuplex relies on properties of the data analytics setting and the LINQ-style pipeline API to simplify code generation compared to general, arbitrary Python programs:

- (1) UDFs are “closed” at the time the high-level API operator (e.g., map or filter) is invoked, i.e., they have no side-effects on the interpreter (e.g., changing global variables or redefining opcodes).
- (2) The lifetime of any object constructed or used when a UDF processes a row expires at the end of the UDF, i.e., there is no state across rows (except as maintained by the framework).
- (3) The pipeline structures control flow: while UDFs may contain arbitrary control flow, they always return to the calling operator and cannot recurse.

Tuplex’s generated code contains a *row classifier*, which processes all rows, and two code paths: the optimized *normal-case* code path, and a *general-case* code path with fewer assumptions and optimizations. The general-case path is part of exception-case execution, and Tuplex uses it to efficiently resolve some exception rows.

Row Classifier. Tuplex must classify all input rows according to whether they fit the normal case. Tuplex generates code for this classification: it checks if each column in a row matches the normal-case structure and types, and directly continues processing the row on the normal-case path if so. If the row does not match, the generated classifier code copies it out to the exception row pool for later processing. This design ensures that normal-case processing is focused on the core UDF logic, rather including exception resolution code that adds complexity and disrupts control flow.

Code Paths. All of Tuplex’s generated code must obey the top-level invariant that execution must match Python semantics. Tuplex traverses the Python AST for each UDF and generates matching LLVM IR for the language constructs it encounters. Where types are required, Tuplex instantiates them using the types derived from the sample, but applies different strategies in the normal-case and general-case code. In the normal-case code, Tuplex assumes the common-case types from the sample always hold and emits no logic to check types (except for the option types used with inconclusive null value statistics, which require checks). The normal-case path still includes code to detect cases that trigger exceptions in Python: e.g., it checks for a zero divisor before any division.

By contrast, the general-case path always assumes the most general type possible for each column. For example, it includes option type checks for all columns, as exception rows may contain nulls in any column. In addition, the general-case path embeds code for any user-provided resolvers whose implementation is a compilable UDF. But it cannot handle all exceptions, and must defer rows from the exception pool that it cannot process. The general-case path

Figure 2: Tuplex’s exception case consists of a compiled general path and a fallback path that invokes the Python interpreter. Exceptions (red) move rows between code paths.

therefore includes logic that detects these cases, converts the data to Python object format, and invokes the Python interpreter inline.

Tuplex compiles the pipeline of high-level operators (e.g., map, filter) into stages, similar to Neumann [47], but generates up to three (fast, slow, and interpreter) code paths. Tuplex generates LLVM IR code for each stage’s high-level operators, which call the LLVM IR code previously emitted for each UDF. At the end of each stage, Tuplex merges the rows produced by all code paths.

Memory Management. Because UDFs are stateless functions, only their output lives beyond the end of the UDF. Tuplex therefore uses a simple slab allocator to provision memory from a thread-local, pre-allocated region for new variables within the UDF, and frees the entire region after the UDF returns and Tuplex has copied the result.

Exception handling. To simulate a Python interpreter execution, the code Tuplex generates and executes for a row must have no observable effects that deviate from complete execution in a Python interpreter. While individual code paths do not always meet this invariant, their combination does. Tuplex achieves this via *exceptions*, which it may generate in three places: when classifying rows, on the normal-case path, and on the general-case code path. Figure 2 shows how exceptions propagate rows between the different code paths.

Rows that fail the row classifier and those that generate exceptions on the normal-case code path accumulate in the exception row pool. When Tuplex processes the exception row pool, it directs each row either to the general-case code path (if the row is suitable for it) or calls out to the Python interpreter. Any rows that cause exceptions on the general-case path also result in a call into the interpreter. An interpreter invocation constitutes Tuplex’s third code path, the *fallback code path*. It starts the UDF over, running the entire UDF code over a Python object version of the row. Finally, if the pipeline developer provided any resolvers, compilable resolvers execute on the general-case code path, and all resolvers execute on the fallback path. If the fallback path still fails, Tuplex marks the row as failed.

Consequently, Tuplex may process a row a maximum of three times: once on the normal-case path, once on the general-case path, and once on the fallback path. In practice, only a small fraction of rows are processed more than once.

4.4 Execution

Tuplex executes pipelines similar to a typical data analytics framework, although customized to handle end-to-end UDF compilation.

Tuplex has a logical planner, which applies logical optimizations (e.g., operator reordering and filter pushdown); a physical planner, which splits the pipeline execution into distinct stages; and a UDF compiler, which handles the actual code generation. However, the typing requirements of Tuplex’s dual-mode processing model permeate all these components. For example, the logical planner also types the UDFs according to the normal-case types deduced from the sample in order to allow for type-aware logical optimizations.

Stages. A *stage* is a sequence of operators, including UDFs, that is bounded on either side by an operator that consumes materialized data from memory or requires generating it. Examples of such operators include inputs, joins, aggregations, and outputs. Stages are also the unit of code generation: Tuplex generates and executes a normal-case and an exception-case code path for each stage. The materialized output of a stage may initially consist only of normal-case result rows, though some operators require immediate production and materialization of resolved exception-case rows too (see §4.5).

Tuplex delineates stages similarly to HyPer [47]. Tuplex makes stages as long as possible to facilitate compiler optimizations, and so that rows are processed through many UDFs while in CPU cache. Ideally, most input rows proceed through a single, highly-optimized stage that ends with the pipeline’s materialized output.

4.5 Joins

Tuplex uses a hash join, which materializes records on one side of the join (the “build” side) and streams rows on the other side to look up into the hash table. Tuplex chooses the smaller side as the build side and terminates a stage at the materialized join input.

This standard design, however, requires adaptation for dual-mode processing. A classic data-parallel join works because the data on both sides of the join is partitioned by the same key. For join $A \bowtie B$ between relations A and B , it suffices to join each $A_i \bowtie B_i$. But in dual-mode execution, each partition of A is itself split into normal-case rows $NC(A_i)$ and exception-case rows $EC(A_i)$, and likewise for B . For correct results, Tuplex must compute each pairwise join:

$$\begin{aligned} NC(A_i) \bowtie NC(B_i) \cup NC(A_i) \bowtie EC(B_i) \cup \\ EC(A_i) \bowtie NC(B_i) \cup EC(A_i) \bowtie EC(B_i) \end{aligned}$$

To compute the joins between normal-case and exception-case rows, Tuplex would have to execute all three code paths for *both* join inputs and materialize the input rows in memory. This conflicts with the goal of long stages that keep caches hot on the normal path and avoid unnecessary materialization. Instead, Tuplex executes all code paths for the build side of the join and resolves its exception rows before executing any code path of the other side. If the build side is B and the result of resolving exception rows of B_i is $R(B_i) = NC(B_i) \cup \text{resolve}(EC(B_i))$, Tuplex then executes $NC(A_i) \bowtie R(B_i)$ as part of a longer stage and without materializing $NC(A_i)$.

4.6 Aggregates

Dual-mode processing works for aggregations as long as the aggregation function is associative. Tuplex separately aggregates normal-case rows and, subsequently, exception-case rows via the general and fallback code paths; in a final merge step, it combines the partial aggregates into a final result. This merging of partial aggregates happens at the end of the stage after resolving exception rows.

Aggregations are compatible with Tuplex’s assumption that UDFs are stateless, as the framework tracks the accumulated state across rows. To make this work, the aggregation operator needs to take a UDF with a row argument and an accumulator argument, and return an updated accumulator. For example, `.aggregate`’s UDF signature is `lambda acc, r: acc + r['col']`, where `acc` is an accumulator (e.g., an integer, a list or a more complicated object like a nested tuple or dictionary). Tuplex’s runtime is responsible for managing the memory of `acc`, and the UDF remains stateless.

4.7 Optimizations

Tuplex applies several optimizations to the normal-case path.

Logical optimizations. Pushing selective operators (e.g., filters, projections) to the start of the pipeline is a classic database optimization. Yet, systems that treat Python UDFs as black box operators cannot apply this optimization across UDFs. Tuplex’s logical planner analyzes UDFs’ Python ASTs to determine which input objects are preserved, dropped, and modified by each UDF. Based on this knowledge, Tuplex then reorders operators to preserve columns only as long as needed. Another, more complex optimization pushes UDFs that modify a column past any operators and UDFs that do not read it. This helps e.g., push UDFs that rewrite non-key columns below joins, which is a good choice if the join is selective. Crucially, this optimization is possible because Tuplex analyzes the Python UDFs.

UDF-specific optimizations. Tuplex applies standard compiler optimizations like constant folding to Python UDFs. In addition, Tuplex applies optimizations specific to UDFs as part of a LINQ-style pipeline. For example, Tuplex rewrites dictionaries with string keys known at compile time into tuples (avoiding string operations); Tuplex flattens nested tuples to avoid pointer indirection; and Tuplex optimizes for the common case in nullable values, i.e., column types can get specialized to `NULL`, `Option[T]` or `T`.

Code generation optimizations. On the normal-case path, Tuplex removes any code related to types that it classified as exceptions. Consider, for example, `lambda m: m * 1.609 if m else 0.0`: with an input sample of mostly non-null floats, Tuplex removes code for integer-to-float conversion, null checks, and the `else` branch from the normal-case path. This reduces the generated code from 17 LLVM IR instructions (5 basic blocks) to 9 IR instructions (1 basic block). If the common-case input is null, Tuplex simplifies the normal-case path to 3 IR instructions that return zero.

Compiler optimizations. Once Tuplex has generated LLVM IR for the normal-case path, it applies several LLVM optimizer passes to the code. In particular, we use the Clang 9.0 pass pipeline equivalent to `-O3` which are applied for all UDFs and operators inside a stage.

However, since Tuplex’s generated code must match Python semantics, not all compiler optimizations are valid. For example, some optimizations to speed up floating point math (equivalent to the `-ffast-math` C compiler flag) change the handling of NaN values in ways that fail to match Python. Tuplex avoids these optimizations.

5 Implementation

We implemented a prototype of Tuplex in about 65k lines of C++. Our prototype uses LLVM 9’s ORC-JIT to compile the generated LLVM IR code at runtime. It is implemented as a C-extension (shared library) which users import as a Python module or from a Jupyter

Dataset	Size	Rows	Columns	Files
Zillow	10.0 GB	48.7M	10	1
Flights	5.9 GB	14.0M	110	24
	30.4 GB	69.0M	110	120
Logs	75.6 GB	715.0M	1	3797
311	1.2 GB	197.6M	1	1
TPC-H (SF=10)	1.5 GB	59.9M	4	1

Table 2: Dataset overview (smaller join tables excluded).

Notebook. Tuplex provides a shell in CPython interactive mode and a web UI with a history server, which developers can use to inspect their pipelines’ execution and any failed rows generated.

Multithreaded Execution. On a single server, our prototype runs executors in a thread pool. Executors process input data partitions in individual tasks, which run identical code. Each thread has its own bitmap-managed block manager for memory allocation. When invoking the fallback path, Tuplex acquires the global interpreter lock (GIL) of the parent Python process.

Distributed Execution. Tuplex’s techniques apply both on a single server and in a distributed setting, where many servers process parts of the input data in parallel. For datasets that require this scale-out data parallelism, our prototype supports executing individual processing tasks in serverless AWS Lambda functions over data stored in S3. Tuplex divides each stage into many data-parallel tasks and runs each task in a Lambda function, which reads its input from S3 and writes its output back to S3. The driver machine generates LLVM IR, initiates, and supervises the Lambdas. Each Lambda receives the LLVM IR code of its task from the driver, lowers the IR to machine code, and executes the machine code over its input data.

Exception handling. Tuplex implements exception control flow on the normal-case and general-case paths via special return codes. We found that this is 30% faster than the “zero-cost” Itanium ABI exception handling [36], and allows more optimization than `setjmp/longjmp` (SJLJ) intrinsics [37].

Limitations. Our prototype supports compiling optimized code for many, but not all Python language features. The prototype currently supports compiling integer, float, string, and tuple operations, as well as essential dictionary and list operations. UDFs can be passed either as lambda functions or regular functions and may contain optional type annotations. The prototype supports variables, simple list comprehensions, control flow, random number generation, and regular expressions. It does not yet support while loops, generator expression, try-except, sets, async expressions, classes, objects, nested functions and external modules. For unsupported language features, Tuplex falls back on running the UDF in the Python interpreter. We believe that support for all missing core Python features could be added to our prototype with additional engineering effort.

Our prototype also does not focus on external modules, which could be compiled but often already come with their own native-code backends. Linking Tuplex’s generated LLVM IR with the LLVM IR code produced by library-oriented compilers such as Weld [52], Numba [32] or Bohrium [31] should be feasible in future work.

6 Evaluation

We evaluate Tuplex with three representative pipelines and with microbenchmarks of specific design features. Our experiments seek to answer the following questions:

- (1) What performance does Tuplex achieve for end-to-end data science pipelines, compared to both single-threaded baselines and widely-used parallel data processing frameworks? (§6.1)
- (2) What is the cost of Tuplex’s code paths, and of exception handling? (§6.2)
- (3) How does Tuplex’s performance compare to off-the-shelf Python compilers, such as PyPy, Cython, and Nuitka; and to state-of-the-art query compilers, such as Weld [52] and Hyper [27]? (§6.3)
- (4) What factors affect Tuplex’s performance, and what is the impact of optimizations enabled by Tuplex’s dual-mode processing model? (§6.4)
- (5) How does Tuplex perform when operating distributedly across many servers? (§6.5)

Setup. In most experiments, Tuplex and other systems run on an `r5d.8xlarge` Amazon EC2 instance (16-core Xeon Platinum 8259CL, 2.50 GHz; hyperthreads disabled) with 256 GB RAM, and 2 NVMe SSDs. The input data is CSV-formatted UTF-8 text. We compare Tuplex against Dask (2021.03) and Spark (PySpark, v2.4.7) on Ubuntu 20.04. All systems use 16-way parallelism. All numbers are averages of at least five runs with warmed-up OS caches.

Our focus is Tuplex’s performance on a multi-core server, a common medium-scale analytics setup [14]. But the systems we compare against support scale-out across servers, so we also compare Tuplex’s prototype AWS Lambda backend to Spark (§6.5).

6.1 End-to-End Performance

We measure Tuplex’s end-to-end performance using three data science pipelines, and with the datasets shown in Table 2.

Zillow. Zillow is a real estate directory website whose listings are uploaded by individual brokers. We scraped 38,570 Boston area listings [59], scaled the data to 10 GB, and cleaned it for performance experiments to avoid failures in Spark and Dask. The two queries extract information like the number of bedrooms, bathrooms, and the price from textual data and filter for all houses (Z1) or condos currently for sale (Z2). Each version involves eleven Python UDFs, which perform value conversions, multiple substring extractions, and several simple lookups, as well as filtering out implausible records. The UDF’s operators can execute as a single, large pipelined stage.

Flights. We modeled this workload after a Trifacta tutorial [17] and extended it by joining with additional airport and airline data from other sources (743 KB [54] and 82 KB [5]). The pipeline has one inner and two left joins, as well as UDFs to reconstruct values from multiple columns which can’t be easily expressed in a SQL query. We ran this pipeline on ten years (2009–2019) of CSV data [6].

Weblogs. Based on a Spark use case [9], this pipeline extracts information from twelve years of Apache web server logs obtained from a U.S. university. It converts the Apache log format into a relational representation, and retains records for potentially malicious requests. We extended the original query by an inner join with a list of bad IPs [45] and anonymize any personally-identifiable URLs by replacing usernames (e.g., “~alice”) with random strings.

311 and TPC-H. We use the Pandas cookbook [15] data cleaning query for 311 service requests, which yields a set of unique ZIP codes, to compare to Weld [52]. Finally, we also run microbenchmarks with

Figure 3: Tuplex outperforms single-threaded and parallel alternatives by 5.8×–18.7× when running the Zillow pipeline over 10G of input data, and comes close to hand-tuned C++.

TPC-H Q6 and Q19 to measure Tuplex’s performance compared to Hyper [27], a state-of-the-art SQL query compiler.

6.1.1 Zillow: String-heavy UDFs. In this experiment, we compare Tuplex to other frameworks using the Zillow pipeline. This pipeline contains eleven UDFs, which use Python idioms for substring search (e.g., “bd” in s, or s.find(“bd”)), string splitting, normalization (s.lower()), and type conversions (int, float).

We consider two row representations: (i) as Python tuples, and (ii) as Python dictionaries (hashmaps). The dictionary representation simplifies code by using column names, but typically comes with a performance overhead. Tuplex allows either representation and compiles both representations into identical native code.

Single-threaded execution. We compare standard CPython (v3.6.9), Pandas (v1.1.5), and hand-optimized C++ (via gcc v10.2) and Scala (v2.12.10) baselines to Tuplex configured with a single executor. Tuplex’s end-to-end optimized code might offer an advantage over CPython and Pandas, which call into individual native-code functions (e.g., libc string functions) but cannot optimize end-to-end. Tuplex should ideally come close to the hand-optimized C++.

Figure 3 shows our results. As expected, the CPython implementation with rows represented as dictionaries is substantially slower (about 2×) than the tuple-based implementation. Pandas, perhaps surprisingly, is about 5.5× slower than tuple-based CPython in Z1, and 65% slower than tuple-based CPython in Z2. While Pandas benefits from a faster CSV parser, an efficient data representation (numpy arrays), and specialized native-code operators for numeric computation, its performance suffers because UDFs require converting between numpy and Python data representations. Z2 filters fewer rows early than Z1, which exacerbates this UDF-related cost. Finally, Tuplex completes processing in 33–43 seconds, a speedup of 5.8×–18.7× over the CPython and Pandas implementations. This is 1.9× faster than a single-threaded baseline written in pure Scala, and 1.3–1.7× slower than the hand-optimized C++ implementation. However, this overstates Tuplex’s overhead: in Tuplex, the compute part of Z1 (i.e., excluding I/O) takes 11.2s, 29% slower than the C++ implementation (8.7s); Z2 sees a 0.5% slowdown (19.8s vs. 19.7s).

Figure 4: Tuplex achieves speedups of 7.6–24.6× over PySparkSQL and Dask on the flights pipeline.

Data-parallel execution. Next, we benchmark Tuplex against widely-used frameworks for parallel processing of large inputs: PySpark (v2.4.7) and Dask (2021.03). We configure each system for 16-way parallelism: PySpark uses 16 JVM threads and 16 Python processes for UDFs; Dask uses 16 worker processes; and Tuplex uses 16 executor threads. We benchmark PySpark both with RDDs [71] and with the more efficient SparkSQL operators [3]. Neither PySpark nor Dask compile UDFs to native code or optimize across UDFs, which indicates that Tuplex should outperform them.

Figure 3 confirms that this is the case: Tuplex outperforms the fastest PySpark setup by 15.5× and Dask by 7.5× in Z1. For Z2, Tuplex is 14.5× faster, as the compiled UDFs process more rows. Compared to the single-threaded execution, Tuplex achieves a speedup of 12.2× when using 16 threads (for Z1). We also ran the pipeline in SparkSQL with Scala UDFs rather than Python UDFs, which keeps computation within the JVM and avoids overheads of calling into Python. Tuplex’s end-to-end optimized code is still 5.8–7.1× faster.

These results confirm that Tuplex’s code generation and end-to-end optimization offer performance gains for UDF-heavy pipelines. In §6.3.1, we compare Tuplex to other Python compilers, and §6.4 drills down into the factors contributing to Tuplex’s performance.

6.1.2 Flights: Joins and Null Values. We repeat the comparison between Tuplex, Spark, and Dask for the flights pipeline. This pipeline contains three joins, and the dataset has “sparse” columns, i.e., columns that mostly contain null values, while others have occasional null values complemented by extra columns (e.g., diverted flights landing at a different airport). Tuplex infers the normal-case null value status for each column from its sample, and defers the more complicated logic needed to resolve exception rows to the general-case code path. 2.6% of input rows violate the normal-case and get handled by the general-case code path in Tuplex. Spark and Dask handle null values inline in UDF execution, and we use PySparkSQL, which compiles the query plan (though not the UDFs) into JVM bytecode. Figure 4 shows the results for two years’ worth of data (5.9 GB) and ten years (30.4 GB). PySparkSQL outperforms Dask by 2.3–2.7× because of its compiled query plan and more efficient join operator. Tuplex, despite its unoptimized join operator, still achieves a 7.6–9× speedup over PySparkSQL (17.4–24.6× over Dask) because it compiles and merges the UDFs, and processes the bulk of the data through a single, end-to-end optimized stage (we break this down in §6.4.2).

6.1.3 Log Processing: Regex and Randomness. We use the weblogs pipeline to investigate how Tuplex’s compiled code compares to special-purpose operators designed to accelerate common UDF functionality in existing frameworks. The pipeline splits an input log line into columns, and then rewrites one of those columns with a random string if it matches a username pattern:

Figure 5: Tuplex outperforms Spark and Dask by 5.1–91× on the weblogs pipeline; all Tuplex variants perform similarly. PySparkSQL only supports per-column regexes.

```
def randomize_udf(x):
    r = [random_choice(LETTERS) for t in range(10)]
    return re_sub('^/~[^/]+', '/~' + ''.join(r), x)
```

We consider three settings for the log line splitting operation:

- (1) natural Python using string operations (strip/split);
- (2) per-column regular expressions; and
- (3) a single regular expression.

Natural Python requires UDFs in all systems, but we also wrote an equivalent query using SparkSQL’s native string functions (*i.e.*, the query executes entirely within the JVM). PySparkSQL also has a native operator for regular expressions (regexp_extract). It only supports per-column regular expressions (second setting), but the operator applies the regular expression in the JVM, rather than in Python. Finally, all systems currently require UDFs when using a *single* regular expression. Tuplex supports all three approaches.

We would expect Python UDFs (both strip/split and regex-based) in Spark and Dask to be slowest. PySparkSQL’s native regex operator and the split-like SQL query should outperform them. A good result for Tuplex would show performance improvements in all three setups, as Tuplex end-to-end compiles and optimizes each setting for this pipeline. The input in our experiment is 75.6 GB of logs (715M rows). For Dask, we excluded 31.7M rows (4.5%, 4 GB) of the data because they triggered a known bug in the inner join [64].

Figure 5 reports the results organized by setting. The PySpark pipelines with two UDFs are slowest at about 80 minutes, while Dask UDFs are roughly 4× faster (18 min). Dask is more efficient because it executes the entire pipeline in Python, avoiding costly back-and-forth serialization between the JVM and Python workers. However, when PySparkSQL keeps the log line splitting in the JVM—either using string functions (PySparkSQL (split)) or via per-column regexes—runtime reduces to about 12 minutes. This happens because SparkSQL can generate JVM bytecode for most of the pipeline (except the randomization UDF) via its whole-stage code generation [69]. Tuplex, on the other hand, completes the pipeline in one minute both using natural Python and with a regular expression. Per-column regular expressions slow Tuplex down by a factor of two, but it still outperforms PySparkSQL by 5.1×; likewise, Tuplex’s split-based pipeline is 10.6× faster than PySparkSQL’s equivalent native SQL query. This difference comes, in part, because Tuplex compiles both UDFs to native code, while PySpark can only use compiled code for line-splitting. When we subtract the anonymization UDF runtime in both systems, Tuplex is still about 8× faster than PySparkSQL. The remaining speedup comes from Tuplex’s end-to-end optimization, and from using PCRE2 regular expressions: in our microbenchmarks, PCRE2 is 8.85× faster than java.util.regex, which Spark uses.

Tuplex’s fastest pipelines (single regex, strip) outperform the best PySpark and Dask setups by 13× and 19.7×. Tuplex supports

Figure 6: Tuplex’s exception resolution adds little overhead (0.3%) with compiled resolvers for Z2 on “dirty” data (25% of rows malformed). The interpreter (*) bar is a lower bound that assumes ideal 16× speedup over single-threaded interpreter. logical optimizations unavailable to Dask and Spark that improve performance further, which we discuss in §6.4.1.

6.2 Exception Handling

Tuplex speeds up processing of common-case rows by deferring exception-case rows to slower code paths. Exception rows arise either because of malformed (“dirty”) input data, or because they don’t match the common case in Tuplex’s sample (*e.g.*, due to type mismatches). We now measure the cost of exception row handling.

We run Z2 on the original, uncleaned Zillow dataset (scaled to 10 GB). 25% of the 56M input rows are exception rows with malformed data. We compare three setups: (i) ignoring and discarding all exception rows; (ii) the developer manually resolving exceptions in UDFs; (iii) using Tuplex’s resolvers (§3), both with compiled resolvers (on the general path) and resolution in the interpreter (fallback path). Our prototype runs a single-threaded Python interpreter for the fallback path, but this could be parallelized, so we assume a hypothetical, ideal 16× speedup to obtain a lower bound on fallback path runtime. Ignoring exception rows should be the fastest, while a good result for Tuplex would show manual and automatic resolution achieve similar performance, and a low overhead for handling exception rows. Figure 6 shows a breakdown of Tuplex’s execution time in each case. Ignoring all exceptions is fastest, since it merely skips the rows. Manual resolution adds an 8% overhead, but requires laborious changes and makes the UDFs much more complex. Tuplex’s compiled resolvers come within 0.3% of the hand-crafted resolution, with the overhead owed to increased LLVM compilation time. When we force all resolution onto the fallback path, however, it adds a 13× overhead, as 25% of rows are now processed in the Python interpreter. This shows that Tuplex’s compiled general-case path is crucial for good performance on exception-heavy workloads.

Processing a single row on the normal path takes 0.8μs. The compiled general path takes 0.28μs per row on average, as most exception rows are discarded early. To measure the full cost, we replaced all exception rows with synthetic data that proceeds all the way through the pipeline; in this case, the compiled general path takes 1.3μs (vs. 299μs/row in the interpreter).

6.3 Comparison To Other Systems

We now compare Tuplex to systems that generate and (JIT-)compile efficient native code for Z1. Z2 yields similar results (omitted).

6.3.1 Python Compilers. We first compare Tuplex to general Python compilers, which compile arbitrary Python programs.

PyPy. PyPy [57] is a tracing JIT compiler that can serve as a drop-in replacement for the CPython interpreter. It detects hot code paths

Figure 7: The PyPy3 general-purpose JIT fails to accelerate the Z1 query, and degrades performance by up to 3×. Dark bars use PyPy, light bars use the CPython interpreter (Fig. 3).

System		Runtime	Compile time
CPython (interpreter)		233.1 s	–
Python compilers	Cython	195.3 s	6.5 s
	Nuitka	192.5 s	9.4 s
	Tuplex	32.3 s	0.2 s
Hand-optimized C++		19.2 s	7.5 s

Table 3: Tuplex runs the Z1 query 6× faster than Cython and Nuitka, and compiles 32–47× faster than alternatives.

(usually loops), JIT-compiles a specialized interpreter and caches the hot paths’ native code. We configured Pandas, PySpark and Dask to use PyPy (v7.3.3) instead of CPython to measure how well PyPy performs on UDFs, and run the Zillow pipeline in the same setups as before. Even though PyPy is still bound to Python’s object representation and has limited scope for end-to-end optimization, the hope is that JIT-compiling the hot code paths will improve performance.

Figure 7 shows that this is actually not the case. PyPy is slower than interpreted Python in all settings, by between 3% and 3.18×; only with PySparkSQL it comes close to interpreted Python. Profiling with cProfile [16] suggests that PyPy has a variable impact on UDF performance: of twelve UDFs, seven are faster (13%–11.6×) with PyPy, and five are 26%–2.8× slower. The one UDF that benefits substantially (11.6×) merely forms a tuple; for others, even superficially similar string-processing UDFs exhibit varying performance. We attribute this to PyPy JIT-compiling and caching only some code paths, but not others. The 3× slowdown for Pandas and Dask is due to PyPy3’s poor performance when invoking C extension modules [63]. Tuplex is 14–24× faster.

Cython and Nuitka. Nuitka and Cython emit C/C++ files that contain unrolled calls to C functions which power the CPython interpreter. Compiling this file into a shared library object produces a drop-in replacement for a Python module. We used Nuitka (v0.6.13) and Cython (v0.29.22) to transpile the Python module to C for Z1 and compile it with gcc v10.2. This eliminates the cost of Python byte code translation and allows the C compiler to optimize the whole pipeline. We run the resulting module over 10 GB of input data, and compare single-threaded runtime to interpreted CPython and Tuplex.

Table 3 shows runtimes and compile times. Nuitka and Cython’s compiled code runs 17% faster than interpreted Python, but is still over 6× slower than Tuplex. Tuplex outperforms Nuitka and Cython because it replaces C-API calls with native code, eliminates dispensable checks and uses a more efficient object representation than Cython and Nuitka, which use CPython’s representation. Cython and Nuitka also have 32–47× higher compile times than Tuplex. They

Figure 8: For the 311 data cleaning pipeline, single-threaded Tuplex comes within 2× of Weld and outperforms all parallel systems. Tuplex outperforms Weld by 2× end-to-end because Tuplex inlines the aggregation in its generated parser.

take about a second to generate code, with the rest of the compile time taken up by the C compiler (gcc v10.2). Tuplex generates LLVM IR, which is faster to compile than higher-level C/C++, and also compiles 37× faster than gcc compiles the C++ baseline.

6.3.2 Data-parallel IR: Weld [52]. Weld is a data-parallel IR that admits optimizations like vectorization or loop fusion across libraries [52]. Weld serves as a backend to ported existing libraries such as Pandas [38], while Tuplex is a complete data analytics system, but both execute compiled native code. We compare Tuplex’s performance to Weld’s on the 311 data cleaning workload [15] and TPC-H Q6 and Q19. Q6 and Q19 perform simple filters and aggregations and are a challenging workload for Tuplex, which shines at string-heavy workloads with row-level UDFs and does not yet support vectorized (SIMD) compilation of UDFs. We compare to Weld v0.4.0; since Weld’s multithreaded runtime was removed in v0.3.0 [51], we compare single-threaded performance. In addition, we preprocessed the inputs to contain only the required columns and converted all dates to integers. In the single-threaded case all string-typed columns in Q6 and Q19 were transformed to integers, because Weld lacks automatic projection pushdown and has limited string processing capabilities. Because Weld does not have a native CSV parser, we preload the Q6/Q19 data into its columnar in-memory format with a single-threaded C++ CSV parser [67]. For the 311 workload, we use Weld’s benchmark code, which uses Pandas to load the data. We measure pure compute time, which measures how good Tuplex’s generated code is, and end-to-end runtime, which measures a realistic data analytics experience. A good result for Tuplex would show competitive compute time and an improved end-to-end runtime.

Figure 8 shows that Tuplex’s compute time (including compilation and sampling) for the 311 data cleaning workload is within 2× of Weld’s, and that end-to-end (total runtime to load the data, compile the query, and execute it), Tuplex runs the workload 2× faster than Pandas+Weld. On TPC-H Q6, Tuplex’s runtime is within 2× of Weld’s for Q6, despite Tuplex’s lack of vectorization and its row-structured data layout in memory (Figure 9a), and Tuplex again outperforms Weld by 1.86× end-to-end (Figure 9b). Tuplex’s end-to-end performance gains come from an optimization available when compiling full pipelines: instead of loading the data first and then running the aggregation, Tuplex *generates* a CSV parser and inlines the aggregation code into it. Weld, by contrast, first loads the data via Pandas to a columnar in-memory representation and then aggregates it via fast SIMD instructions. The results for Q19 are similar: due to vectorization, Weld outperforms Tuplex (without logical optimizations, 10.9s) by 2×. However, Tuplex can apply logical optimizations and

Figure 9: For TPC-H Q6/19, Tuplex’s generated code (without vectorization or indexes) is competitive with Weld’s vectorized code and within 3–8× of Hyper’s index-based execution. End-to-end, Tuplex outperforms Weld by 2× (due to its generated parser) and Hyper by 5–7× (by avoiding index creation). push down filters as Tuplex’s optimizer is aware of both UDFs and the overall query structure. This awareness leads to a 3× speedup over Weld, even though Tuplex lacks vectorization.

6.3.3 SQL query compiler: Hyper [27]. Tuplex is designed for analytics over large, non-indexed data sets. In classic SQL databases, query compilation is well-established. While Tuplex seeks to support a broader use case (Python UDFs) than SQL queries, we compare to the Hyper query compiler [27, 47] to establish a baseline for Tuplex’s performance on classic SQL queries. We use Tableau’s latest HyperAPI [35] (0.0.12366) to run TPC-H Q6 with 16 threads. Hyper relies on indexes for performance [44]: we expect Q6 to run an order of magnitude faster when indexes are used, as they allow to skip most of the data compared to a pure scan-based version. This comes at the upfront cost of creating the indexes, however.

Tuplex’s scan-based query execution is indeed 3–8× slower than Hyper’s index-based execution (Figure 9a). Tuplex’s Python code is also more expensive to compile (120ms) than directly parsing a simple, well-structured SQL query like Q6, as Tuplex must perform additional steps like type inference and tracing. Finally, Figure 9b shows that Tuplex outperforms Hyper by 5–7× on end-to-end runtime, since Tuplex avoids upfront index creation and interleaves the aggregation with data loading through its generated parser.

6.3.4 Discussion. Tuplex by design cannot use some optimizations available to Weld or Hyper, because Tuplex adheres strictly to Python semantics and must forego optimizations that would violate these semantics (e.g., via `-ffast-math`). Further, Tuplex generates code that still contains instructions to check for exceptions, while Weld and Hyper only work on perfectly clean data.

6.4 Tuplex Performance Breakdown

The largest contributor to Tuplex’s speedup over Spark and Dask is compiling Python UDFs to native code, but specific design choices improve Tuplex’s performance by up to 3×.

We measure the impact of specific design choices and optimizations with the flights pipeline, using 4-way concurrency and with Tuplex configured to avoid swapping. Figure 10 summarizes the impact of each factor on **flights** (5.9 GB input data) with and without LLVM optimizers enabled, plotting times only for the compute part of the pipeline (i.e., excluding I/O). There are two high-level take-aways: first, logical optimizations and stage fusion are important; and second, our optimizations give additional purchase to the LLVM

Figure 10: Factor analysis for the flights pipeline: Tuplex optimization and LLVM optimizers together realize speedups.

optimizers. We mention results for other pipelines where relevant; these are end-to-end numbers including I/O.

6.4.1 Logical Optimizations. Tuplex compiles Python UDFs with full knowledge of their ASTs. This allows Tuplex to apply standard optimizations like filter and projection pushdowns and operator reorderings *through* UDFs—in contrast to Spark or Dask, which treat UDFs as black-boxes. We illustrate the impact such logical optimizations have with the weblogs and flight pipelines; the Zillow pipeline has few logical optimization opportunities.

In the **flights** pipeline, projection pushdown helps drop many of the 110 input columns early. Tuplex achieves a 2× speedup thanks to this logical optimization when we disable LLVM optimizers, but the benefit grows to 3× with LLVM optimizers enabled. This is caused by LLVM eliminating code that processes data eventually dropped and its ability to reorder basic blocks for inlined functions.

The **weblogs** pipeline contains a join with a list of malicious IPs and a mapColumn operator that anonymizes some records. Applying mapColumn to output rows of the (selective, i.e., filtering) join requires anonymizing fewer rows. But Spark or Dask cannot move a UDF-applying mapColumn through a join, while Tuplex can, thanks to its understanding of columns read and modified in the UDF. With this optimization, Tuplex takes 27 seconds (2× faster than the unordered result we reported in Figure 5). If we manually reorder the operators in PySparkSQL, it also runs 2.5× faster (305 seconds), but remains 11.3× slower than Tuplex.

6.4.2 Stage Fusion. Systems that treat UDFs as black-box operators are unable to end-to-end optimize across them. In Spark and Dask, a UDF operator is an *optimization barrier*, while Tuplex makes stages—the unit of optimized compilation—as large as possible. To measure the impact of this design, we manually insert optimization barriers in the flights pipeline, forcing Tuplex to use additional stages. We consider Tuplex with optimization barriers that mimic Spark’s optimization constraints; and Tuplex with stage fusion (i.e., only the build side of a join is a barrier, cf. §4.5). For each, we disable and enable LLVM optimizers to measure any cross-UDF optimization enabled. Without LLVM optimizers, Tuplex takes 27.2 seconds without stage fusion and 21.1 seconds with stage fusion (22% improvement); with LLVM optimizers, runtimes drop to 17.1 and 9.9 seconds (42% improvement). Stage fusion thus enables optimization potential that improves runtime by an extra 20%.

6.4.3 Optional Types off the Normal Path. Dual mode processing allows Tuplex to optimize the normal-case path by deferring complexity to the exception-case path. We measure the impact of shifting rare null values to the general-case code path (§4.7). In **flights**, this optimization reduces the pipeline’s compute time by

Setup	Spark (64 executors)	Tuplex (64 Lambdas)
100 GB	209.03 sec ($\sigma = 10.53$)	31.5 sec ($\sigma = 8.25$)
1 TB	1791.51 sec ($\sigma = 4.38$)	351.1 sec ($\sigma = 22.10$)

Table 4: In a distributed scale-out experiment, Tuplex’s Lambda backend outperforms a Spark cluster by 5.1–6.6 $\times$.

14–19%, albeit at the cost of increasing compile time by 2 seconds, which reduces end-to-end benefit. (Larger datasets would realize more of the benefit, as compile time is a constant.)

6.5 Distributed, Scale-Out Execution

While our focus has been on the single-machine performance of our Tuplex prototype, some systems we compare to (PySpark and Dask) support distributed execution. To verify that Tuplex’s performance gains are not merely a consequence of avoiding overheads associated with distributed operation, we compare these systems with Tuplex’s experimental distributed execution over AWS Lambda functions.

We compare our prototype’s Lambda backend with a maximum concurrency of 64 simultaneously running requests to a Spark cluster with 64 executors. We use Lambdas with 1.5 GB of memory. The Spark cluster runs on 32 m5.large instances that each run two executors with 1 core and 2 GB of memory per executor. This gives Spark an advantage, as it has more memory and the cluster runs continuously, while Tuplex provisions a Lambda container for each task. In addition, while Tuplex’s Lambda backend writes to S3, Spark merely collects results on its driver node, as writing to S3 requires extra infrastructure [23, 65]. We run the Zillow pipeline over scaled datasets of 100 GB and 1 TB, with data stored in 256 MB chunks in AWS S3. To verify that the compute speed of m5.large VMs is comparable to 1.5 GB Lambda functions, we ran a microbenchmark over one 256MB chunk. It takes 3.99 seconds on an m5.large VM, while our code within a Lambda function takes 4.00 seconds on average, with some variance (min: 3.68 sec, max 9.99 sec).

Table 4 shows the results. For Spark, we show numbers for the tuple-based pipeline; the dictionary and SparkSQL versions are 10–20% slower. Tuplex completes the pipeline in 31.5 and 351 seconds for 100 GB and 1 TB, 5.1 $\times$ and 6.6 $\times$ faster, respectively, than the fastest Spark setup. This difference comes from Tuplex’s compute speed, which outweighs the overheads associated with Lambdas (HTTP request, queueing, container provisioning, etc.). In terms of direct monetary cost, Tuplex is competitive at 4¢ for 100 GB (Spark: 3.7¢) and 55¢ for 1 TB (Spark: 32¢), while also avoiding the setup and provisioning time costs, idle costs, and EBS storage costs that Spark incurs on top of the EC2 VM costs. This suggests that Tuplex can be competitive both on a single server and in scale-out settings.

7 Discussion and Experience

Tuplex’s primary objective is high performance for pipelines that include Python UDFs. But the dual-mode execution model may also help Tuplex users avoid some long-standing challenges of pipeline development and debugging [26, 55]. Key to this is Tuplex’s guarantee that pipelines never fail because of malformed input rows: instead, Tuplex does its best to complete the pipeline on valid, normal-case rows and reports statistics about failed rows to the user. It is difficult to quantify the impact of failure-free pipelines on developer productivity. However, in our anecdotal experience implementing pipelines we found Tuplex preferable for several reasons:

- (1) Although our evaluation data sets are fairly “clean”, they contain a small number of anomalous rows, which often caused hard-to-debug failures in Spark and Dask.
- (2) Representing rows as tuples instead of dictionaries improves PySpark performance, but the numerical indexing took painstaking work to get right. Tuplex avoids the speed-usability tradeoffs and has the same performance for tuples and dictionaries.
- (3) Making null values work with Dask/Pandas required using special datatypes (e.g., `np.int64`), rather than native Python types, as Pandas fails on None values.
- (4) The semantics of special-purpose operators designed to help developers avoid UDFs differ from Python code. For example, SparkSQL’s `regex_extract` returns an empty string when there is no match, rather than NULL as a Python user might expect (Python’s `re` returns None in this case). The weblog dataset has two anomalous rows, which cause SparkSQL to silently return incorrect results, while Tuplex correctly reported them.
- (5) We compared to Weld using the Pandas cookbook’s sub-sampled 311 dataset [15] (99k rows) scaled 2,000 $\times$ in §6.3.2, but Tuplex works out-of-the-box on the full NYC 311 dataset [49] (22.4M rows), while Weld, PySpark, PySparkSQL, and Dask all fail and require changes to the UDF code for the realistic dataset.

We spent substantial time tracking down edge cases in framework documentation for other systems, while Tuplex’s Python UDFs behaved as expected. We also found that Tuplex’s reporting of exceptions and failed rows helped us track down bugs in our pipelines.

Tuplex’s dual mode processing requires a representative sample. Like with any sampling approach, an unrepresentative sample can lead Tuplex to deduce an incorrect common case. If the sample itself produces only exceptions, Tuplex warns the user either to revise the pipeline or increase the sample size.

8 Conclusion

Tuplex is a new data analytics system that compiles Python UDFs to optimized, native code. Tuplex’s key idea is dual-mode processing, which makes optimizing UDF compilation tractable. Tuplex is available as open-source software at <https://tuplex.cs.brown.edu>.

Acknowledgments

We are grateful to Shriram Krishnamurthi and the Brown Systems and Database groups for feedback on drafts of this paper. We also thank Colby Anderson, Khosrow Arian, Ben Givertz, Andy Ly, Raghu Nimmagadda, William Riley, Yunzhi Shao, and Andrew Wei for their contributions to Tuplex’s implementation. This research was supported by NSF awards DGE-2039354 and IIS-1453171, funding from Google, and a Brown University Richard B. Salomon Award. This research was also sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator under Cooperative Agreement FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Yanif Ahmad and Christoph Koch. 2009. DBToaster: A SQL Compiler for High-Performance Delta Processing in Main-Memory Databases. *Proceedings of the VLDB Endowment* 2, 2 (Aug. 2009), 1566–1569. <https://doi.org/10.14778/1687553.1687592>
- [2] Anaconda, Inc. 2020. Will Numba Work For My Code? (2020). <http://numba.pydata.org/numba-doc/latest/user/5minguide.html#will-numba-work-for-my-code>
- [3] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, et al. 2015. Spark SQL: Relational Data Processing in Spark. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. 1383–1394.
- [4] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D.S. Seljebotn, and K. Smith. 2011. Cython: The Best of Both Worlds. *Computing in Science Engineering* 13, 2 (3 2011), 31–39. <https://doi.org/10.1109/MCSE.2010.118>
- [5] United States Department of Transportation Bureau of Transportation Statistics. 2020. Carrier history lookup table. (2020). https://www.transtats.bts.gov/Download_Lookup.asp?Lookup=L_CARRIER_HISTORY
- [6] United States Department of Transportation Bureau of Transportation Statistics. 2020. Reporting Carrier On-Time Performance (1987-present). (2020). https://www.transtats.bts.gov/Fields.asp?Table_ID=236
- [7] C. Chambers and D. Ungar. 1989. Customization: Optimizing Compiler Technology for SELF, a Dynamically-Typed Object-Oriented Programming Language. In *Proceedings of the ACM SIGPLAN 1989 Conference on Programming Language Design and Implementation*. 146–160. <https://doi.org/10.1145/73141.74831>
- [8] Andrew Crotty, Alex Galakatos, Kayhan Dursun, Tim Kraska, Carsten Binnig, Ugur Cetintemel, and Stan Zdonik. 2015. An Architecture for Compiling UDF-centric Workflows. *Proceedings of the VLDB Endowment* 8, 12 (2015), 1466–1477.
- [9] Databricks. 2019. W3LI: Apache Logs Lab Dataframes (Python). (2019). <https://databricks-prod-cloudfront.cloud.databricks.com/public/4027ec902e339c93eaa8714f173bfcf/2799933550853697/4438435960036599/2202577924924539/latest.html>
- [10] Databricks, Inc. 2021. Handling bad records and files. Databricks for SQL Developers Documentation. (21 1 2021). <https://docs.databricks.com/spark/latest/spark-sql/handling-bad-records.html>
- [11] James M Decker, Dan Moldovan, Andrew A Johnson, Guannan Wei, Vritant Bhardwaj, Gregory Essertel, and Fei Wang. 2019. Snek: Overloading Python Semantics via Virtualization. (7 2019). <https://www.cs.purdue.edu/homes/rompf/papers/essertel-preprint201907.pdf>
- [12] G. Dot, A. Martinez, and A. González. 2015. Analysis and Optimization of Engines for Dynamically Typed Languages. In *2015 27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. 41–48. <https://doi.org/10.1109/SBAC-PAD.2015.20>
- [13] Gilles Duboscq, Thomas Würthinger, Lukas Stadler, Christian Wimmer, Doug Simon, and Hanspeter Mössenböck. 2013. An intermediate representation for speculative optimizations in a dynamic compiler. In *Proceedings of the 7th ACM workshop on Virtual Machines and Intermediate Languages*. 1–10.
- [14] Gregory Essertel, Ruby Tahboub, James Decker, Kevin Brown, Kunle Olukotun, and Tiark Rompf. 2018. Flare: Optimizing apache spark with native compilation for scale-up architectures and medium-size data. In *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 799–815.
- [15] Julia Evans. 2020. Pandas Cookbook: Chapter 7 – Cleaning up messy data. (6 8 2020). <https://github.com/jvns/pandas-cookbook/blob/master/cookbook/Chapter%207%20-%20Cleaning%20up%20messy%20data.ipynb>
- [16] Python Software Foundation. 2021. The Python Profilers: cProfile. (12 4 2021). <https://docs.python.org/3/library/profile.html#module-cProfile>
- [17] Lars Grammel. 2015. Wrangling US Flight Data - Part 1. (12 5 2015). <https://www.trifacta.com/blog/wrangling-us-flight-data-part-1/>
- [18] Kay Hayen. 2018. Nuitka. (2018). <http://nuitka.net/>
- [19] Anna Herlihy. 2016. PYLLVM: A compiler from a subset of Python to LLVM-IR. (May 2016). <https://pycon.org.il/2016/static/sessions/anna-herlihy.pdf>
- [20] Fabian Hueske, Mathias Peters, Matthias J Sax, Astrid Rheinländer, Rico Bergmann, Aljoscha Krettek, and Kostas Tzoumas. 2012. Opening the black boxes in data flow optimization. *Proceedings of the VLDB Endowment* 5, 11 (2012), 1256–1267.
- [21] Eunji Jeong, Sungwoo Cho, Gyeong-In Yu, Joo Seong Jeong, Dong-Jin Shin, and Byung-Gon Chun. 2019. JANUS: Fast and Flexible Deep Learning via Symbolic Graph Execution of Imperative Programs. In *Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation*. 453–468.
- [22] Eunji Jeong, Sungwoo Cho, Gyeong-In Yu, Joo Seong Jeong, Dong-Jin Shin, Taebum Kim, and Byung-Gon Chun. 2019. Speculative Symbolic Graph Execution of Imperative Deep Learning Programs. *SIGOPS Oper. Syst. Rev.* 53, 1 (July 2019), 26–33. <https://doi.org/10.1145/3352020.3352025>
- [23] Cy Jervis. 2017. Introducing S3Guard: S3 Consistency for Apache Hadoop. Cloudera Blog. (03 8 2017). <https://blog.cloudera.com/introducing-s3guard-s3-consistency-for-apache-hadoop/>
- [24] Li Jin. 2017. Introducing Pandas UDF for PySpark - The Databricks Blog. [https://databricks.com/blog/2017/10/30/introducing-vectorized-udfs-for-](https://databricks.com/blog/2017/10/30/introducing-vectorized-udfs-for-pyspark.html)
- [25] pyspark.html. (10 2017).
- [25] Johannes Kepler University (JKU). 2019. The Truffle Language Implementation Framework. (2019). <http://www.ssw.uni-linz.ac.at/Research/Projects/JVM/Truffle.html>
- [26] Sital Kedia, Shuojie Wang, and Avery Ching. 2016. Apache Spark @Scale: A 60 TB+ production use case. (31 8 2016). <https://code.fb.com/core-data/apache-spark-scale-a-60-tb-production-use-case/>
- [27] Alfons Kemper and Thomas Neumann. 2011. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *Proceedings of the 27th IEEE International Conference on Data Engineering (ICDE)*. 195–206.
- [28] Timo Kersten, Viktor Leis, Alfons Kemper, Thomas Neumann, Andrew Pavlo, and Peter Boncz. 2018. Everything you always wanted to know about compiled and vectorized queries but were afraid to ask. *Proceedings of the VLDB Endowment* 11, 13 (2018), 2209–2222.
- [29] Yannis Klonatos, Christoph Koch, Tiark Rompf, and Hassan Chafi. 2014. Building Efficient Query Engines in a High-Level Language. *Proceedings of the VLDB Endowment* 7, 10 (June 2014), 853–864. <https://doi.org/10.14778/2732951.2732959>
- [30] Sanjay Krishnan, Michael J. Franklin, Kenneth Goldberg, and Eugene Wu. 2017. BoostClean: Automated Error Detection and Repair for Machine Learning. (11 2017). [arXiv:cs.DB/1711.01299](https://arxiv.org/abs/1711.01299)
- [31] Mads R.B. Kristensen, Simon A.F. Lund, Troels Blum, and James Avery. 2016. Fusion of Parallel Array Operations. In *Proceedings of the 2016 International Conference on Parallel Architectures and Compilation*. 71–85. <https://doi.org/10.1145/2967938.2967945>
- [32] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. 2015. Numba: A LLVM-based Python JIT Compiler. In *Proceedings of the 2nd Workshop on the LLVM Compiler Infrastructure in HPC*. Article 7, 6 pages. <https://doi.org/10.1145/2833157.2833162>
- [33] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A Compiler Infrastructure for the End of Moore's Law. (2020). [arXiv:cs.PL/2002.11054](https://arxiv.org/abs/2002.11054)
- [34] Chris Leary and Todd Wang. 2017. XLA: TensorFlow, compiled. TensorFlow Dev Summit. (2017). <https://youtu.be/kA0anJczHA0>
- [35] Tableau Software LLC. 2021. Tableau Hyper API. (07 4 2021). https://help.tableau.com/current/api/hyper_api/en-us/index.html
- [36] LLVM Project. 2019. Itanium ABI Zero-cost Exception Handling. (Aug. 2019). <https://releases.llvm.org/8.0.1/docs/ExceptionHandling.html#itanium-abi-zero-cost-exception-handling>
- [37] LLVM Project. 2019. Setjmp/Longjmp Exception Handling. (Aug. 2019). <https://releases.llvm.org/8.0.1/docs/ExceptionHandling.html#setjmp-longjmp-exception-handling>
- [38] Wes McKinney et al. 2011. pandas: a foundational Python library for data analysis and statistics. *Python for High Performance and Scientific Computing* 14, 9 (2011).
- [39] Micro Focus International, PLC. 2020. Vertica 9.2.x: Capturing Load Rejections and Exceptions. Vertica Analytics Platform Version 9.2.x: Documentation. (2020). <https://www.vertica.com/docs/9.2.x/HTML/Content/Authoring/AdministratorsGuide/BulkLoadCOPY/CapturingLoadExceptionsAndRejections.htm>
- [40] Kevin Modzelewski. 2017. Pyston 0.6.1 released, and future plans. (31 1 2017). <https://blog.pyston.org/2017/01/31/pyston-0-6-1-released-and-future-plans/>
- [41] Kevin Modzelewski. 2018. Introducing Pyston: an upcoming, JIT-based Python implementation. (2018). <https://blogs.dropbox.com/tech/2014/04/introducing-pyston-an-upcoming-jit-based-python-implementation/>
- [42] Kevin Modzelewski. 2020. Pyston v2: 20% faster Python. (28 10 2020). <https://blog.pyston.org/2020/10/28/pyston-v2-20-faster-python/>
- [43] Yann Moisan. 2018. Spark performance tuning from the trenches. (29 5 2018). <https://medium.com/teads-engineering/spark-performance-tuning-from-the-trenches-7cbde521cf60>
- [44] Tobias Mühlbauer, Wolf Rödiger, Robert Seilbeck, Angelika Reiser, Alfons Kemper, and Thomas Neumann. 2013. Instant loading for main memory databases. *Proceedings of the VLDB Endowment* 6, 14 (2013), 1702–1713.
- [45] Myip.ms. 2020. Blacklist IP Addresses Live Database (Real-time). (2020). <https://myip.ms/browse/blacklist>
- [46] Christian Navasca, Cheng Cai, Khanh Nguyen, Brian Demsky, Shan Lu, Miryung Kim, and Guoqing Harry Xu. 2019. Gerenuk: thin computation over big native data using speculative program transformation. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP)*. 538–553.
- [47] Thomas Neumann. 2011. Efficiently compiling efficient query plans for modern hardware. *Proceedings of the VLDB Endowment* 4, 9 (2011), 539–550.
- [48] Khanh Nguyen, Lu Fang, Christian Navasca, Guoqing Xu, Brian Demsky, and Shan Lu. 2018. Skyway: Connecting managed heaps in distributed big data systems. *ACM SIGPLAN Notices* 53, 2 (2018), 56–69.
- [49] NYC OpenData. 2020. 311 Service Requests from 2010 to Present. (2020). <https://data.cityofnewyork.us/Social-Services/311-Service-Requests-from-2010-to-Present/erm2-nwe9>
- [50] Oracle, Inc. 2021. GraalVM. (2021). <https://www.graalvm.org/>

- [51] Shoumik Palkar. 2018. PR #308: Remove Multithreaded Backend. (23 8 2018). <https://github.com/weld-project/weld/pull/383>
- [52] Shoumik Palkar, James Thomas, Deepak Narayanan, Pratiksha Thaker, Rahul Palamuttam, Parimajan Negi, Anil Shanbhag, Malte Schwarzkopf, Holger Pirk, Saman Amarasinghe, Samuel Madden, and Matei Zaharia. 2018. Evaluating End-to-end Optimization for Data Analytics Applications in Weld. *Proceedings of the VLDB Endowment* 11, 9 (May 2018), 1002–1015. <https://doi.org/10.14778/3213880.3213890>
- [53] Shoumik Palkar and Matei Zaharia. 2019. Optimizing data-intensive computations in existing libraries with split annotations. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP)*. 291–305.
- [54] Arash Partow. 2020. The Global Airport Database. (2020). <https://www.partow.net/miscellaneous/airportdatabase/>
- [55] Tejas Patil and Jing Zheng. 2017. Using Apache Spark for large-scale language model training. (7 2 2017). <https://code.fb.com/core-data/using-apache-spark-for-large-scale-language-model-training/>
- [56] PostgreSQL Global Development Group. 2015. Error logging in COPY. PostgreSQL Wiki. (12 5 2015). https://wiki.postgresql.org/wiki/Error_logging_in_COPY
- [57] Armin Rigo, Maciej Fijalkowski, Carl Friedrich Bolz, Antonio Cuni, Benjamin Peterson, Alex Gaynor, Hakan Ardoe, Holger Krekel, and Samuele Pedroni. 2021. PyPy. (2021). <http://pypy.org/>
- [58] Nadav Rotem, Jordan Fix, Saleem Abdulrasool, Garret Catron, Summer Deng, Roman Dzhabarov, Nick Gibson, James Hegeman, Meghan Lele, Roman Levenstein, et al. 2018. Glow: Graph lowering compiler techniques for neural networks. (2018). [arXiv:cs.PL/1805.00907](https://arxiv.org/abs/1805.00907)
- [59] ScrapeHero. 2017. How to Scrape Real Estate Listings from Zillow.com using Python and lxml. (4 Oct. 2017). <https://www.scrapehero.com/how-to-scrape-real-estate-listings-on-zillow-com-using-python-and-lxml/>
- [60] Amir Shaikhha, Yannis Klonatos, Lionel Parreaux, Lewis Brown, Mohammad Dashti, and Christoph Koch. 2016. How to Architect a Query Compiler. In *Proceedings of the 2016 International Conference on Management of Data*. 1907–1922. <https://doi.org/10.1145/2882903.2915244>
- [61] Juliusz Sompolski, Marcin Zukowski, and Peter Boncz. 2011. Vectorization vs. Compilation in Query Execution. In *Proceedings of the 7th International Workshop on Data Management on New Hardware*. Association for Computing Machinery, New York, NY, USA, 33–40. <https://doi.org/10.1145/1995441.1995446>
- [62] Ruby Y. Tahboub, Grégory M. Essertel, and Tiark Rompf. 2018. How to Architect a Query Compiler, Revisited. In *Proceedings of the 2018 International Conference on Management of Data*. 307–322. <https://doi.org/10.1145/3183713.3196893>
- [63] The PyPy Team. 2021. PyPy: Performance. (2021). <https://www.pypy.org/performance.html>
- [64] trstovall. 2019. [bug] dask.dataframe.DataFrame.merge fails for inner join. (28 3 2019). <https://github.com/dask/dask/issues/4643>
- [65] Gil Vernik, Michael Factor, Elliot K. Kolodner, Effi Ofer, Pietro Michiardi, and Francesco Pace. 2017. Stocator: An Object Store Aware Connector for Apache Spark. In *Proceedings of the 2017 Symposium on Cloud Computing*. 653. <https://doi.org/10.1145/3127479.3134761>
- [66] Jiannan Wang, Sanjay Krishnan, Michael J Franklin, Ken Goldberg, Tim Kraska, and Tova Milo. 2014. A sample-and-clean framework for fast and accurate query processing on dirty data. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 469–480.
- [67] David Wilson. 2019. csvmonkey - Header-only vectorized, lazy-decoding, zero-copy CSV file parser. (19 5 2019). <https://github.com/dw/csvmonkey>
- [68] Eugene Wu, Samuel Madden, and Michael Stonebraker. 2012. A demonstration of DBWipes: clean as you query. *Proceedings of the VLDB Endowment* 5, 12 (2012), 1894–1897.
- [69] Reynold Xin and Josh Rosen. 2015. Project Tungsten: Bringing Apache Spark Closer to Bare Metal. (28 4 2015). <https://databricks.com/blog/2015/04/28/project-tungsten-bringing-spark-closer-to-bare-metal.html>
- [70] Yuan Yu, Michael Isard, Dennis Fetterly, Mihai Budiu, Úlfar Erlingsson, Pradeep Kumar Gunda, and Jon Currey. 2008. DryadLINQ: A System for General-Purpose Distributed Data-Parallel Computing Using a High-Level Language. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*.
- [71] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation (NSDI)*.
- [72] M. Zukowski, M. van de Wiel, and P. Boncz. 2012. Vectorwise: A Vectorized Analytical DBMS. In *Proceedings of the 28th IEEE International Conference on Data Engineering*. 1349–1350.