ORACLE

Intern with Oracle Labs MLE!

Are you looking for an exciting internship opportunity? If so, this might be the perfect fit for you. The MLE team at Oracle Labs has an opening for an internship to help extend MLE with new exciting functionality!

Oracle

Oracle, a global provider of enterprise cloud computing, is empowering businesses of all sizes on their journey of digital transformation. Oracle Cloud provides leading-edge capabilities in software as a service, platform as a service, infrastructure as a service, and data as a service.

Oracle's application suites, platforms, and infrastructure leverage both the latest technologies and emerging ones – including artificial intelligence, machine learning, blockchain, and Internet of Things – in ways that create business differentiation and advantage for customers. Continued technological advances are always on the horizon.

Oracle Labs

Oracle Labs is the advanced research and development arm of Oracle. We focus on the development of technologies that keep Oracle at the forefront of the computer industry. Oracle Labs researchers look for novel approaches and methodologies, often taking on projects with high risk or uncertainty, or that are difficult to tackle within a product-development organization. Oracle Labs research is focused on real-world outcomes: our researchers aim to develop technologies that will someday play a significant role in the evolution of technology and society. For example, chip multithreading and the Java programming language grew out of work done in Oracle Labs.

Multilingual Engine (MLE)

The Multilingual Engine (MLE) [1] research project investigates how to leverage programming language runtimes in database systems (DBMS). Our hypothesis is that application development and data science can benefit from running code as close to the data as possible. For example, Python workloads for training machine learning models can run directly in the DBMS, using the DBMS as a compute cluster with efficient access to data. Similarly, the best place to run data-centric applications can be the database system itself, completely eliminating performance concerns due to network round trips and reducing infrastructure costs. The focus of our work is to enable Oracle Database to execute such workloads written in modern and popular languages and frameworks. The foundation for the project is GraalVM [2], Oracle Labs' high-performance, polyglot programming language runtime. A first outcome of our vision is the JavaScript support in Oracle Database 23c.

Additionally, we leverage Just-In-Time (JIT) compilation to improve the performance of database query processing. We explore making queries on relational tables and document collections faster using code generation and JIT compilation, all based on GraalVM and the Truffle framework.

Internships in the MLE project offer the opportunity to work with state-of-the-art technology at the crossroads of database systems and programming language runtimes. The MLE project conducts research with a strong focus on practical applicability.

Internship Details

We offer various topics depending on the candidate's skills and interests. Here are some of the projects that can be explored during the internship:

Efficient columnar data export for in-database data science

This internship aims to implement fast data export from the database engine into Apache Arrow, a standard columnar format for in-memory data. This will enable extremely efficient access to table data for Python ML pipelines running in the DBMS.

Reliable lock-free shared-memory data structures

The goal of this project is to provide users of MLE with fast concurrent data structures such as a key-value store, that can be shared across database sessions. Such shared-memory data structures will enrich the MLE programming model and unlock new use cases.

Compilers for tree ensemble inference in DB

Boosted tree models are the foundation for many practical ML applications. Compiling tree models to machine code can speed up inference. In this project, we will adapt an existing tree model compiler for in-DB inference by leveraging GraalVM.

Python parallel computing in RDBMS

Python uses a multiprocessing model to parallelize computation. We will explore how to map this model to the multiprocessing model of the Oracle Database for parallel data processing.

Accelerating JSON processing in the Database

This project will explore ways to accelerate JSON processing by leveraging Just-in-Time compilation. It includes generating specialized code for JSON encoding and decoding operations and compiling JSON path expressions.

Optimizing transactional workloads under a closed-world assumption

In this project we will perform program analysis to construct a graph of database accesses. Using this graph, we will optimize the execution of the program under the assumption that the graph contains all possible operations with the database. Optimizations that work under this assumption, such as transaction chopping, allow extracting more concurrency from the workload while maintaining the same level of isolation.

Image Caption 1. Our beautiful offices at Prime Tower in Zurich.

References

[1] https://docs.oracle.com/en/database/oracle/oracle-database/23/mlejs/introduction-to-mle.html

[2] https://www.graalvm.org/

Required Skills

The person stepping into this role is expected to complete the internship using a wide and diverse set of skills.

- Basic understanding of databases and SQL (having completed advanced database systems courses is a plus)
- Experience with C and Java programming. Experience with JavaScript and/or Python is a plus
- Experience with Linux (e.g., bash scripts)
- An average grade of at least 5.0 in the master studies is a plus

Internship Facts

The duration of the internship can vary based on the candidate's constraints. The usual duration is 6 months. We pay a competitive salary. The research topics listed are informative, we are open to suggestions depending on your skills and qualifications. By sending in your application, you opt-in for processing your personal information.

If this sounds like you, we hope to meet you!

How to Apply

After reviewing the Oracle Labs internship program, please provide us with the following information to apply:

- Your CV or link to your home page containing your curriculum
- Description of your motivation and area(s) of interest
- Availability and preferred internship duration
- Preferred location

If available, please apply to this topic via the internship portal of your university. Alternatively, send an email with the aforementioned information to labs-hiring_ww@oracle.com

"During my 3-month internship with the MLE team, I worked with state-of-the-art technology and conducted research whose (future) benefit to Oracle's products was very apparent for me. Whenever I faced any hurdles, I had great mentors and colleagues to help me out with their vast knowledge and experience. Overall, I had a great time working at Oracle Labs."

Tom Papke
TU Munich student, 3-months intern

About Us

Innovation starts with inclusion at Oracle. We are committed to creating a workplace where all kinds of people can be themselves and do their best work. It's when everyone's voice is heard and valued, that we are inspired to go beyond what's been done before. That's why we need people with diverse backgrounds, beliefs, and abilities to help us create the future, and are proud to be an affirmative-action equal opportunity employer.

Oracle is an Equal Employment Opportunity Employer. All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, national origin, sexual orientation, gender identity, disability and protected veterans status, age, or any other characteristic protected by law. Oracle will consider for employment qualified applicants with arrest and conviction records pursuant to applicable law.

Life at Oracle

An Oracle career can span industries, roles, countries and cultures, giving you the opportunity to tackle new roles and challenges, while blending work and life. Oracle has thrived through 40+ years of change by innovating and operating with integrity while delivering for the top companies in almost every industry. To nurture the talent that makes this happen, we work hard to build a vibrant and inspiring workplace that celebrates diverse, hardworking teams where everyone can contribute. We take care of each other, and value giving back to the community. We have flexible work arrangements and offer benefits including generous paid parental leave and comprehensive premium medical insurance.

We will ensure that individuals with disabilities are provided reasonable accommodation to participate in the job application or interview process, to perform crucial job functions, and to receive other benefits and privileges of employment. Please contact us to request accommodation.

At Oracle, we don't just respect differences—we celebrate them. We believe that innovation starts with inclusion and to create the future we need people with diverse backgrounds, perspectives, and abilities. That's why we're committed to creating a workplace where all kinds of people can do their best work. It's when everyone's voice is heard and valued that we're inspired to go beyond what's been done before.

More information on Oracle's stance on diversity and inclusion