

How I Learned to Stop Worrying and Love Compilers

Eric Sedlar
Oracle Labs

Abstract

The modern platforms that we want to use to manage our data are far more complex to program efficiently than the machines we used in the past. Every computer we run on is a massively parallel machine with many architectural “surprises” for programmers who are unaware of the way the underlying hardware architecture works. A simple way to think about this is that optimal programs must specify how & where the data should move, not just what computations should be performed and in what order. Architecture-oblivious software that leaves the decisions about data movement to a low-level coherence protocol is becoming much less efficient, relatively speaking.

After an extended flirtation with using imperative programming frameworks such as Map-Reduce and NoSQL, many people are returning back to declarative languages like SQL, where the language compiler & runtime are free to make most of the data movement decisions for the programmer. Another way to think about a SQL compiler is that it includes an “algorithm picker” and the runtime includes libraries of useful algorithm implementations (which contain the data movement specifications). This talk will discuss the needs and opportunities for expanding the domain of algorithm-picking languages like SQL. Doing so will require integration with managed-language runtime compilers (e.g. Java or Javascript compilers) that are integrated with the SQL compiler not just to provide efficiency gains during query execution, but also to use managed language runtime profiling to help in algorithm selection as well as assembly-level compilation decisions.

Categories and Subject Descriptors

H.2.3 Database Languages, D.3.4 Compilers

Keywords

Hardware-conscious software; SQL

Short Bio

Eric Sedlar Eric Sedlar is Vice President and Technical Director of Oracle Labs (formerly Sun Labs). This position entails figuring out how to transfer research results from Labs research into Oracle products & services, as well as setting overall technical direction for new research projects in Oracle Labs. Eric manages over 160 fulltime researchers working on over a dozen separate research areas.



Eric's own research interests are in domain-specific languages and acceleration of database operations both via new hardware and using JIT compilation.

Previously, he led the effort for XML-native storage inside Oracle, starting with Oracle 9iR2. Eric has held various architecture and development management positions at Oracle since starting there in 1990. He holds over 68 patents, and has served on standards organizations for Oracle in the W3C and IETF. He co-authored the Best Paper at SIGMOD 2010 on architecture-sensitive search trees.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

SIGMOD'14, June 22–27, 2014, Snowbird, UT, USA.

ACM 978-1-4503-2376-5/14/06.

<http://dx.doi.org/10.1145/2588555.2602133>