Synthesizing Database Engines

Tim Gubner (CWI) and Peter Boncz (CWI)

November 13, 2019

So far a few hundred database systems have been created by hand. Implementing them has been an extremely time-consuming and labor-intensive process. Consequently, risky choices are rarely taken and systems tend to evolve around best-practices. However, by avoiding risk it is impossible to explore the design space fully as radically different designs might oppose known best-practices.

We argue that instead of practicing manual database architecture, it would be great to be able to synthesize them automatically. The ability to generate systems, does not only save significant amount of time, but also allows quick prototyping of radically new ideas. In the presentation, we explore the idea of a *Database Synthesizer*. We describe a common representation, VOILA, which is able to expresses algorithmic details while hiding specific execution details. We show that the state-of-the-art execution strategies, data-centric compilation and iterator-based vectorized execution, can be synthesized from VOILA. Last but not least, we highlight future steps and research challenges.