

Hardware Integrator for RDBMS and Machine Learning Accelerators

A hardware architecture to integrate hardware accelerators and databases

Inventors at Georgia Tech have developed a hardware architecture called Striders to provide a seamless integration between hardware accelerators for machine learning and relational database engines. Striders can extract, cleanse, and process the data stored in databases for the machine-learning accelerator to use. Further, the inventors devised a fixed length Instruction Set Architecture (ISA) for the Striders to target a range of RDBMS engines such as PostgreSQL and MySQL (InnoDB) that have similar backend page layouts.

Summary Bullets

- Faster and more energy efficient execution of learning algorithms
- Ensure compatibility between hardware accelerators and data in RDBMS

Solution Advantages

- Faster and more energy efficient execution of learning algorithms
- Ensure compatibility between hardware accelerators and data in RDBMS

Potential Commercial Applications

- Machine learning
- Data processing
- Unify machine-learning, hardware acceleration, and databases

Background and More Information

Hardware accelerators are designed to make some functions more efficient than they would otherwise be. Currently, in order to use hardware accelerators with large datasets in Relational Database Management Systems (RDBMS), a user would have to manually extract, copy, and reformat the dataset to fit the accelerator's criteria.

Inventors

- Hadi Esmaeilzadeh
Previous Assistant Professor – Georgia Tech College of Computing
- Divya Mahajan
Doctoral Candidate – Georgia Tech College of Computing

IP Status

:

Publications

, -

Images

Visit the Technology here:

[Hardware Integrator for RDBMS and Machine Learning Accelerators](#)

<https://s3.sandbox.research.gatech.edu//index.php/print/pdf/node/3474>