

SQREAM DB FOR AD-TECH

Better, more accurate data science models with the power of SQream DB and the GPU

THE SHORT VERSION




SQream DB uses GPU technology to improve the performance of columnar queries by at least 20x on large data sets, while reducing the hardware required to perform the query. Typically a single 2U server equipped with a GPU is equivalent to a 42U rack full of servers.

WHY GPU TECHNOLOGY?



A single GPU PCIe card may have up to 4992 processors. These GPU processors are specifically designed to perform high-volume and high-velocity numerical computations on both fixed and floating point values. SQream has patented and proprietary technology that uses a GPU card as a massively parallel processor tailored to process complex SQL queries.



DATA SCIENCE IN AD-TECH




Data scientists are true scientists in practice in that they perform experiments to prove which hypothetical models perform the best. A model is iterated several times, on each iteration new parameters are supplied to the model to test accuracy and validity. Each iteration operates on a dataset that is the result set of a query or queries.

Reducing the latency of a complex query on a large dataset can increase the productivity of data scientists, implicitly reducing costs.

In Ad-Tech, data science models the behavior of an anonymous user by using cookies and device IDs to classify a user in a segment, with the goal to present an ad most suited to the user. **When the model is better, the fit is better for the user; the bid price is higher, and a higher price increases the bid/win ratio.** The advertiser can spend their budget with a higher likelihood of conversion, and the publisher earns the most money on their inventory.

The time spent in model training effects the frequency a model can be applied to production real-time bidding. Lower query latency shortens the model training time. If you can produce a new model faster; you can refresh the cookie profiles more frequently with **more accurate segmentation**.



QUERY LATENCY AT SCALE

Complex queries contain multiple filters, type conversions, complex *predicates*, exotic *join* semantics, *subqueries* and many of them.

When running this kind of query on large data sets (>100 terabytes in billions of rows in several tables), the number of numerical computations performed is a product of the complexity of the query predicates and the number of rows to be processed.

Even when distributed, a conventional query engine using CPUs alone cannot deliver the result within an acceptable period. The query latency is huge, ranging from many minutes to hours.

```
SELECT array_agg(players), player_teams
FROM (
  SELECT DISTINCT t1.t1player AS players, t1.player_teams
  FROM (
    SELECT
      p.playerid AS t1id,
      concat(p.playerid,':', p.playername, ' ') AS t1player,
      array_agg(pl.teamid ORDER BY pl.teamid) AS player_teams
    FROM player p
    LEFT JOIN plays pl ON p.playerid = pl.playerid
    GROUP BY p.playerid, p.playername
  ) t1
  INNER JOIN (
    SELECT
      p.playerid AS t2id,
      array_agg(pl.teamid ORDER BY pl.teamid) AS player_teams
    FROM player p
    LEFT JOIN plays pl ON p.playerid = pl.playerid
    GROUP BY p.playerid, p.playername
  ) t2 ON t1.player_teams=t2.player_teams AND t1.t1id <> t2.t2id
  ) innerQuery
GROUP BY player_teams
```

SQream DB can execute the same query on the same data set with a latency of seconds to minutes.

HOW SQREAM DB CAN HELP

SQream DB is built from the ground up, to make the best use of available resources, including the revolutionary power of the GPU.

Standard SQL is used to query the data using industry standard drivers like JDBC, ODBC and Python – meaning most tools like Jupyter notebooks and Spark SQL are supported.

SQream DB combines performance, flexibility and ease-of-use, empowering your data science and making discovery insights in your data fast, allowing you to focus on the core of your business, not on the infrastructure.

