Clustered TDB: A Clustered Triple Store for Jena

Owens Alisdair, et al.

Presented By: Giovanni Thenstead

Outline

- Definitions
- Clustered TDB Overview
- Considerations
- Evaluations
- IJCAI Review

Definitions

- Cluster
  - We consider a cluster to be a collection of closely related computation nodes acting as a whole.

- Hotspot
  - The event in which a machine is the bottleneck in a cluster as a result of handling a near uneven amount of data or computation load.

Clustered TDB Overview

- Distributed prototype-variant of the Jena Tuple Database (TDB).
- Supports virtual servers (aka v-nodes).
  - This improves fault tolerance and balancing
- Tuple data are accessed by triple indexes with 64-bit indexing on each RDF attribute (SPO, POS, OSP).
  - 8 bits reserved for datatype
  - Remaining 56 bits reserved for disk (address, v-node)-pair, or as inline.

Clustered TDB Overview

- Uses Query Coordinators
  - Decreases network congestion and improves routing logic to reduce hotspots.
    - Done by storing statistical information
    - Holds routing tables (using Node/NodeID mappings) of each physical node (Data Nodes)
- Borrows from Jena's philosophy
  - Hopes to achieve fast read/write access compared to other standalone TDBs.
  - ARQ backend optimizes SQL query requests (via an execution plan) and performs variable-binding where possible.
- Extends Jena's philosophy
  - Supports merging and splitting parallel operations.

Considerations

- The data used was synthetically generated.
  - 375 million tuples used
  - 4,000 unique properties

Evaluation

- Loading benchmarks
  - Datasets considered: Standalone, 1-machine cluster, 2-machine cluster, 3-machine cluster.

Table 1. Load rates

System Type	Average Load Rate (triples/s)
Standalone	6,946
CTDB1	4,276
CTDB2	8,973
CTDB3	12,536

- Comments
  - Good scalability as more machines are added to cluster.

Evaluation

- Reading
  - Datasets considered: 1 and 5 users accessing the cluster simultaneously.

Table 2. Reading individual triple patterns (1 user)

System Type	SPO Index (ms)	POS Index (ms)	OSP Index (ms)
Standalone	439	21530	512
CTDB1	361	1733	483
CTDB2	187	1958	264
CTDB3	178	1821	363

Table 3. Reading individual triple patterns (5 users)

System Type	SPO Index (ms)	POS Index (ms)	OSP Index (ms)
Standalone	1904	65602	2205
CTDB1	1784	6648	2053
CTDB2	664	6879	672
CTDB3	647	5617	628

- Comments
  - Reading seems to scale well with more users.

IJCAI Review

- Relevance: 4
  - A scalable and efficient distributed TDB solution is desirable, but Clustered TDB includes too much complexity and requires more resources than may be necessary.
- Significance: 5
  - The author's deserve credit for trying to think outside-the-box, however, I'm afraid that a lot of emphasis is being placed on increasing the hardware expense in favor of an effective software/algorithmic solution. Also, it does not help that the authors' have admitted that their index lookups can be larger than the data themselves—this could potentially lead to poor memory-cache coherency.
- Technical Soundness: 4
  - There are technical incompatibles in regards to the storing of attributes with high cardinalities, and also in regards to possible hotspot issues that their de-clustering technique were unable to address.
- Novelty: 4
  - Borrows a lot from current DBMS and TDB technologies.

IJCAI Review

- Quality of Evaluation: 5
  - Though the evaluation did not cover a broad scope of scalability issues, the assessments were acceptable enough for a Clustered TDB prototype.
- Clarity: 4
  - The paper was fairly-well presented: surveying various TDBs and techniques. However, I almost lost interest while reading through the evaluation section as it was becoming more technical and ambiguous in some sense.
- Confidence Score: 4
  - In my opinion, this paper was more useful as a refresher on various TDB and DBMS technologies but not as a viable distributed TDB solution.
- Overall Score: 4.29

References

- Owens, Alisdair et al. “Clustered TBD: A Clustered Triple Store for Jena.” <http://eprints.ecs.soton.ac.uk/16974/1/www2009fixedref.pdf>

Questions?