Low Latency Cloud Data Management through Consistent Caching and Polyglot Persistence

Felix Gessert

Database and Information Systems Group
University of Hamburg
gessert@informatik.uni-hamburg.de

The ongoing increase of complexity, mobility and scale in modern applications triggered a paradigm shift towards distributed, cloud-based data management. The expanding field of NoSQL and cloud data stores encompasses a rich variety of systems that deal with non-functional requirements of these applications such as latency, throughput, availability and elastic scalability. However, two central problems remain unsolved. First, the performance of mobile and web applications is governed almost exclusively by latency. Since the recent shift to smarter clients and single-page applications, dynamic database content is mostly requested in end devices directly. This makes data requests extremely latency critical, as they block the user experience . Thus, performance cannot be solved at the database level alone but end-to-end latency has to be addressed, too. Second, the heterogeneity and complexity of different data stores make it tremendously difficult for application developers to choose an appropriate system and reason about its performance and functionality implications. The situation is frequently complicated when no one-size-fits-all solution satisfies all requirements. Until now, the overhead and required know-how to manage multiple database systems prevents many applications from employing polyglot persistence.

We introduce an integrated solution to both the latency and diversity problem as ORESTES [1], a database-as-a-service middleware capable of exposing different data stores through a uniform REST interface and database-independent data model [2]. To solve the latency problem we propose the *Cache Sketch* [3]. It is the first approach to exploit the web's expiration-based caching model and its globally distributed content-delivery infrastructure which were previously considered irreconcilable with dynamic workloads. Cache Sketches guarantee rich tunable consistency (Δ-atomicity [6, 7]) using Bloom filters to create compact representations of potentially stale records to shift the task of cache coherence to clients. Furthermore, the number of invalidations on caches that support them (e.g., CDNs) is minimized. With different age-control policies the Cache Sketch achieves very high cache hit ratios with arbitrarily low stale read probabilities. The *Constrained Adaptive TTL Estimator* complements the Cache Sketch by a statistical framework for inferring cache expiration dates (TTLs) that optimize the trade-off between Cache Sketch size, cache hit ratio and the number of invalidations. The *YCSB Monte-Carlo Caching Simulator* offers a generic framework for simulating the performance and consistency characteristics of any caching and replication topology. Simulations as-well-as real-world benchmarking provide empirical evidence for the efficiency of the Cache Sketch and the

considerable latency reductions it achieves. To provide even stronger safety guarantees we propose *Scalable Cache-Aware Transactions* [4] that attain optimistic ACID transactions over a wide range of unmodified data stores, relying on Cache Sketches to minimize the abort rates of optimistic concurrency control.

Instead of prescribing the use of one particular data store, we propose the *Polyglot Persistence Mediator* (PPM) [5] that automates polyglot persistence based on service level agreements (SLAs) defined over functional and non-functional requirements. In a three-step process, tenants first annotate schemas with SLAs (e.g., $latency_{read} < 30ms$). In the second step, the schema annotations are recursively scored against available data stores, yielding a routing model comprised of a mapping from schema elements to data stores. In the third step, the PPM in ORESTES employs the routing model to delegate requests to appropriate data stores, manages replication and collects metrics for scorings. Preliminary experimental results show drastic performance improvements for scenarios with high write throughput and complex queries.

In our ongoing work we are extending the Cache Sketch approach to query result caching. In particular we are designing a scalable stream processing system to detect required query result invalidations and Cache Sketch additions. The TTL estimator is extended to an integrated reinforcement learning process that predicts TTLs of single records as well as query results. We are also broadening the scope of the cache-aware transaction scheme by providing transparent selection of either general-purpose optimistic concurrency control or specialized transaction protocols for certain workloads like write-only transactions. For the PPM we are currently introducing active request scheduling and multi-tenant workload management, as well as improved scoring through machine learning techniques. The research results around ORESTES also form the technological basis of a backend-as-a-service startup called *Baqend*.

References

1. F. Gessert and F. Bücklers, ORESTES: ein System für horizontal skalierbaren Zugriff auf Cloud-Datenbanken, in Informatiktage, 2013.
2. F. Gessert, S. Friedrich, W. Wingerath, M. Schaarschmidt, and N. Ritter, Towards a Scalable and Unified REST API for Cloud Data Stores, Workshop Data Management in the Cloud, 2014, Bd. 232, S. 723734.
3. F. Gessert, M. Schaarschmidt, W. Wingerath, S. Friedrich, and N. Ritter, The Cache Sketch: Revisiting Expiration-based Caching in the Age of Cloud Data Management, in Datenbanksysteme für Business, Technologie and Web (BTW), 2015.
4. F. Gessert, F. Bücklers, and N. Ritter, Orestes: A scalable Database-as-a-Service architecture for low latency, in CloudDB Workshop, ICDE, 2014, S. 215222.
5. M. Schaarschmidt, F. Gessert, and N. Ritter, Towards Automated Polyglot Persistence, in Datenbanksysteme für Business, Technologie and Web (BTW), 2015.
6. S. Friedrich, W. Wingerath, F. Gessert, and N. Ritter, NoSQL OLTP Benchmarking: A Survey, in Workshop Data Management in the Cloud, 2014, Bd. 232, S. 693704.
7. W. Wingerath, S. Friedrich, and F. Gessert, Who Watches the Watchmen? On the Lack of Validation in NoSQL Benchmarking, in Datenbanksysteme für Business, Technologie and Web (BTW), 2015.