

Database caching is the practice of keeping an amount of data in memory, to reduce the cost of accesses to the main storage, and thus improve the performance of a database machine. We particularly focus on two properties of database caches. Firstly, how different cache replacement policies decide what data is kept in memory and secondly, what options exist to allow parallel accesses to the cached set by multiple threads. With a limited resource of access logs from real NoSQL databases in production, we will measure the performance of the discussed replacement policies. And we will measure the in-memory performance of the data structures that enable parallelism.