

Runtime-Optimized Analytics

Anastasia Ailamaki^{1,2}

¹*Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland*

²*RAW Labs SA, Switzerland*

Abstract

The ever-increasing demand for diverse real-time analysis on exponentially growing data has brought a series of new system design challenges: First, we can no longer afford to pre-load the data in a database in order to support interactive analytics. Second, with the semiconductor advancement predicted by the end of Dennard scaling, hardware in servers becomes increasingly heterogeneous. Third, the need for throughput is increased as a function of the number of concurrent queries issued by applications and users, but current work sharing techniques do not scale. Fourth, data pipelines are made of heterogeneous tools, each optimized for each processing step, but cross-tool communication introduces high overheads. Finally, we need real-time processing over fresh data (aka Hybrid Transactional Analytical Processing or HTAP), but interference between heterogeneous workloads results in suboptimal performance. The common theme is increasing heterogeneity which is impossible to address efficiently with system design decision made ahead of time, as at design time we know too little too early. Runtime decisions about both mechanisms and heuristics, on the other hand, always lead to efficient processing because optimal processing depends on the use case properties (dat, workload, hardware, concurrency). I will discuss novel just-in-time (JIT) systems which make and actuate decisions at runtime, and explain how the individual JIT solutions synthesise a real-time intelligence paradigm that helps resolve most system performance challenges.

