

Silentium! Run–Analyse–Eradicate the Noise out of the DB/OS Stack

Wolfgang Mauerer^{1,2}, Ralf Ramsauer³, Edson R. Lucas F.⁴, Daniel Lohmann⁵,
Stefanie Scherzinger⁶



Abstract: When multiple tenants compete for resources, database performance tends to suffer. Yet there are scenarios where guaranteed sub-millisecond latencies are crucial, such as in real-time data processing, IoT devices, or when operating in safety-critical environments. In this paper, we study how to make query latencies deterministic in the face of noise (whether caused by other tenants or unrelated operating system tasks). We perform controlled experiments with an in-memory database engine in a multi-tenant setting, where we successively eradicate noisy interference from within the system software stack, to the point where the engine runs close to *bare-metal* on the underlying hardware.

We show that we can achieve query latencies comparable to the database engine running as the sole tenant, but without noticeably impacting the workload of competing tenants. We discuss these results in the context of ongoing efforts to build custom operating systems for database workloads, and point out that for *certain* use cases, the margin for improvement is rather narrow. In fact, for scenarios like ours, existing operating systems might just be *good enough*, provided that they are expertly configured. We then critically discuss these findings in the light of a broader family of database systems (e.g. including disk-based), and how to extend the approach of this paper accordingly.

Keywords: Low-latency databases; tail latency; real-time databases; bounded-time query processing; DB-OS co-engineering

1 Introduction

The operating system is frequently considered boon and bane for the development of scalable service stacks. While general-purpose operating systems (like Linux) provide a great deal of hardware support, drivers and system abstractions, they have also been identified as a cause of jitter in network bandwidth, disk I/O, or CPU [Ar09; SDQ10; Xu13] when operating software services in cloud environments, where multiple tenants compete for resources. Naturally, this also affects the performance of cloud-hosted database engines [Ki15].

¹ Ostbayerische Technische Hochschule Regensburg, Germany wolfgang.mauerer@othr.de

² Siemens AG, Corporate Research and Technology, Munich, Germany

³ Ostbayerische Technische Hochschule Regensburg, Germany ralf.ramsauer@othr.de

⁴ Universität Passau, Germany edson.lucas@uni-passau.de

⁵ Leibnitz Universität Hannover, Germany lohmann@sra.uni-hannover.de

⁶ Universität Passau, Germany stefanie.scherzinger@uni-passau.de

Unacceptable noise and long-tailed latency distributions, but also the recent advances in hardware technology, have renewed interest in building database-specific operating systems. While historically, database and operating-systems research have been highly interwoven, the communities have parted ways in the past, and are just now rediscovering potential synergy effects (e.g. [Ca20; Mü20]). This has sparked immense interest in devising novel system architectures [KSL13], especially for database-centric operating systems kernels (e.g., [Ca20; Gi13; Gi19; MS19; Mü20]) that aim at deterministic performance. However, implementing an OS kernel is a herculean effort with tremendous follow-up costs, requiring substantial and largely duplicate effort for otherwise generic tasks, such as writing and maintaining device drivers, file systems, and infrastructure code, among others.

About This Paper. We take a fresh look at standard operating systems for low-latency/high determinism workloads, as they arise in real-time scenarios. Similar problems arise in cloud settings, where latency effects along the data path add up and can lead to substantial systemic problems, as Dean and Barroso have pointed out [DB13]. Rather than designing a new kernel from scratch⁷ to avoid noise and jitter, we follow an orthogonal approach, employing existing open-source components: Identify the root causes, analyse, and then address them as far as possible within the *existing* components. If necessary, enhance.

By vertical, cross-cutting engineering, we tailor the stack towards the needs of database engines, eradicate interference, and ultimately, reduce any noise-induced latencies in query evaluation. Our first results show that in many cases, a large degree of jitter is avoidable by the well-considered and purposeful employment of existing architectural measures – actually measures originally developed for other domains, such as embedded real-time systems. We present controlled experiments with an in-memory database engine running in a multi-tenant scenario on a number of different system software stack scenarios.

We focus on in-memory database engines as a specific (and deliberately narrow) use case, as they are often employed in domains for which deterministic latencies are essential [BL01], and thus considered a particularly convincing use case for developing specific operating-systems or even a bare-metal database stack [Bi20; Ca20; Gi13]. In this realm, our experimental setup, which is available as a Docker image for easy reproduction, can also serve as a baseline for researchers building special-purpose operating systems to compare their results against. In particular, we claim the following contributions:

- We perform controlled experiments with an in-memory database engine running on custom system software stacks based on existing open source components. By careful cross-cutting engineering, we modify this stack to eradicate interference, and to ultimately reduce any noise-induced latencies in query evaluation.
- We show that we can achieve the same performance using available operating systems as compared to running the database workload (near) *bare-metal*.

⁷ Whether to build a new operating system from scratch or whether to extend existing systems to cater to data processing needs has been an ongoing debate for decades [Gr78].

- We show that we can achieve the same performance in a multi-tenant scenario as compared to a database engine executing as the sole tenant without competing load.
- We voice doubts whether these specific scenarios can benefit from operating systems custom-designed towards database workloads, as they are currently being proposed.
- We discuss the potential generalisability of our approach to disk-based database engines, and systems involving I/O. In particular, we discuss opportunities that call for the joint efforts of the operating systems and database communities.

Structure. Our paper is structured as follows. We give an overview in Section 2, survey related work in Section 3, and present our experiments in Section 4. Their consequences are discussed in a more general context in Section 5. We conclude in Section 6.

2 Overview

We start with a brief summary of possible perturbations of an executing database workload by neighbourly noise, followed by an overview of the system software stack scenarios considered in this paper. In this section and beyond, by the term *kernel* we refer to the operating systems kernel (not the database kernel).

2.1 Sources of Noise

The three major sources of noise as observed by an unprivileged userspace workload (as compared to system services or the kernel) are (1) other processes and system services that compete for CPU usage, (2) CPU performance optimisations (caches, pipelines, . . .) that can usually not be disabled or controlled, and (3) contention of implicitly shared resources (memory bus etc.). The signature of such *systemic* noise is not necessarily distinguishable from the *intrinsic* noise of the application, that is, variations in run-time caused by data-dependent code paths, application-specific optimisations, and so forth.

Processes and system services. Multi-tasking operating systems manage M schedulable entities that compete for N processors, with $M \gg N$. Linux uses a *completely fair scheduling* (CFS) [Ma10] policy for regular processes, but also includes support for (soft) real-time scheduling via FIFO and round robin. The kernel can preempt most userland activities (depending on the preemption model statically configured at kernel build time), for instance upon the arrival of interrupts. It can also place *kernel threads* into the schedule that perform activities on behalf of the kernel (for instance, to support migrating processes across CPUs, to perform post-interrupt actions, etc.), and enjoy higher priority than regular processes, regardless whether these are governed by real-time policies. The interplay of these factors creates noise compared to an uninterrupted, continuous flow of execution of a single job.

CPU noise. Even given the uninterrupted execution of code on a CPU, pipelined and superscalar execution of code may lead to different temporal behaviour than would be

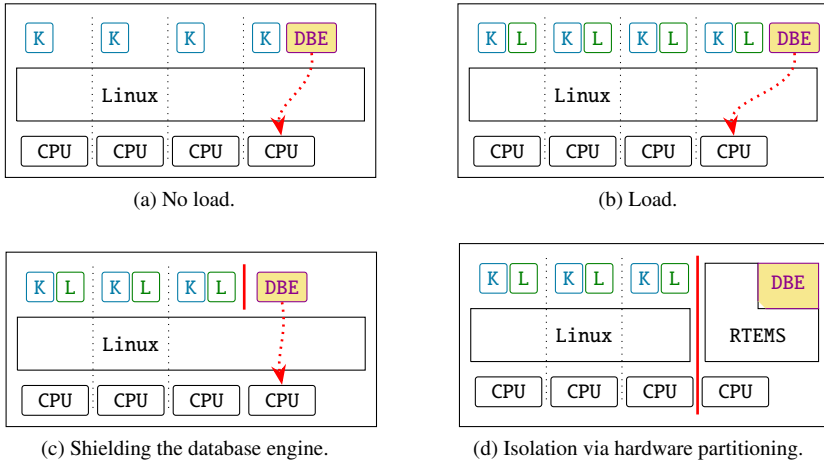


Fig. 1: System software stack scenarios.

achieved by a straightforward execution of assembly instructions, which manifests itself as another source of noise. Also, caching mechanisms (most importantly, the cache hierarchy that comes into play with memory references, but possibly also mechanisms like translation lookaside buffers used in virtual-to-physical address translation) cause (widely) varying latencies in accessing memory. This effectively adds noise.

Shared resources. Workloads executing on different CPUs are not entirely isolated from each other, but interact via shared resources (cache, memory, etc.) that are accessed via system buses. This even holds despite a possible logical partitioning of system components that we discuss later. While the overall situation (for instance, handling competing requests for bus usage) is deterministic from a system-global view, delays caused by competing requests manifest as noise when viewed from the perspective of an individual process.

2.2 Experimental System Configurations

The configurations of the system software stack, as used in our experiments, are visualised in Figure 1. For now, we treat the in-memory database engine (DBE) as a black box.

No Load. In the *no-load* scenario (Figure 1a), a single database engine executes on an otherwise quiet multicore system. The database payload is pinned to one CPU (c.f. the dashed arrow), to avoid perturbations, for instance caused by the scheduler moving the process across CPUs. However, standard system services, as limited to the bare necessity, and kernel threads required by the operating system proper (“K” in the figure) can execute on all CPUs, including the CPU dedicated to the database workload.

Load. In the *load* scenario (Figure 1b), additional tenants put the system under maximum strain. We simulate this payload (marked “L”) by running synthetic workloads on each CPU.

While the database workload is again pinned to one particular CPU, it is scheduled by the operating system alongside kernel threads and the described payload. In the load scenario, we assume the viewpoint of a cloud provider maximising the utilisation of the available resources, while serving all tenants equally. We therefore refrain from assigning the database workload higher priority than the payload generated by the competing tenants.

Load/FIFO. A variation of the *load* scenario uses standard Linux mechanisms to set a real-time scheduling policy for the database workload. All load processes fall under scheduling policy “other”, and compete for CPU resources as managed by the Linux standard scheduler. We place the database task in the real-time scheduling group `SCHED_FIFO`, so it can preempt any other userland tasks that execute on the CPU. However, the database task can still be preempted by the kernel, or by incoming interrupts.

Shielding. Another approach towards isolating the database workload from noise is to use CPU shielding (Figure 1c). This distributes all existing tasks and kernel threads on a given CPU to the rest of the system, and prevents utilisation of the shielded CPU by the standard scheduling for processes that are not explicitly assigned to this CPU.

We additionally make sure that incoming external interrupts only arrive at other CPUs. Nevertheless, main memory, buses, caches, etc. remain shared resources in the system, and accesses can induce additional noise that goes beyond the pure CPU noise. Additionally, the kernel can still preempt the single running userland task (for instance, when timers expire), and latencies can arise from administrative duties performed by the kernel on such occasions, or in the context of system calls issued by the task.

One set of measurements combines shielding and real-time priorities. This limits the kernel’s abilities to preempt the running userspace task. However, some caution needs to be administered: Not only the ability to preempt a running task, but also the amount of work performed in kernel context when a preemption occurs influences latency variance, and this amount is highly dependent on specific (static) kernel configuration settings.⁸ Isolation in this scenario is based on guarantees provided by the Linux kernel. This implies trust in a complex, monolithic code base, which is undesirable for safety-critical scenarios.

Partitioning. The strongest form of isolation that we consider in this paper (Figure 1d) relies on the Jailhouse hypervisor [Ra17]. Jailhouse can partition system hardware resources by establishing independent and strictly isolated computing domains. Jailhouse leverages extensions of the underlying system architecture which include essential virtualisation mechanisms for system partitioning, such as segregation of CPUs, memory and devices, as

⁸ Linux provides a tick-less mode, which eliminates periodic interventions by a regular timer (at frequencies ranging from 100Hz to 1,000Hz, depending on compile-time settings), but which may cause overhead on other occasions, because maintenance of data structures performed during such ticks must be performed “en block”.

well as additional extensions that allow to control the utilisation of shared resources, such as caches or system buses.

Jailhouse comes at a negligible performance overhead, as it does neither (para-)virtualise or emulate resources, nor schedule its partitions (*guests*) among CPUs. The virtual machine monitor only interferes in case of critical exceptions and access violations. This architecture can find application in multi-tenant database scenarios, described in [MKN12], and in particular, safety-critical scenarios, which require spatial-temporal isolation between tenants.

Bare-Metal Operation. Data center, cloud and high performance data processing systems often employ x86 server class CPUs, and we have argued before that such use cases benefit from bounded tail latencies. Other important use cases that require determinism are found in embedded systems, which are typically equipped with ARM CPUs. Consequently, our investigation addresses both, x86 and ARM.

Using a simplistic ARM core that is just capable enough for realistic database deployment reduces systemic noise that stems from multicore effects, as found on server-class x86 CPUs [PH90] to the bare minimum, (e.g. long pipelines, large caches, and strong interference on buses). This allows us to explore the *intrinsic* variations of our database workload.

We employ the in-memory database engine DBToaster [Kol4] (see also Section 4), a highly portable serverless database engine that requires a C++/STL run-time environment, but no other libraries or system services. Plain C++ can be executed without relying on an OS proper with moderate effort, but the STL requires (at least conceptual) support for threads and preemptive locking, as well as a full memory allocator. These requirements do not create a need to deploy it on top of a fully-fledged general-purpose operating system, such as Linux or Windows, but we deem the implementation efforts large enough to warrant a tiny operating system. Thus, we ported the database engine to RTEMS (real-time executive for multiprocessor systems) [BS14], a mature, tailorable embedded real-time operating system (with a 25-year development history) that finds deployment in systems ranging from IoT devices to Mars orbiters. Similar to unikernel approaches [Br15; Ma13], RTEMS and the database engine are linked together into one single executable. This binary can be either booted as stand-alone operating system on a bare-metal system, or (given low-level changes like the use of a custom bootloader and adaptations of the RTEMS kernel to Jailhouse) be executed in parallel to Linux on a partitioned system (as visualised in Figure 1d).

To reduce operating system noise as far as possible, we essentially limit RTEMS to providing only a console driver, and execute the database engine in a single thread, which eliminates the need for a scheduler. This configuration is supposed to reduce any OS noise to the bare minimum, and is *comparable to a bare-metal*⁹ main-loop style binary.

⁹ Bare-metal operation refers to code that runs without distinction between payload and OS close to the hardware, without intermediary layers. This in contrast to, for example, Ref. [RF18], which uses the term to denote code that runs without containers or virtual machines, but still relies on heavyweight, multi-million-LoC OS kernels.

3 Related Work

In this paper, we focus on in-memory database engines. We refer to [GS92] for an early overview of their architecture, and to [Fa17] for a more recent survey.

In real-time scenarios, where in-memory database engines traditionally play an important role [BL01], deterministic latencies are crucial. However, aspects such as consistency of query answers given transactional workload, or alternative tuple consumption strategies, are of no concern to the work presented in our paper, since we treat the database largely as a black box, and are interested in the overall system software stack.

What is indeed highly relevant for us is the existing work on worst-case execution time (WCET) of queries in in-memory databases [Bu05], which considers control-flow graphs through the code (in fact, in the presentation of our experiments, effects of different paths through control flow graphs during query processing actually become visible).

Also close to our work in both methodology and context is research on the influence of NUMA effects, focusing on in-memory database engines in particular. It is known that assigning threads to CPUs improves database performance, due to caching effects [Do18; Ki15; KSL13]. Similar studies of assigning database workloads to computational units can be found throughout database research, for instance in Refs. [DAM17; Po12]. Our experiments also assign threads to dedicated CPUs, and we benefit from data caching, but our motivation differs, as we *isolate* the database workload from harmful noise.

Databases operating in multi-tenant environments are another focus of our work. This differs from many benchmarks conducted in database research, where database workload often runs in isolation, while multi-tenant environments are closer to real-world conditions. Similarly, an overview over *performance isolation* for cloud databases is provided in [Ki15].

A systematic discussion of multi-tenant in-memory databases is provided in [MKN12]: From the viewpoint of a cloud provider, guaranteeing narrow service-level-agreements is a challenge, since the provider must cater to all tenants, while utilising the hardware resources. This mindset is also found in engineering for mixed-criticality systems [BD17; Ve07], where a critical workload (in our case, the database engine) must be shielded from noise (in our case, competing tenants), without cutting into the performance of the remaining workload.

In designing multi-tenant database engines, shielding tenants can be implemented on several levels in the system software stack. Aulbach et al. [Au08] enable multi-tenancy on the level of the database schema; by appropriately mapping between the tenants' schemas and the internal schema, tables may be transparently shared between tenants. By rewriting queries, the authors ascertain isolation between tenants in an otherwise standard database engine.

Narasayya et al. [Na13] also aim at resource isolation, for the database-a-service provider Microsoft SQL Azure. They explore virtual machine mechanisms in userland without relying on mechanisms provided by the kernel (and, consequently, not benefiting from the

guarantees provided by the OS kernel – for instance, some isolations are not possible in userland, such as access to shared buses and other resources).

Noll *et al.* [No18] discuss how to accelerate concurrent workloads inside a single database engine by partitioning caches. This feature is not targeted at multi-tenant databases *per se*, but applicable in general. However, this feature is specific to Intel CPUs. Further, it is not directly subject to control from userland, but exposed to applications by the `sysfs` pseudo-filesystem interface of the Linux kernel. Our x86-based RTEMS measurements in a Jailhouse cell actually use the same infrastructure to assign a portion of the cache to the system performing the measurements, which reduces variations in memory access times.

The general idea of using existing OS-level isolation mechanisms to reduce the amount of inference between latency (or otherwise) sensitive database workloads and the rest of a system has also been pursued by Rehmann *et al.* [RF18]: The authors use Docker containers to isolate database instances from system and competing payload noise. Their work essentially implements limiting the CPU quota available to tasks, and pinning database-relevant operations to specific CPUs in the system. Especially the latter is similar to some of our experiments, albeit we additionally include scheduling prioritisation and control the system noise on pinned CPUs with various measures. Thus, we make use of a richer toolset to achieve stronger levels of isolation, as our measurements show. In fact, containers are conceptually not intended to isolate a given workload from other workloads, but to provide a specific, probably restricted view of the system to a given workload.

Currently, there is renewed interest in building database-specific operating systems, partly motivated by such problems as unpredictability in performance. For instance, the MXKernel project [Mü20] proposes an alternative to the classic thread model, to cater to the demands of large-scale data processing. The DBOS initiative [Ca20] goes so far as to envision managing database-internal data structures inside the OS kernel. Further, there are suggestions to share the database cost model with the operating system [Gi13], to allow for more transparency and to ultimately arrive at better scheduling decisions.

Recent developments in modern hardware, and in particular modern memory technology, motivate database architects to re-evaluate the entire DBMS systems architecture and in-memory data structures [AP17; Re19; St07]. Over the years, research in this area has delivered promising propositions, e.g. [APD15; Ch18; GTS20; Le20]. In contrast, we evaluate how far *existing* technology will take us, given careful, cross-cutting engineering.

4 Experiments

We next describe the setup of our experiments, and then present our results. Our Docker image¹⁰, which we describe in Appendix 7, allows for inspection and reproduction.

Database Engine. We conduct our experiments with the in-memory database engine DBToaster [Ko14]. DBToaster can compile SQL queries to C++ code, which we then compile

¹⁰ Available online from <https://github.com/1fd/btw2021>.

(in a second step) for our target platform. The resulting executable is a single-threaded database engine that incrementally updates a SQL view given a tuple stream. DBToaster is thus a SQL-to-code compiler, designed to maintain materialised SQL views with low refresh latencies. Typical application scenarios would be in stream processing, such as algorithmic trading, network monitoring, or clickstream analysis¹¹. The DBToaster system and its theory have been prominently published (e.g. [KLT16; Ko10; Ko14; NDK16]).

We have created our own fork of the DBToaster code base (which is open source), with minor modifications for our experiments (e.g., buffering measurement data in memory, rather than writing directly to standard output). Our fork is part of our reproduction package.

Data and Queries. We consider two benchmark scenarios from the DBToaster experiments in [Ko13]. To be able to discuss the run-time results in greater detail, we focus on only a subset of queries. In particular, we exclude queries that display a high level of *intrinsic* variability in their latencies, where the computational effort between tuples can vary greatly, for instance because of nested correlated sub-queries and multi-joins. These queries are *per se* not well-suited for stream processing. The queries considered by us are listed in Figure 2.

Finance queries. The queries over financial data process a tuple stream with stock market activity; we chose three queries which use different relational operators: Query *coun-tone* (C1) is designed by us and serves as a minimal baseline. DBToaster can incrementally evaluate this query with constant-time overhead per tuple. The queries *axfinder* (AXF) and *pricespread* (PSP) each compute a join, a selection, aggregation, and in the case of *axfinder* also a group-by on the input stream. Here, we use the exact same query syntax as in [Ko13], as DBToaster has certain restrictions (e.g., no LEFT OUTER join). To be able to execute these queries on hardware devices with very limited memory, we use a base data set of 100 tuples¹², over which we iterate 5k times, yielding 500k data points. Since the query predicates do not filter on time-stamps, this does not affect query semantics.

TPC-H queries. We generated TPC-H data with the *dbgen* data generator, set to scale factor 4. We chose the queries Q6, Q1, and Q11a (shown in Figure 2) from the DBToaster experiments in [Ko13]. The queries perform selections, aggregations, and in the case of Q11a also a join.

Execution Platforms. For x86 reference measurements, we use a Dell PowerEdge T440. The T440 is equipped with a single 12 core Intel[®] Xeon[®] Gold 5118 CPUs and 32 GiB of main memory. For measurements on Linux, we use kernel version 5.4.38 (vanilla kernel as provided by kernel.org) as baseline, with the Preempt_RT real-time preemption patch.

Since delays are caused by parallel access to shared execution units and resources, symmetric multithreading (SMT) is a source of undesired high latencies and noise in real-time systems. Consequently, we deactivate SMT on our target, in accordance with the original DBToaster experiments. Furthermore, we deactivate Intel[®] Turbo Boost[®], as sporadic variations of

¹¹ See the project homepage at https://dbtoaster.github.io/home_about.html, last accessed January 2021.

¹² <https://github.com/dbtoaster/dbtoaster-experiments-data/blob/master/finance/tiny/finance.csv>

C1	<pre>SELECT count (1) FROM bids;</pre>
axfinder	<pre>SELECT b.broker_id, SUM(a.volume+(-1*b.volume)) AS axfinder FROM bids b, asks a WHERE b.broker_id = a.broker_id AND ((a.price+((-1) * b.price)>1000) OR (b.price+((-1) * a.price)>1000)) GROUP BY b.broker_id;</pre>
TPCH Q11a	<pre>SELECT ps.partkey, SUM(ps.supplycost * ps.availqty) AS query11a FROM partsupp ps, supplier s WHERE ps.suppkey = s.suppkey GROUP BY ps.partkey;</pre>
TPCH Q1	<pre>SELECT returnflag, linestatus, SUM(quantity) AS sum_qty, SUM(extendedprice) AS sum_base_price, SUM(extendedprice*(1-discount)) AS sum_disc_price, SUM(extendedprice*(1-discount)*(1+tax)) AS sum_charge, AVG(quantity) AS avg_qty, AVG(extendedprice) AS avg_price, AVG(discount) AS avg_disc, COUNT(*) AS count_order FROM lineitem WHERE shipdate<=DATE('1997-09-01') GROUP BY returnflag, linestatus;</pre>
pricesread	<pre>SELECT SUM(a.price + (-1*b.price)) AS psp FROM bids b, asks a WHERE (b.volume > 0.0001 * (SELECT SUM(b1.volume) FROM bids b1)) AND (a.volume > 0.0001 * (SELECT SUM(a1.volume) FROM asks a1));</pre>
TPCH Q6	<pre>SELECT SUM(l.extendedprice*l.discount) AS revenue FROM lineitem l WHERE l.shipdate>=DATE('1994-01-01') AND l.shipdate<DATE('1995-01-01') AND (l.discount BETWEEN (0.06-0.01) AND (0.06+0.01)) AND l.quantity<24;</pre>

Fig. 2: SQL queries used in the experiments (queries from [Kol3], with the exception of C1).

the core frequency result in non-deterministic execution times for identical computational paths. We configure the CPUs in the highest possible P-State (performance setting) that guarantees a stable core frequency of 2.29 GHz.

For the shielding scenario, we try to remove all operating system noise from the target CPU. The Linux kernel provides multiple mechanisms for this purpose, of which we choose CPU namespaces that can be dynamically reconfigured during system operation.¹³

For the partitioned Jailhouse setup, we release one single CPU and 1 GiB of main memory from Linux, and assign them to a new computational domain. On that domain, we boot the RTEMS + DBToaster binary¹⁴ that runs in parallel to Linux. We use Intel’s Cache Allocation Technology (CAT), part of Intel’s Resource Director Technology, to partition last-level caches and exclusively assign 5 MiB of Level 3 Cache (L3\$) to the RTEMS + DBToaster domain. This mitigates noise (cache pollution) of neighboured CPUs, as the L3\$ is shared across all cores [In15].

For the ARM reference platform, we use a BeagleBone Black with a single-core Sitara AM3358, a 32 bit ARM Cortex-A8 processor and 512 MiB of main memory. In contrast to the powerful Intel server CPU, such ARM processors are typically found in embedded or industrial applications. We boot the RTEMS + DBToaster application directly on bare-metal.

¹³ Other mechanisms like CPU isolation at boot-time would provide a slightly higher level of isolation, but must be statically configured at boot-time, limiting the flexibility of the setup.

¹⁴ Getting DBToaster to run on RTEMS was not straightforward; along our trials, several fixes were proposed to open source systems, such as a decade-old bug revealed in GCC, as well as a bug identified in RTEMS.

Methodology. DBToaster logs a time-stamp for every N input tuples processed. This allows us to compute the *latency per N input tuples processed*, averaged over N tuples. While averaging is a sensible and established choice for throughput measurements to minimise overhead of the measurement intervention, we are interested in a precise characterisation of system noise vs. intrinsic variation of the core processing code, and therefore resort to measuring processing times on a *per-tuple* basis ($N = 1$).

We distinguish between two units of measurements: (1) time stamps obtained by the standard POSIX API (`clock_gettime` with `CLOCK_MONOTONIC`). This allows for nanosecond resolution, but also inflicts considerable overheads in the microsecond range, and introduces a noise level that is on par with the processing time proper for some of the simpler queries. Therefore, we extend DBToaster with the optional capability of (2) using x86 time stamp counter (TSC) ticks. While there are several problems and pitfalls associated with using the TSC on SMP configurations, and while the obtained measurement values cannot be converted to walltime without further ado [Ma10], TSCs are one of the highest-precision clock sources available on x86 hardware, and can be read from userspace without transition to kernelmode.¹⁵

As is a standard approach in settings like ours [BL01], we start measuring time once the input is in memory. In particular, we pre-load all tuples prior to stream processing, to exclude noise caused by I/O. Of course, in any real-world setting, the tuples would be read over peripheral communication channels, such as ethernet. To further avoid noise in our measurements, we have modified the code generated by DBToaster such that these time-stamps are cached in memory during query evaluation, in a pre-allocated array, rather than being continuously written to the standard output console.

Simulating tenant load. We simulate further tenants executing on the same system using the standard utility `stress-ng`, running 6 synthetic workloads.¹⁶ In Figure 1, we depict `stress-ng` running as additional load on the CPUs that are annotated with “L”.

4.1 Results

4.1.1 Noise and Determinism: Finance Queries

We begin our discussion of results for the finance queries. The time series in Figure 3 show observed latencies for processing each out of 500k input tuples. Red, labelled triangles mark the minima and maxima. Since almost all measured values fall into a comparatively

¹⁵ Using a high-resolution, low overhead time source is not necessary on our ARM reference platform because the time required to obtain a time stamp is negligible in comparison to the average processor performance, and our operating system has a flat memory and privilege model – that is, there is no distinction between kernel- and usermode on our near-bare-metal measurements on this platform.

¹⁶ (1) Binary search on a sorted array (exercises random memory access and processor caches), (2) matrix multiplication (to stress memory cache and floating point units), (3) compressing/decompressing random data (exercising CPU, cache, and memory), (4) randomly spread memory read and writes (to thrash the CPU cache), (5) sequential, random and memory mapped read/write operations (to exercise the I/O subsystem), and (6) timer interrupts at the frequency of 1 MHz (to induce continuous kernel/userspace transition due to interrupt handling).

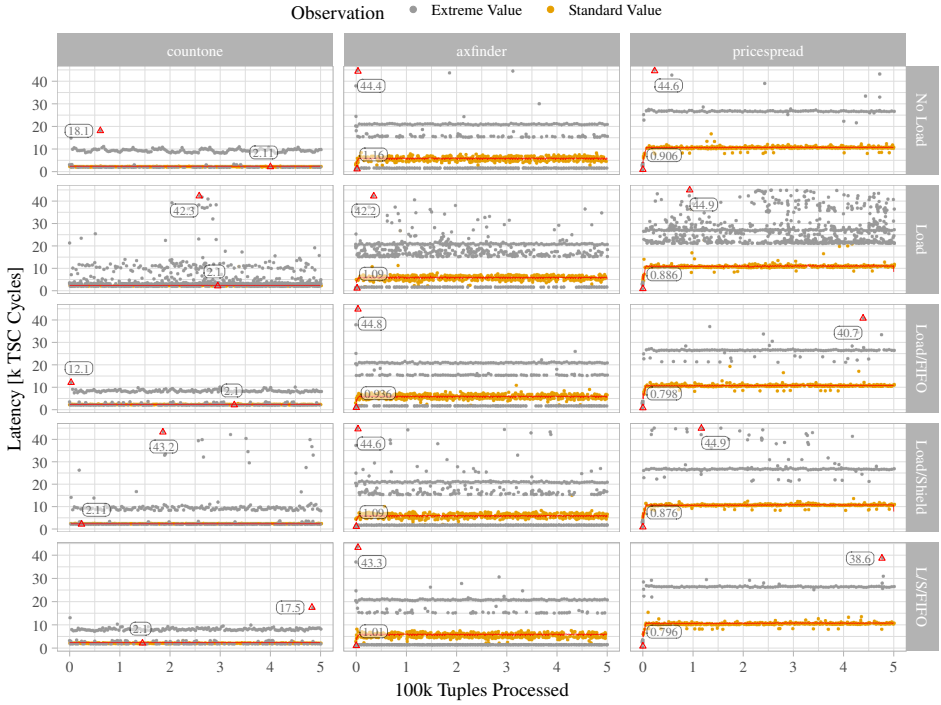


Fig. 3: Latency time series for finance queries on x86, using the high-speed time stamp counter (TSC).

narrow standard range, which would lead to massive over-plotting and loss of information, we colour all “extreme” measurement points that fall in the bottom 0.05% percentile, or that exceed the 99.95% percentile, in grey. We consider all other data points (marked ochre) as the normal observations. Note that such outliers have no noticeable influence when it comes to performance measurements, which usually concern query throughput, based on temporal averages, but are of paramount importance for real-time, bounded latency scenarios. For instance, the experiments in [Ko14] consider the query refresh times for processing batches of 1,000 tuples, and we compute a sliding mean window over 1,000 tuples as a consistency check; the resulting red line nicely reproduces the original DBToaster experiments [Ko14].

Each subplot of a given column corresponds to one system software stack scenario from Section 2. Almost all latencies are centred around the sliding mean value. However, a few outliers exceed the mean by a factor of about four.¹⁷

We have also tracked the average performance of the simulated other tenants, and found that it was essentially identical regardless of the measurement setup, which shows that improved

¹⁷ In the scenario discussed in this paper, the *maximum observed latency* is essential. Exceeding a threshold in industrial control scenarios might have severe consequences, from lost capital over destroyed machinery to bodily harm or loss of life, which can never be compensated by the fact the this does not happen *on average*.

determinism for a given workload does not necessarily decrease average throughput for non-real-time loads. Detailed data are available in the reproduction package.

While we consider queries of different intrinsic complexity, there is no direct relation between query complexity and noise – however, there *is* a relation between query complexity and *average* performance, as visible in the increasing latencies of the red line from left (simpler query) to right (more complicated query).

Query *countone* merely counts the number of input tuples processed so far (and is refreshed for each input tuple), whereas the other finance queries compute joins. As can be expected, the average latency for *countone* is distinctly lower. For the other queries, we can observe densely populated discrete “horizontal bands” that group the majority of all observed values. They correspond, based on an analysis of profiling data, to the main execution paths taken by the DBToaster-generated code (two main execution paths are a consequence of the “orderbook adapter” that distinguishes between the two types of input data, bids and asks). Also, when DBToaster-internal dynamic data structures grow in size (such as when buffering tuples for computing hash joins), additional DBToaster-intrinsic latencies incur.

The vertical spread of observations around these bands is an obvious visual noise measure. By comparing against the “Load” scenario, it is visually apparent that the different isolation mechanisms substantially reduce the observed jitter, typically to the level of an otherwise unloaded system. The strongest form of isolation, CPU shielding plus realtime scheduling (L/S/FIFO), produces latency distributions that are not only comparable, but even better in terms of maximum values than in the “No Load” scenario. The amount of noise decreases in order Load/Shield, Load/FIFO, and Load/Shield/FIFO. It might surprise that a shielded CPU performs worse than a CPU with additional load, but with a real-time prioritised task of interest. Recall that there is a complex interaction of kernel features as outlined in Sec. 2.2, and that, for instance, a larger set or possible preemption points, together with delayed kernel administrative work in a shielded scenario, may well compensate the advantages gained by exclusive CPU access.

While the measurements show a noticeable reduction of noise when using more advanced isolation techniques, the reduction of maximal latencies comprises only a factor of two.

4.1.2 Noise and Determinism: TPC-H Queries

The latency measurement results for the TPC-H queries are shown in Figure 4. While the general observations are identical – measured values concentrate around a few horizontal “bands”,¹⁸ and noise decreases with the various forms of systemic isolation – the behaviour of the queries under high load differs considerably from the “No load” and isolated case:

¹⁸ Notice that such bands may also be caused by system effects, and are then not necessarily present in all measurement combinations: For instance, the Load/Shield scenario in Figure 4 contains a band that disappears when FIFO scheduling is activated. Bands present in all scenarios are typically, but not necessarily, caused by the payload software. Such detail observations are not possible in measurements that average over observations.

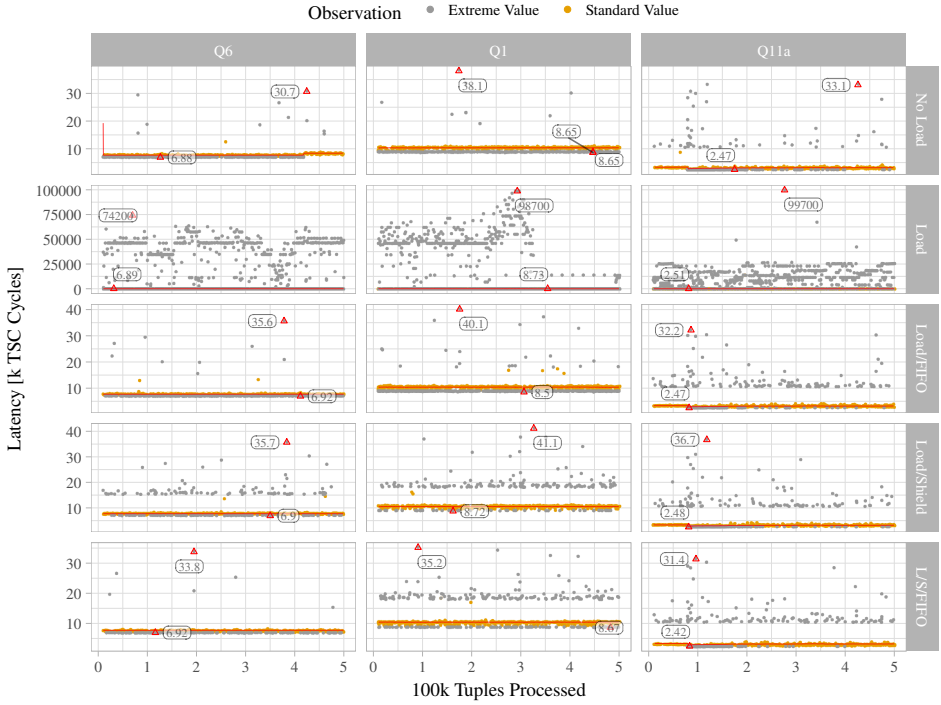


Fig. 4: Latency time series for TPC-H queries on x86, using the high-speed time stamp counter (TSC).

The difference in maximal latencies comprises more than three decimal orders of magnitude (observe the different scales of the vertical axes in the plots), and a similar statement can be made for the width of the spread around the running mean value, the latter again plotted with a red line. While such high variance has grave consequences for real-time systems, it is not even observable when throughput measurements are averaged.

So far, we have relied on visual means for characterising noise. For a quantitative measure, consider the set of observed latencies $\{\Delta t_i\}$ (each data point in Fig. 4 corresponds to one value of Δt_i). While we have focused on x86 so far, we will also consider ARM-based systems below. These platforms vary widely in their performance, and absolute values consequently require interpretation. It is therefore pertinent to consider *relative* deviations from the expected response time, which allows us to compare across platforms.

To this end, we define *spreads*, which are *not* influenced by the absolute processing speed. The *maximum spread* is given by $\max(\{\Delta t_i\})/\text{med}(\{\Delta t_i\})$, and *minimum spread* by $\text{med}(\{\Delta t_i\})/\min(\{\Delta t_i\})$, where $\text{med}(\cdot)$ denotes the median of the argument set. The quantities characterise the system-global relative span between a “typical” observed value, and the most extreme outliers in both directions. The results shown in the row labelled “TSC” of Fig. 5 quantitatively underlines this: Spread in the “Load” scenario typically exceeds




Fig. 5: Span of observations relative to median query latency of the TPC-H queries in Figure 4. Clock: Measured using `clock_gettime`. TSC: Measured using the high-resolution time stamp counter.

the spread in the isolation scenarios by orders of magnitude. The differences between the various isolation scenarios (and the “No Load” case) are much less pronounced, but can still encompass a factor of two or three (note the log-transformation of the y axis).

Recall that we distinguish between two units of measurements for latencies, (1) wallclock time in nanoseconds, and (2) x86 time-stamp counter ticks. The row labeled “Clock” of Figure 5 highlights another issue related to this fact that is mostly technical, but nonetheless requires careful consideration: how to perform the measurement itself. It shows the relative span for the *identical* measurement as considered in the other row, but this time using per-tuple latency measurements based on the POSIX API call `clock_gettime` offered as service by the Linux kernel (and often replaced by the lower-precision variant `gettimeofday`, in a good fraction of published performance measurements). Especially the maximum span can differ considerably among measurement variants. For TPC-H query 11a, it is even the major source of noise, as the right part of Fig. 5 shows.

Fig. 6 illustrates, for a subset of the isolation mechanisms, the increase in spread and noise distribution for clock-based measurements. It particularly highlights that even the mean throughput value (red line) is substantially influenced by the increased overhead.

4.2 The Role of CPU Noise

To a major extent, the previous experiments concern the control of noise introduced by the operating system and the presence of other tasks that compete for CPU time and other shared resources. Especially the scenario using CPU isolation, combined with a real-time scheduling policy, eliminates a substantial fraction of this noise. We now question how much of the remaining noise is caused by the executing CPU itself, and can thus be seen as an effective lower bound on any systemic noise.

We thus run our binaries as close to the bare-metal as possible, which reduces OS overhead. We perform these measurements on an ARM processor that we deem powerful enough to execute reasonable database operations, but that uses substantially fewer performance

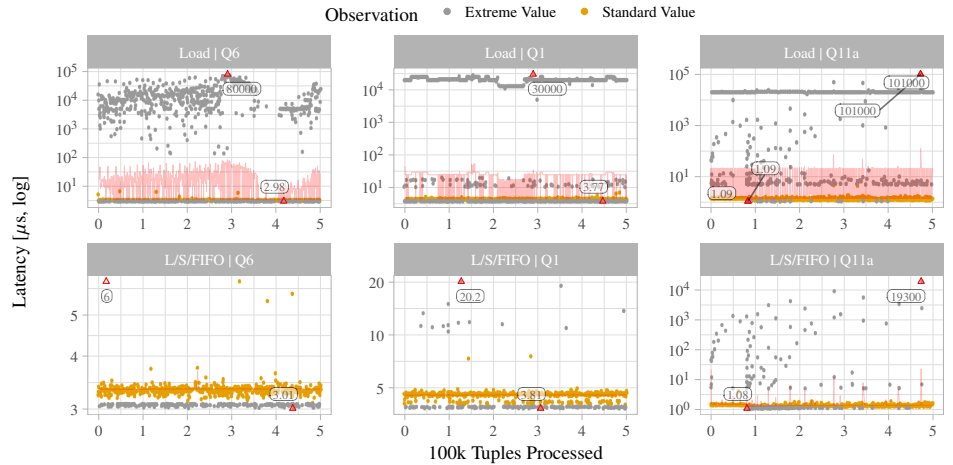


Fig. 6: Latency time series for TPC-H queries on x86, using a kernel-provided clock.

optimisations than x86 server-class CPUs (and, thus, suffers from less intrinsic noise). Our choice for an ARM CPU is not just driven by simplicity, though: Processors of this type are the most frequent choice in embedded systems and IoT devices, where low latency data processing is a common requirement (for instance, think of sensor-based systems that derive action decisions by combining previously measured values stored in a database with current data points). Our measurements are therefore representative for this large class of systems that we expect will gain even more importance in future applications. Of course, measurements on CPUs with drastically different capabilities cannot be directly compared, and this is not our desire: Instead, it is important to consider the *relative* difference between average and maximal latencies, and the span within measurements, as discussed below.

Fig. 7 shows latency time series for three finance queries. Again, observations centre around horizontal bands induced by the main execution paths, but the overall jitter is limited. The reduction compared to Jailhouse on x86 is quantified in Fig. 8.




Fig. 7: Latency time series for finance queries on an ARM system (BeagleBone Black) using RTEMS. Red, labelled triangles represent minima and maxima (not necessarily unique).




Fig. 8: Span of observations relative to median query latency of finance queries on bare-metal, on a high-end (x86) and low-end (ARM) CPU.

The summary for a second set of measurements shown in the bottom part of Fig. 8 represents bare-metal results obtained on the x86 CPU, but this time driven by an RTEMS kernel running inside a Jailhouse cell. Since the system is equipped with a total of 12 cores (compared to the single-core ARM), and only one of the cores is needed to run the database workload, we extend the measurement with an additional aspect that quantifies the aptitude of the setup to decouple latency-critical database operations performed by one tenant from other, perhaps throughput-oriented operations performed by other tenants. The spread is, as Fig. 8 shows, almost identical between the scenarios. This is also reflected in the time series shown in Fig. 9, which demonstrates that the results of the two configurations do not deviate in any substantial way. Since the isolation provided by Jailhouse does not only address execution timing, but also extends to other security and privacy related aspects of database workload processing, we deem this configuration a suitable basis for multi-tenant systems with strong separation guarantees.

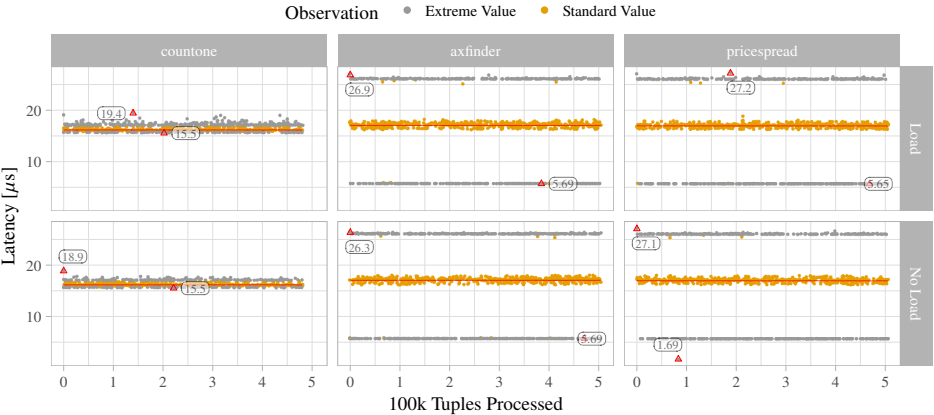


Fig. 9: Latency time series for finance queries on an RTEMS-based near bare-metal CPU provided by the Jailhouse hypervisor.

5 Consequences

Our experimental results show that – at least for our use-case of an in-memory database engine – building a database stack on a plain-vanilla Linux with custom settings can already compete with running close to bare-metal on the hardware. In our experiments we reached a state where the major source of noise turned out to be the interruptions to measure time (and noise) itself – any other sources for system-software induced jitter had been eradicated. In the following, we discuss some of these results in a broader context.

There’s life in the old dogs, yet. Our results suggest that it may not be necessary to design dedicated DB-aware operating systems *from scratch*. Rather, a prudent strategy to extend and enhance existing systems selectively may pay off equally well, and provide faster results. This is a recurring experience in the systems community: About a decade ago, the upcoming “multicore challenge” was supposed to render existing system-software designs unviable, and new kernel designs were deemed necessary [Ba09; Bo08; WA09]. It later turned out that existing system software could be scaled-up almost equally well by a systematic examination of their bottlenecks, which then could be fixed by employing standard techniques of parallel programming combined with *a few* novel abstractions [Bo10]. In a similar run–analyse–fix–approach, we have shown that existing system software, such as Linux, might just be *good enough* for many more database use cases, given proper configuration and adaptation. Of course, this does not invalidate the ongoing research on novel operating systems customised for database engines. Instead, the lesson to be learned here is that studying the *actual* reasons behind noise observed with existing operating systems is important. Only if we can pinpoint and understand the root causes, can we think of innovative solutions to these problems.

The cure might come by foreigners. Basically all of the measures we applied have originally been introduced into Linux to improve determinism and worst-case latencies not in database engines, but for the domain of embedded real-time systems: SCHED_FIFO, CPU shields, interrupt redirection, and PREEMPT_RT were developed and introduced to make Linux a suitable platform for mixed-criticality [BD17; Ve07] workloads in the real-time domain. This is also the main motivation behind partitioning hypervisors, such as Jailhouse. In our understanding, multi-tenant database engines that need to provide isolation and a guaranteed quality of service could (and should) be considered as (soft) real-time systems. Hence, operating-system solutions originally developed for the real-time domain might be a promising solution vector for the development of time-critical database systems. This underlines the necessity to rejoin the database and system communities.

Many challenges remain. Even though our results are promising, it should be pointed out that we eradicated only the CPU noise, and deliberately ignored I/O noise. Compared to the (narrow) case of a pure in-memory database, disk-based or otherwise I/O-intensive database stacks can be treated using the same measures as employed in this paper, but would face different challenges. In fact, most of Linux’s built-in real-time measures do not interact well with the I/O subsystem without further ado; it is often assumed that

the time-critical part can be decoupled from all I/O activities. While we have not yet examined disk-based database engines in this respect, we expect this to become a larger challenge that probably requires more invasive changes to the existing software stack. Direct I/O [Pe14] might be a promising way to approach this. Likewise, low-latency [Le19] or deterministic [SC03; Ya15] I/O scheduling have received a fair amount of attention outside the database community. External input via networks must consider additional stochastic parameters (e.g. unpredictable arrival times of data packets) that add to the complexity of the investigation. Real-time [KW05], or time-sensitive networking, and userland-based low-latency interaction with networking hardware can also be applied in database use-cases, albeit details must be left to future work.

What's next. Consider, as a specific and current example, how in-memory database engines can be extended with disk support – as, for instance, happens in the extension of Hyper to Umbra, where the authors propose to use SSDs for storage [NF20]. Especially parallel combinations of multiple SSDs promise RAM-like access performance.

However, parallel SSDs must be managed and driven. Database engines frequently aim at controlling block devices (at least to schedule access) from userland, since they have more complete usage pattern information than the OS. Yet this approach inevitably suffers from (at least) the amount of noise we have observed in our measurements, and advanced functionalities like RAID require substantial engineering effort. Operating systems provide such services as a commodity, but lack integration with the database query optimiser and its cost model. Additionally, operating systems are commonly optimised for throughput, so considerable tail latencies can be expected without adaptations. However, we are optimistic that moderate extensions of existing kernel mechanisms will combine the benefits of already existing infrastructure with little noise. This is important since increased determinism is beneficial to finding optimal query plans.

For all of the challenges listed above, we are optimistic that the required changes will be comparatively small compared to developing a new operating system from scratch.

6 Conclusion

We have shown that proper use of standard mechanisms of full-featured OSes can achieve database query latencies comparable to running an in-memory database engine directly on raw hardware. We reach a point where measuring time becomes the largest source of noise. By addressing challenges beyond CPU noise, we plan to bridge to the domain of real-time systems, and leverage techniques established for mixed-criticality systems which we apply to the database domain. After all, the underlying ideas match our scenario: One workload (the database engine) is to be shielded, without impairing the other workloads. We are confident that the respective research communities will enjoy many mutual benefits.

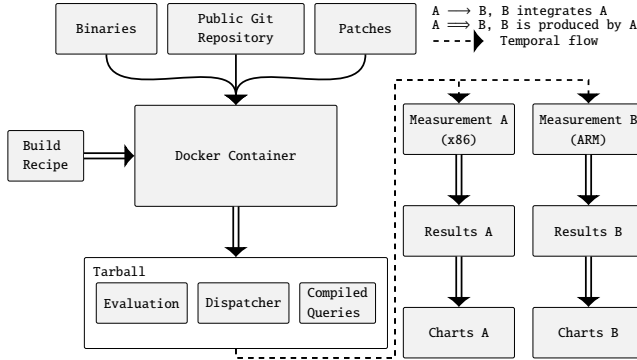


Fig. 10: Components and workflow of the reproduction package.

7 Appendix: Reproduction Package

Our publicly available reproduction package is based on Docker¹⁹. The process is illustrated in Figure 10: A Docker build recipe produces the docker container, and scripts that run therein produce a tarball with executables for all measurements in this paper. By transferring this tarball to a target, the experiments can be automatically executed, and charts generated.

We use binary sources for distribution-level software, and build other components (DBToaster, embedded compilers, RTEMS board support packages, ...) from source using the latest released state, augmented with local patches, to address issues found during this work that relate to the RTEMS kernel, the Jailhouse hypervisor, the GNU C compiler, and DBToaster itself (we include patches as explicit diff files to make any deviations from upstream sources explicit without relying on git history inspection). Additionally, we do not rely on the long-term availability of external sources by providing a pre-built Docker image. It contains all sources and dependencies, and enables re-building the exact same binaries from source that we use for the measurements (of course, our peers may choose to build the Docker image from scratch, depending on the latest binaries).

Finally, we provide all raw measurement results for all system combinations considered in the paper, and all post-processing scripts to evaluate and visualise the data.

Acknowledgements. This work was supported by the iDev40 project and the German Research Council (DFG) under grant no. LO 1719/3-1. The information and results set out in this publication are those of the authors and do not necessarily reflect the opinion of the ECSEL Joint Undertaking. The iDev40 project has received funding from the ECSEL Joint Undertaking (JU) under grant no. 783163. The JU receives support from the European Union’s Horizon 2020 research and innovation programme. It is co-funded by the consortium members, grants from Austria, Germany, Belgium, Italy, Spain and Romania. We thank the DBToaster team, and Jan Kiszka for guidance on difficult technical issues related to Jailhouse on x86 systems.

¹⁹ Available online from <https://github.com/1fd/btw2021>.

References

- [AP17] Arulraj, J.; Pavlo, A.: How to Build a Non-Volatile Memory Database Management System. In: Proceedings of the 2017 ACM SIGMOD International Conference on Management of Data. SIGMOD'17, pp. 1753–1758, 2017.
- [APD15] Arulraj, J.; Pavlo, A.; Dulloor, S. R.: Let's Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems. In: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data. SIGMOD'15, pp. 707–722, 2015.
- [Ar09] Armbrust, M.; Fox, A.; Griffith, R.; Joseph, A. D.; Katz, R.; Konwinski, A.; Lee, G.; Patterson, D.; Rabkin, A.; Stoica, I.; Zaharia, M.: Above the Clouds: A Berkeley View of Cloud Computing, tech. rep., University of California at Berkeley, 2009.
- [Au08] Aulbach, S.; Grust, T.; Jacobs, D.; Kemper, A.; Rittinger, J.: Multi-Tenant Databases for Software as a Service: Schema-Mapping Techniques. In (Wang, J. T.-L., ed.): Proceedings of the ACM SIGMOD International Conference on Management of Data. SIGMOD'08, pp. 1195–1206, 2008.
- [Ba09] Baumann, A.; Barham, P.; Dagand, P.-E.; Harris, T.; Isaacs, R.; Peter, S.; Roscoe, T.; Schüpbach, A.; Singhanian, A.: The multikernel: A New OS Architecture For Scalable Multicore Systems. In: Proceedings of the 22nd ACM Symp. on Operating Systems Principles. SOSP'09, pp. 29–44, 2009.
- [BD17] Burns, A.; Davis, R. I.: A Survey of Research into Mixed Criticality Systems. ACM Comput. Surv. 50/6, 2017.
- [Bi20] Bittman, D.; Alvaro, P.; Mehra, P.; Long, D. D. E.; Miller, E. L.: Twizzler: a Data-Centric OS for Non-Volatile Memory. In: 2020 USENIX Annual Technical Conference. USENIX ATC'20, pp. 65–80, 2020.
- [BL01] Buchmann, A. P.; Liebig, C.: Distributed, Object-Oriented, Active, Real-Time DBMS: We Want It All - Do We Need Them (at) All? In: Annual Reviews in Control, pp. 147–155, 2001.
- [Bo08] Boyd-Wickizer, S.; Chen, H.; Chen, R.; Mao, Y.; Kaashoek, F.; Morris, R.; Pesterev, A.; Stein, L.; Wu, M.; Dai, Y.; Zhang, Y.; Zhang, Z.: Corey: An Operating System for Many Cores. In: 8th Symposium on Operating System Design and Implementation. OSDI'08, pp. 43–57, 2008.
- [Bo10] Boyd-Wickizer, S.; Clements, A. T.; Mao, Y.; Pesterev, A.; Kaashoek, M. F.; Morris, R.; Zeldovich, N.: An Analysis of Linux Scalability to Many Cores. In: 9th Symposium on Operating System Design and Implementation. OSDI'10, 2010.
- [Br15] Bratterud, A.; Walla, A.-A.; Haugerud, H.; Engelstad, P. E.; Begnum, K.: IncludeOS: A Minimal, Resource Efficient Unikernel For Cloud Services. In: 2015 IEEE 7th International Conference on Cloud Computing Technology and Science. CloudCom'15, pp. 250–257, 2015.

- [BS14] Bloom, G.; Sherrill, J.: Scheduling and Thread Management with RTEMS. SIGBED Rev. 11/1, pp. 20–25, 2014.
- [Bu05] Buchmann, A. P.: Real-Time Databases. In (Rivero, L. C.; Doorn, J. H.; Ferragine, V. E., eds.): Encyclopedia of Database Technologies and Applications. Idea Group, pp. 524–529, 2005.
- [Ca20] Cafarella, M. J.; DeWitt, D. J.; Gadepally, V.; Kepner, J.; Kozyrakis, C.; Kraska, T.; Stonebraker, M.; Zaharia, M.: DBOS: A Proposal for a Data-Centric Operating System. CoRR abs/2007.11112/, 2020.
- [Ch18] Chandramouli, B.; Prasaad, G.; Kossmann, D.; Levandoski, J.; Hunter, J.; Barnett, M.: FASTER: A Concurrent Key-Value Store with In-Place Updates. In: Proceedings of the 2018 International Conference on Management of Data. SIGMOD’18, pp. 275–290, 2018.
- [DAM17] Dominico, S.; de Almeida, E. C.; Meira, J. A.: A PetriNet Mechanism for OLAP in NUMA. In: Proceedings of the 13th International Workshop on Data Management on New Hardware. DaMoN’17, pp. 1–4, 2017.
- [DB13] Dean, J.; Barroso, L. A.: The Tail at Scale. Commun. ACM 56/2, pp. 74–80, 2013.
- [Do18] Dominico, S.; de Almeida, E. C.; Meira, J. A.; Alves, M. A. Z.: An Elastic Multi-Core Allocation Mechanism for Database Systems. In: 2018 IEEE 34th International Conference on Data Engineering. ICDE’18, pp. 473–484, 2018.
- [Fa17] Faerber, F.; Kemper, A.; Larson, P.-Å.; Levandoski, J. J.; Neumann, T.; Pavlo, A.: Main Memory Database Systems. Found. Trends Databases 8/1-2, pp. 1–130, 2017.
- [Gi13] Giceva, J.; Salomie, T.-I.; Schüpbach, A.; Alonso, G.; Roscoe, T.: COD: Database / Operating System Co-Design. In: Sixth Biennial Conference on Innovative Data Systems Research, Online Proceedings. CIDR’13, 2013.
- [Gi19] Giceva, J.: Operating System Support for Data Management on Modern Hardware. IEEE Data Eng. Bull. 42/1, pp. 36–48, 2019.
- [Gr78] Gray, J. N.: Notes on Data Base Operating Systems. In: Operating Systems: An Advanced Course. Springer Berlin Heidelberg, pp. 393–481, 1978.
- [GS92] Garcia-Molina, H.; Salem, K.: Main Memory Database Systems: An Overview. IEEE Trans. on Knowl. and Data Eng. 4/6, pp. 509–516, 1992.
- [GTS20] Götze, P.; Tharanatha, A. K.; Sattler, K.-U.: Data Structure Primitives on Persistent Memory: An Evaluation. In: Proceedings of the 16th International Workshop on Data Management on New Hardware. DaMoN’20, 2020.
- [In15] Intel Corporation: Improving Real-Time Performance by Utilizing Cache Allocation Technology, <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/cache-allocation-technology-white-paper.pdf> (last accessed February 2021), 2015.

- [Ki15] Kiefer, T.; Schön, H.; Habich, D.; Lehner, W.: A Query, a Minute: Evaluating Performance Isolation in Cloud Databases. In: Performance Characterization and Benchmarking. Traditional to Big Data, pp. 173–187, 2015.
- [KLT16] Koch, C.; Lupei, D.; Tannen, V.: Incremental View Maintenance For Collection Programming. In: Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. PODS’16, pp. 75–90, 2016.
- [Ko10] Koch, C.: Incremental Query Evaluation in a Ring of Databases. In: Proceedings of the 29th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems. PODS’10, pp. 87–98, 2010.
- [Ko13] Koch, C.; Ahmad, Y.; Kennedy, O. A.; Nikolic, M.; Nötzli, A.; Lupei, D.; Shaikhha, A.: DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views, tech. rep., EPFL, 2013.
- [Ko14] Koch, C.; Ahmad, Y.; Kennedy, O.; Nikolic, M.; Nötzli, A.; Lupei, D.; Shaikhha, A.: DBToaster: Higher-Order Delta Processing for Dynamic, Frequently Fresh Views. VLDB J. 23/2, pp. 253–278, 2014.
- [KSL13] Kiefer, T.; Schlegel, B.; Lehner, W.: Experimental Evaluation of NUMA Effects on Database Management Systems. In: Datenbanksysteme für Business, Technologie und Web (BTW) 2025. BTW’13, pp. 185–204, 2013.
- [KW05] Kiszka, J.; Wagner, B.: RTnet-A Flexible Hard Real-Time Networking Framework. In: 2005 IEEE Conference on Emerging Technologies and Factory Automation. Vol. 1, IEEE, pp. 456–464, 2005.
- [Le19] Lee, G.; Shin, S.; Song, W.; Ham, T. J.; Lee, J. W.; Jeong, J.: Asynchronous I/O Stack: A Low-latency Kernel I/O Stack for Ultra-Low Latency SSDs. In: USENIX Annual Technical Conference. USENIX ATC’19, pp. 603–616, 2019.
- [Le20] Lersch, L.; Schreter, I.; Oukid, I.; Lehner, W.: Enabling Low Tail Latency on Multicore Key-Value Stores. Proc. VLDB Endow. 13/7, pp. 1091–1104, 2020.
- [Ma10] Mauerer, W.: Professional Linux Kernel Architecture. John Wiley & Sons, 2010.
- [Ma13] Madhavapeddy, A.; Mortier, R.; Rotsos, C.; Scott, D.; Singh, B.; Gazagnaire, T.; Smith, S.; Hand, S.; Crowcroft, J.: Unikernels: Library Operating Systems for the Cloud. SIGARCH Comput. Archit. News 41/1, pp. 461–472, 2013.
- [MKN12] Mühe, H.; Kemper, A.; Neumann, T.: The Mainframe Strikes Back: Elastic Multi-Tenancy Using Main Memory Database Systems On a Many-Core Server. In: Proceedings of the 15th International Conference on Extending Database Technology. EDBT’12, pp. 578–581, 2012.
- [MS19] Müller, M.; Spinczyk, O.: MxKernel: Rethinking Operating System Architecture for Many-Core Hardware. In: 9th Workshop on Systems for Multi-core and Heterogenous Architectures. EuroSys’19, 2019.

- [Mü20] Mühlig, J.; Müller, M.; Spincyk, O.; Teubner, J.: mxkernel: A Novel System Software Stack for Data Processing on Modern Hardware. *Datenbank-Spektrum* 20/3, pp. 223–230, 2020.
- [Na13] Narasayya, V.; Das, S.; Syamala, M.; Chandramouli, B.; Chaudhuri, S.: SQLVM: Performance Isolation in Multi-Tenant Relational Database-as-a-Service. In: 6th Biennial Conf. on Innovative Data Systems Research. *CIDR'13*, 2013.
- [NDK16] Nikolic, M.; Dashti, M.; Koch, C.: How to Win a Hot Dog Eating Contest: Distributed Incremental View Maintenance with Batch Updates. In: *Proceedings of the 2016 International Conference on Management of Data. SIGMOD'16*, pp. 511–526, 2016.
- [NF20] Neumann, T.; Freitag, M. J.: Umbra: A Disk-Based System with In-Memory Performance. In: 10th Conference on Innovative Data Systems Research, *Online Proceedings. CIDR'20*, 2020.
- [No18] Noll, S.; Teubner, J.; May, N.; Böhm, A.: Accelerating Concurrent Workloads with CPU Cache Partitioning. In: *IEEE 34th International Conference on Data Engineering. ICDE'18*, pp. 437–448, 2018.
- [Pe14] Peter, S.; Li, J.; Zhang, I.; Ports, D. R. K.; Woos, D.; Krishnamurthy, A.; Anderson, T.; Roscoe, T.: Arrakis: The Operating System is the Control Plane. In: *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation. USENIX'14*, pp. 1–16, 2014.
- [PH90] Patterson, D. A.; Hennessy, J. L.: *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers Inc., 1990, ISBN: 1558800698.
- [Po12] Porobic, D.; Pandis, I.; Branco, M.; Tözün, P.; Ailamaki, A.: OLTP on Hardware Islands. *Proc. VLDB Endow.* 5/11, pp. 1447–1458, 2012.
- [Ra17] Ramsauer, R.; Kiszka, J.; Lohmann, D.; Mauerer, W.: Look Mum, no VM Exits! (Almost). In: *Proceedings of the 13th Annual Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPRT '17)*. 2017.
- [Re19] van Renen, A.; Vogel, L.; Leis, V.; Neumann, T.; Kemper, A.: Persistent Memory I/O Primitives. In: *Proceedings of the 15th International Workshop on Data Management on New Hardware. DaMoN'19*, 12:1–12:7, 2019.
- [RF18] Rehmann, K.-T.; Folkerts, E.: Performance of Containerized Database Management Systems. In: *Proceedings of the Workshop on Testing Database Systems. DBTest'18*, 2018.
- [SC03] Swaminathan, V.; Chakrabarty, K.: Energy-conscious, deterministic I/O device scheduling in hard real-time systems. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 22/7, pp. 847–858, 2003.
- [SDQ10] Schad, J.; Dittrich, J.; Quiané-Ruiz, J.-A.: Runtime Measurements in the Cloud: Observing, Analyzing, and Reducing Variance. *PVLDB Endow.* 3/1, pp. 460–471, 2010.

- [St07] Stonebraker, M.; Madden, S.; Abadi, D. J.; Harizopoulos, S.; Hachem, N.; Helland, P.: The End of an Architectural Era: (It's Time for a Complete Rewrite). In: Proceedings of the 33rd International Conference on Very Large Data Bases. VLDB'07, pp. 1150–1160, 2007.
- [Ve07] Vestal, S.: Preemptive Scheduling of Multi-Criticality Systems with Varying Degrees of Execution Time Assurance. In: Proceedings of the 28th IEEE International Real-Time Systems Symposium. RTSS'07, pp. 239–243, 2007.
- [WA09] Wentzlaff, D.; Agarwal, A.: Factored Operating Systems (fos): The Case for a Scalable Operating System for Multicores. ACM SIGOPS Operating Systems Review 43/2, pp. 76–85, 2009.
- [Xu13] Xu, Y.; Musgrave, Z.; Noble, B.; Bailey, M.: Bobtail: Avoiding Long Tails in the Cloud. In: 10th USENIX Symposium on Networked Systems Design and Implementation. NSDI'13, pp. 329–341, 2013.
- [Ya15] Yang, S.; Harter, T.; Agrawal, N.; Kowsalya, S. S.; Krishnamurthy, A.; Al-Kiswani, S.; Kaushik, R. T.; Arpaci-Dusseau, A. C.; Arpaci-Dusseau, R. H.: Split-Level I/O Scheduling. In: Proceedings of the 25th Symposium on Operating Systems Principles. SOSP'15, pp. 474–489, 2015.