

MLOS: An Infrastructure for Automated Software Performance Engineering

Carlo Curino, Neha Godwal, Brian Kroth, Sergiy Kuryata, Greg Lapinski, Siqi Liu, Slava Oks, Olga Poppe, Adam Smiechowski, Ed Thayer, Markus Weimer, Yiwen Zhu
first.last@microsoft.com

ABSTRACT

MLOS is a Data Science powered infrastructure and methodology to democratize and automate Software Performance Engineering. MLOS enables continuous, instance-level, robust, and trackable systems optimization.

1 INTRODUCTION

Software Performance Engineering (SPE) is hard because performance depends on the software and hardware environment where the system is run *and* what workload it is processing. This situation is even more difficult, since these influences are not static in time. In response to these challenges, real-world SPE has developed as an amorphous collection of experiments and tribal best practices aimed at tuning parameters and policies to satisfy non-functional performance requirements (e.g., latency, throughput, utilization, etc.). Consider as an example the design and tuning of a hash table. Implementing a working in-memory hash table is relatively simple. But tuning it requires a deep understanding of the workload, complex hardware capabilities and interactions, and local as well as global performance/resource trade-offs of the entire system—as shown in our experiments in §3.

Our key observation is that: while *business logic development* is a necessary and high-value activity, *manual tuning of parameters/heuristics* is an often low-yield (yet important) activity for engineers. The key reasons for the current situation are: 1) lack of integrated tools to explore a large parameter space by running, and tracking experiments, 2) a need to continuously re-tune systems as hardware/software/workload (hw/sw/wl) context evolve, and 3) the manual nature of the process, that forces one-size-fits-all configurations.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DEEM'20, June 14, 2020, Portland, OR, USA

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-8023-2/20/06...\$15.00
<https://doi.org/10.1145/3399579.3399927>

As a consequence, instead of being a pervasive and continuous activity (like unit testing), SPE is relegated to elite teams with local expertise and ad-hoc knowledge and tools. The cost of SPE can only be justified where performance is critical—Microsoft and Intel employ 100s of performance engineers that jointly tune core systems for every major hw/sw release. The alternative is accepting sub-optimal performance. Intel's performance engineers shared that, most system software is typically 30-40% away from what the hardware can achieve¹. This situation is exaggerated by increasingly diverse hardware available in the cloud and on the edge.

Thanks to advancements in data management and machine learning, a fundamental shift in how systems are built and tuned is possible and required. In short, we argue in favor of *melding Data Science into the practice of Software Engineering to make SPE mainstream*. Beyond the maturation of Data Science (DS), MLOS² is enabled by the move to the cloud. Cloud operations provide us with: workload observability, increased incentives to optimize software, and an unprecedented experimentation platform.

We present MLOS, an infrastructure and engineering practice aimed at systematizing and automating the performance tuning of a system via Data Science. With MLOS, we empower engineers with tooling for data-driven optimizations. This improves the end-to-end process of development and long term maintenance costs, increases documentation and reproducibility of performance tuning, and leads to more performant, efficient, well tuned systems—discussed in §2.

We observe that the SPE process closely mimics the DS process, and MLOS brings the rigor and tools of DS to it. Like SPE, DS relies on experimentation, and changing parameters as often as necessary. However, in existing systems software, many parameters are hard-coded and long build times prohibit rapid experimentation. We believe this is a major impediment to data-driven, automated SPE. So, one key architectural decision we make in MLOS is to separate the core system and the data-driven tuning of it. In §2.1 we describe the infrastructure in more detail, highlighting this fundamental architectural choice.

¹That is the typical boost that their specialized SPE team achieves when engaged to tune previously unoptimized software for a new hardware release.

²Short for Machine Learning [Optimized|Operating] Systems.

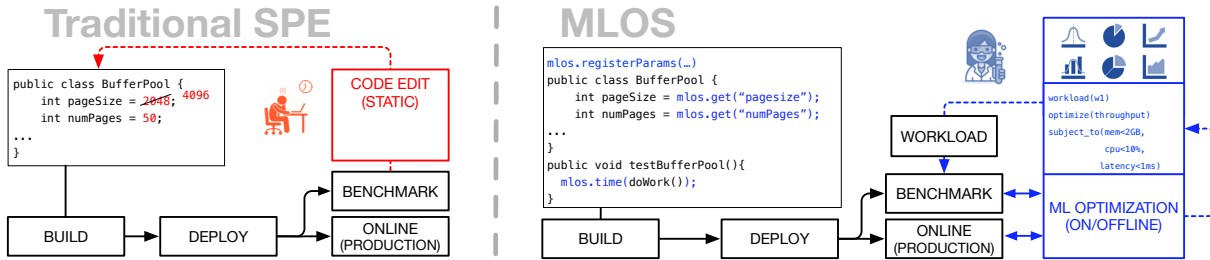


Figure 1: The MLOS Developer Experience

In building MLOS, we are faced with hard challenges: (1) Creating a “natural” development experience, enabling engineers to take advantage of advanced and rigorous DS processes. Annotations+Automation allow for a rich DS experience with minimal effort. (2) Introducing observability/tunability without disrupting the performance of the target system. We achieve this using a side-agent and channel to create low-cost communications with a DS experience. (3) Global optimization across thousands of parameters is an open problem. We cope with this by decomposing the problem (as business logic does using APIs) into microbenchmarks for sub-components and by introducing a Resource Performance Interface (RPI)—the SPE equivalent of an API.

We highlight the results of applying our initial MLOS implementation to tackle (1) and (2) by optimizing core data structures in SQL Server in §3.

At Microsoft, we are developing and using MLOS within SQL Server engineering and are in the process of open-sourcing all of the general purpose system components [2]. Furthermore, we are collaborating with several academic institutions with the intention to create an initial class and lab.

2 THE MLOS DEVELOPER EXPERIENCE

MLOS aims to democratize and automate SPE. Specifically, we leverage existing Data Science technologies, including: notebooks [1], model training [13, 20], tracking [24] and deployment [5]. We concentrate our net-new development on what we believe to be a key gap that prevents system developers from using DS for SPE efficiently: the lack of observability and tunability of low level system internals.

Figure 1 highlights the difference between traditional SPE and MLOS. For traditional SPE, the developer’s “tuning loop” involves: hard-coding “magic” constants (or at best exposing them as start-up parameters), recompiling/redeploying (or restarting) the system, running an experiment, and analyzing the results. This process is lengthy, manual, and error-prone.

By contrast, MLOS enables externally observable and dynamically tunable constants, and allows the developer to “drive” and track experiments in a fashion akin to a DS experience. The loop is faster and more carefully tracked (e.g.,

we record all experiment conditions, including OS/HW counters). Moreover, MLOS enables automated parameter tuning, allowing for continuous and dynamic re-tuning, as well as instance-level optimization (i.e., optimize for a specific hw/sw/wl context).

Our primary concern in designing MLOS is to create a *natural development experience*. The use of auto-parameters (the term we use to refer to DS-tuned system parameters) must be natural and easier than hard-coding a constant. The first step in this journey is providing APIs and annotation mechanisms that are native to the system language. In SQL Server (C++), we leverage C# language *Attributes* to declare certain parameters as tunable, and associate metrics to critical section of the component code for observability. As we extend MLOS to more languages, we will select mechanisms idiomatic to each language (e.g., Java annotations).

Through code-gen mechanisms (§ 2.1), each of these parameters and their associated metrics is made observable and tunable external to the core system without impacting the delicate performance inner-loop of the system—our second core challenge. This externalization allows us to provide an agile notebook-driven experience (similar to a standard DS workbench). This enables a software engineer with rudimentary DS skills to visualize, analyze, model, and optimize component parameters. We require the developer to provide (micro)benchmarks, but minimal input regarding the choice of optimization function (e.g., maximize the throughput of a component), subject to certain constraints (e.g., memory/cpu/latency caps) and the choice of workload(s).

In many cases simple AutoML or well chosen default optimization pipelines will likely provide a significant initial lift over the manual configuration. The reason is simple: we compare a static/global configuration of parameters to a one dynamically tuned for a given hw/sw/wl. Another value-add from MLOS is that the developer need only provide a few application level metrics specific to their component (e.g., timing of a critical section) and MLOS automatically gathers a large amount of contextual information (e.g., OS/HW counters). These are leveraged to gain additional insight and can help in modeling a system. Finally, in MLOS we follow DS best practices and integrate with tools such as MLFlow [24]

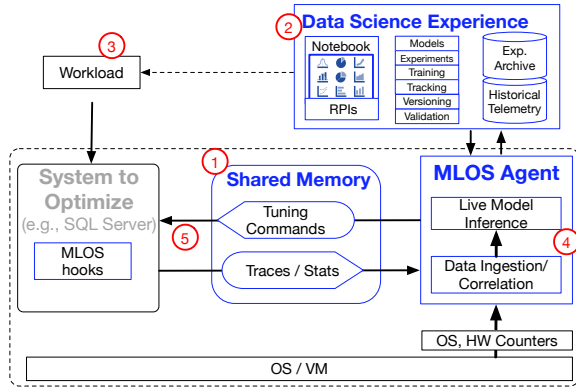


Figure 2: The MLOS Architecture

for versioning and tracking of all models/experiments. On its own this provides rigor and reproducibility and makes the SPE process continuous rather than a one-off optimization that grows stale over time. To this end we expect that in most cases the “Data Science” aspects of MLOS will deliver gains, more than its ML sophistication.

Resource Performance Interface (RPI) A developer can optimize either the end-to-end system (e.g., TPC-E on SQL Server) or a specific component (e.g., microbenchmark on spinlocks, as we do in §3). Locally optimizing individual components out of *context* may not align with end-to-end optimization since those components typically share and compete for resources in the larger system. For this reason we introduce the notion of a *Resource Performance Interface (RPI)*. Conceptually an RPI allows the developer to define the *context* of the component in the larger system in the form of resource and performance expectations. Practically, this describes an acceptable “envelope” of resources for a given component stressed by a chosen workload. Note that the values here may be initially specified by the developer, or learned from an existing system (e.g., during build-test runs). RPIs provide the grounding for systematic performance regression testing, at the component level. This is key to decompose the performance behavior of a large system, and it is the moral equivalent of an API for the non-functional performance aspects of a system. Importantly, the RPI is *not* specified as part of the system code, but as part of the DS experience. The rationale for this is that the same piece of business logic, might be used in different parts of a system or in different environments and could be optimized differently. More precisely this means that the same API could be associated with multiple RPIs depending on the context of usage. Next, we describe how we deliver this developer experience.

2.1 The MLOS Architecture

A fundamental choice in MLOS is to enable DS-driven optimization without fundamentally modifying the underlying system. The practical motivation comes from our initial target

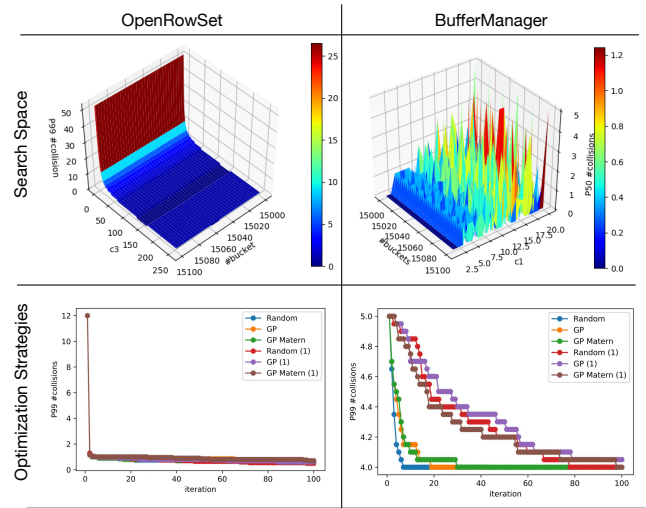


Figure 3: Tuning of two hash-tables in SQL Server

case: SQL Server. This massive codebase is hard to modify, and the size of it makes the build process very costly, slowing the experimental SPE loop. Next we describe the mechanics of how we decouple the system from the modeling and tuning of it. The key challenge here is to provide observability and tunability with minimal impact to the underlying system performance—a performance variant of the Socratic oath.

MLOS is depicted in Figure 2 and operates as follows:

- (1) Given an annotated code base (per §2), MLOS performs code gen of: a) hooks in the system, b) a low latency shared memory communication channel, and c) the MLOS Agent (a daemon operating side-by-side the main system). This is a high-performance mechanism to externalize all tuning.
- (2) A Data Science experience is enabled on top of the application metrics and OS/HW counter data collected. Developers can leverage popular data science tools for visualization, modeling, and tracking of experiments.
- (3) The Data Science experience also provides generic APIs to drive system-specific workloads, allowing us to focus on collecting data points of interest interactively or automatically during CI/CD and complement observations of live systems.
- (4) Models, Rules, and Optimizations are deployed via the DS experience into the MLOS Agent, where they are hosted for online inferencing based on live and contextual conditions (e.g., observed load on the server, internal system metrics).
- (5) Finally commands are sent via the shared memory channel and enacted by the MLOS hooks (e.g., updated the value for the maximum spin count of a spinlock).

This approach creates the independence required for fast experimentation, and reduces to a minimum the impact on the system inner-loop. Clearly not all parameters are easily tuned dynamically, as some would incur costly re-initialization operations. As we gain experience with MLOS in practice we will develop simple development patterns and libraries of




Figure 4: HW/OS Counters to tune in a resource-aware fashion

components that are amenable to dynamic tuning. This is why we focus our first application of MLOS to core datastructures and functions in SQL Server: SQLOS [18].

3 EVALUATION

Our first experiments focus on tuning the parameters of hashtables in SQL Server. Figure 3 shows the effects of tuning two instances (OpenRowSet and BufferManager). We observe: 1) DS-driven optimization can improve against SQL Server highly tuned parameters by 20% to 90% (the initial point in the strategy graphs), 2) the optimization surface for the same component may vary with the workload (smooth for OpenRowSet and very jagged for BufferManager) and 3) Random Search (RS) performs very competitively with Bayesian Optimization (BO), shown using Gaussian Processes (GP), GP with Mattern 3/2 kernels, and 4) optimizing multiple parameters at a time can be faster than one at a time (lines marked as (1) in graph). In other experiments (omitted due to space) BO was more efficient, indicating that the choice of optimization mechanism is non-trivial and our agile DS experience is indeed important.

Next we showcase how the ability of MLOS to gather HW/OS counters during experiments allows us to adapt to different operating conditions. Leveraging a similar Hashtable experiment Figure 4 shows that a larger hashtable (more buckets, hence more memory), reduces collisions (app metric) leading to better latency. Interestingly, up to 5MB of allocation this also offers a good reduction of cpu-cycles and cache-misses (HW metrics), but this effect disappears beyond that. This indicates that at least 5MB allocation is advisable to save CPU resources in most conditions, but that past that the tradeoff of collisions/memory dominates. Importantly, automatic collection of HW/OS counters make it easy to observe these tradeoffs, and the MLOS optimizer allows the developer to focus on the desired objective (e.g., a tradeoff of resources and performance) abstracting away the parameter fiddling. Finally, we note that MLOS can perform this tuning continuously adapting to changing hw/sw/wl conditions.

Our last experiment showcases how the workload may influence optimal parameters choice. Figure 5 shows the performance of a spinlock when changing the maximum spinning before backoff. We show 7 workloads each comprised




Figure 5: Optimal spinlock polling length varies by workload.

of several threads performing little work, and one performing a larger and larger number of operations under the lock. Subtle changes in the workload (or hardware platform) can substantially affect the optimal choice of parameters.

Discussion These experiments show great promise for component-level optimizations. A key open challenge: how to tackle the impractical parameter space for global optimization. Our current practical solution is to leverage the power of RPIs combined with developer intuition. More research in this space is needed.

4 RELATED WORK

MLOS and Software 2.0 [3] are related, but not the same. MLOS does not make a distinction between developers who write code ("Developer 1.0") and those that curate data and train models ("Developer 2.0"). Instead, MLOS augments systems developers with enough data science tools and methodology to meet the challenges of the day. Similarly, MLOS isn't geared at problems that can only be solved by replacing functions with models (e.g. image recognition). Instead, it fuses traditional solutions (e.g. OS/HW measurement of benchmarks) to problems solvable with them (e.g. hashtable optimization) with data science to organize and optimize them for the context in which they are applied. Recent work [8, 12, 15, 17] on replacing sub-components of a DBMS with models is incredibly promising, but as of the time of this writing still brittle for our production applications.

Another key area of related work, is tools that tune external parameters of a system [9, 19, 22, 25]. Closer points of comparison are approaches such as [14, 16, 21, 23] that leverage ML to guide/improve existing query optimizers. All of these efforts are complementary to MLOS, which focuses on making system internals observable and tunable. We are in fact, considering to port some of them onto the MLOS infrastructure.

Past efforts have also tried to address this problem at a library [6, 10] or compiler level [4, 7, 11]. In some ways these are early applications of MLOS style optimization for those areas. While complementary, though don't focus on exposing the DS for the developer to improve with as well.

5 CONCLUSION

Software Performance Engineering (SPE) is a very valuable but costly activity. With MLOS we fill an infrastructure gap that will enable developers to leverage the tools and practices of Data Science (DS).

REFERENCES

- [1] ADS. <https://github.com/Microsoft/azuredatastudio>.
- [2] MLOS. <https://github.com/microsoft/MLOS>.
- [3] Andrej Karpathy. Software 2.0. <https://medium.com/@karpathy/software-2-0-a64152b37c35>, 2017.
- [4] A. H. Ashouri, W. Killian, J. Cavazos, G. Palermo, and C. Silvano. A survey on compiler autotuning using machine learning. *ACM Computing Surveys (CSUR)*, 51(5):1–42, 2018.
- [5] J. Bai, F. Lu, K. Zhang, et al. Onnx: Open neural network exchange. <https://github.com/onnx/onnx>, 2019.
- [6] V. Carbune, T. Coppey, A. Daryin, T. Deselaers, N. Sarda, and J. Yagnik. Smartchoices: Hybridizing programming and machine learning. *arXiv preprint arXiv:1810.00619*, 2018.
- [7] D. Chen, T. Moseley, and D. X. Li. Autofdo: Automatic feedback-directed optimization for warehouse-scale applications. In *2016 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 12–23. IEEE, 2016.
- [8] J. Ding, U. F. Minhas, H. Zhang, Y. Li, C. Wang, B. Chandramouli, J. Gehrke, D. Kossmann, and D. Lomet. Alex: an updatable adaptive learned index. *arXiv preprint arXiv:1905.08898*, 2019.
- [9] S. Duan, V. Thummala, and S. Babu. Tuning database configuration parameters with ituned. *Proceedings of the VLDB Endowment*, 2(1):1246–1257, 2009.
- [10] J. Eastep, D. Wingate, and A. Agarwal. Smart data structures: an online machine learning approach to multicore data structures. In *Proceedings of the 8th ACM international conference on Autonomic computing*, pages 11–20, 2011.
- [11] G. Fursin, Y. Kashnikov, A. W. Memon, Z. Chamski, O. Temam, M. Namolaru, E. Yom-Tov, B. Mendelson, A. Zaks, E. Courtois, et al. Milepost gcc: Machine learning enabled self-tuning compiler. *International journal of parallel programming*, 39(3):296–327, 2011.
- [12] S. Idreos, K. Zoumpatianos, B. Hentschel, M. S. Kester, and D. Guo. The data calculator: Data structure design and cost synthesis from first principles and learned cost models. In *Proceedings of the 2018 International Conference on Management of Data*, pages 535–550.
- [13] M. Interlandi, S. Matushevych, S. Amizadeh, S. Zahirazami, and M. Weimer. Machine learning at microsoft with ml.net. 2018.
- [14] A. Jindal, H. Patel, A. Roy, S. Qiao, Z. Yin, R. Sen, and S. Krishnan. Peregrine: Workload optimization for cloud query engines. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 416–427, 2019.
- [15] T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, pages 489–504, 2018.
- [16] R. Marcus, P. Negi, H. Mao, N. Tatbul, M. Alizadeh, and T. Kraska. Bao: Learning to steer query optimizers. *arXiv:2004.03814*, 2020.
- [17] R. Marcus, P. Negi, H. Mao, C. Zhang, M. Alizadeh, T. Kraska, O. Papaemmanouil, and N. Tatbul. Neo: A learned query optimizer. *Proceedings of the VLDB Endowment*, 12(11):1705–1718, 2019.
- [18] S. Oks. Platform layer for sql server. <https://docs.microsoft.com/en-us/archive/blogs/slavao/platform-layer-for-sql-server>.
- [19] A. Pavlo, M. Butrovich, A. Joshi, L. Ma, P. Menon, D. Van Aken, L. Lee, and R. Salakhutdinov. External vs. internal: An essay on machine learning agents for autonomous database management systems. *IEEE Data Engineering*, 11:1910–1913, 2019.
- [20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
- [21] T. Siddiqui, A. Jindal, S. Qiao, H. Patel, et al. Cost models for big data query processing: Learning, retrofitting, and our findings. *arXiv preprint arXiv:2002.12393*, 2020.
- [22] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang. Automatic database management system tuning through large-scale machine learning. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1009–1024, 2017.
- [23] C. Wu, A. Jindal, S. Amizadeh, H. Patel, W. Le, S. Qiao, and S. Rao. Towards a learning optimizer for shared clouds. *Proceedings of the VLDB Endowment*, 12(3):210–222, 2018.
- [24] M. Zaharia, A. Chen, A. Davidson, A. Ghodsi, S. A. Hong, A. Konwinski, S. Murching, T. Nykodym, P. Ogilvie, M. Parkhe, et al. Accelerating the machine learning lifecycle with mlflow. *IEEE Data Eng. Bull.*, 41(4):39–45, 2018.
- [25] Y. Zhu, J. Liu, M. Guo, Y. Bao, W. Ma, Z. Liu, K. Song, and Y. Yang. Bestconfig: tapping the performance potential of systems via automatic configuration tuning. In *Proceedings of the 2017 Symposium on Cloud Computing*, pages 338–350, 2017.