

Buffer Pool Aware Query Scheduling via Deep Reinforcement Learning (Extended Abstract)

Chi Zhang¹, Ryan Marcus², Anat Kleiman¹, Olga Papaemmanouil¹
¹Brandeis University ²MIT CSAIL and Intel Labs
¹{chizhang, akleiman, opapaemm}@brandeis.edu, ²ryanmarcus@csail.mit.edu

ABSTRACT

In this extended abstract, we propose a new technique for query scheduling with the explicit goal of reducing disk reads and thus implicitly increasing query performance. We introduce SmartQueue, a learned scheduler that leverages overlapping data reads among incoming queries and learns a scheduling strategy that improves cache hits. SmartQueue relies on deep reinforcement learning to produce workload-specific scheduling strategies that focus on long-term performance benefits while being adaptive to previously-unseen data access patterns. We present results from a proof-of-concept prototype, demonstrating that learned schedulers can offer significant performance improvements over hand-crafted scheduling heuristics. Ultimately, we make the case that this is a promising research direction in the intersection of machine learning and databases.

1. INTRODUCTION

Query scheduling, the problem of deciding which of a set of queued queries to execute next, is an important and challenging task in modern database systems. Query scheduling can have a significant impact on query performance and resource utilization while it may need to account for a wide number of considerations, such as cached data sets, available resources (e.g., memory), per-query performance goals, query prioritization, or inter-query dependencies (e.g., correlated data access patterns).

In this work, we attempt to address the query scheduling problem by leveraging overlapping data access requests. Smart query scheduling policies can take advantage of such overlaps, allowing queries to share cached data, whereas naive scheduling policies may induce unnecessary disk reads. For example, consider three queries q_1, q_2, q_3 which need to read disk blocks (b_1, b_2) , (b_4, b_5) , and (b_2, b_3) respectively. If the DBMS's buffer pool (i.e., the component of the database engine that caches data blocks) can only cache two blocks at once, executing the queries in the order of $[q_1, q_2, q_3]$ will result in reading 6 blocks from disk. However, if the queries are

executing in the order $[q_1, q_3, q_2]$, then only 5 blocks will be read from disk, as q_3 will use the cached b_2 . Since buffer pool hits can be orders of magnitude faster than cache misses, such savings could be substantial.

In reality, designing a query scheduler that is aware of the current buffer pool is a complex task. First, the exact data block read set of a query is not known ahead of time, and is dependent on data and query plan parameters (e.g., index lookups). Second, a smart scheduler must balance short-term rewards (e.g., executing a query that will take advantage of the current buffer state) against long-term strategy (e.g., selecting queries that keep the most important blocks cached). One could imagine many simple heuristics, such as greedily selecting the next query with the highest expected buffer usage, to solve this problem. However, a hand-designed policy to handle the complexity of the entire problem, including different buffer sizes, shifting query workloads, heterogeneous data types (e.g., index files vs base relations), and balancing short-term gains against long-term strategy is much more difficult to conceive.

Here, we showcase a prototype of SmartQueue, a deep reinforcement learning (DRL) system that automatically learns to maximize buffer hits in an adaptive fashion. Given a set of queued queries, SmartQueue combines a simple representation of the database's buffer state, the expected reads of queries, and deep Q-learning model to order queued queries in a way that garners long-term increases in buffer hits. SmartQueue is fully learned, and requires minimal tuning. SmartQueue custom-tailors itself to the user's queries and database, and learns policies that are significantly better than naive or simple heuristics. In terms of integrating SmartQueue into an existing DBMS, our prototype only requires access to the execution plan for each incoming query (to assess likely reads) and the current state of the DBMS buffer pool (i.e., its cached data blocks).

We present our system model and formalized our learning task in Section 2. We present preliminary experimental results from a proof-of-concept prototype implementation in Section 3, related work in Section 4, and in Section 5 we highlight directions for future work.

2. THE SMARTQUEUE MODEL

SmartQueue is a learned query scheduler that automatically learns how to order the execution of queries to minimize disk access requests. The core of SmartQueue includes a deep reinforcement learning (DRL) agent [3] that learns a query scheduling policy through continuous interactions with its environment, i.e., the database and the incoming

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well allowing derivative works, provided that you attribute the original work to the author(s) and AIDB 2020. *2nd International Workshop on Applied AI for Database Systems and Applications (AIDB'20)*, August 31, 2020, Tokyo, Japan.

queries. This DRL agent is not a static model, instead it *continuously* learns from its past scheduling decisions and *adapts* to new data access and caching patterns. Furthermore, as we discuss below, using a DRL model allows us to define a reward function and scheduling policy that captures long-term benefits vs short-term gains in disk access.

Our system model is depicted in Figure 1. Incoming user queries are placed into an execution queue and SmartQueue decides their order of execution. For each query execution, the database collects the required *data blocks* of each input base relation, where a data block is the smallest data unit used by the database engine. Data blocks requests are first resolved by the buffer pool. Blocks found in the buffer (*buffer hits*) are returned for processing while the rest of the blocks (*buffer misses*) are read from disk and placed into the buffer pool (after possible block evictions). Higher buffer hit rates (and hence lower disk access rates) can enormously impact query execution times but require strategic query scheduling, as execution ordering affects the data blocks cached in the buffer pool.

One tempting solution to address this challenge could involve a greedy scheduler which executes the query that will re-use the maximum number of cached data blocks. While this simple approach would yield short term benefits, it ignores the long-term impact of each choice. Specifically, while the next query for execution will maximally utilize the buffer pool contents, it will also lead to newly cached data blocks, which will affect future queries. A greedy approach fails to identify whether these new cached blocks could be of any benefit to the unscheduled yet queries.

SmartQueue addresses this problem by training a deep reinforcement learning agent to make scheduling decisions that maximize long term benefits. Specifically, it uses a model that simultaneously estimates and tries to improve a weighted average between short-term buffer hits and the long-term impact of query scheduling choices. In the next paragraphs, we discuss the details of our approach: (a) the input features vector that capture data access requests (*Query Bitmap*) and buffer state (*Buffer Bitmap*), and (b) the formalized DRL task.

Buffer Bitmap. One input to the DRL model is the state of the buffer pool, namely which blocks are currently cached in memory. Buffer state B is represented by a bitmap where rows represent base relations and columns represent data blocks. The (i, j) entry is set to 1 if the j -th block of relation i is cached in the buffer pool and is set to zero otherwise. Since the number of blocks of any given relation can be very high and different for each relation, each row vector F_i is downsized by calculating a simple moving average over the number of its blocks entries. Specifically D_i is the downsized row of a relation i and F_i is the full size row, we have:

$$B_{ij} = \lfloor \frac{|F_i|}{|D_i|} \rfloor \times \sum_{k=j \times \lfloor \frac{|D_i|}{|F_i|} \rfloor}^{(j+1) \times \lfloor \frac{|D_i|}{|F_i|} \rfloor} F_{ik} \quad (1)$$

Query Vector. The second input to the DLR model is the data block requests of each query in the queue. Specifically, given a query q , we generate a vector that indicates the data blocks to be accessed by q for each base relation in the database. To implement this, SmartQueue collects the query plan of q , and approximates the probability of

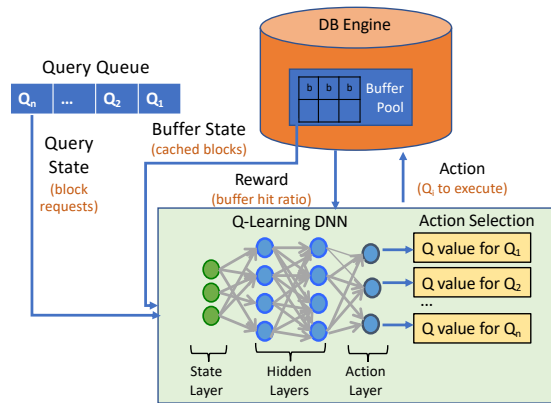


Figure 1: SmartQueue's system model

each table's data block being accessed. Our approach handle requests of index file and base relations similarly, as both type of blocks will be cached into the buffer pool. The query vector is downsized in the same way as the buffer bitmap.

Full table scans for a base relation i indicate that all data blocks of the given relation will be accessed, and therefore each cell of the i -th row vector has the value of 1. For indexed table scans, we calculate the number of tuples to be accessed based on the selectivity of the index scan. If the index scan is feeding a loop-based operator (i.e., nested loop join) the selectivity is adapted accordingly to account for any iterations over the relation. We assume the relation is uniformly stored across data blocks and therefore, if $x\%$ tuples of a base relation are to be selected from an indexed operation, we set the access probability of each data block of the relation to $x\%$. Similarly, we assume that the indexed operation reads $x\%$ of the index's blocks. We note that much more sophisticated probabilistic models could be used, but for this preliminary work we use this simple approximation.

Deep Q-Learning. SmartQueue uses deep Q-learning [17] in order to decide which query to execute next. As with any deep reinforcement learning system, SmartQueue is an agent that operates over a set of states S (buffer pool states) and a set of actions A per state (candidate queries to executed next). SmartQueue models the problem of query scheduling as a Markov Decision Process (MDP) [38]: by picking one query from the queue to execute, the agent transitions from the current to a new buffer pool state (i.e., data blocks cached). Executing a new query on the current buffer state, provide the agent with a reward. In our case, the reward of an action is the buffer hit ratio of the executed query calculated as $\frac{\text{buffer hits}}{\text{total block requests}}$.

The goal of the agent is to learn a *scheduling policy* that maximizes its total reward. This is a continuous learning process: as more queries arrive and the agent makes more scheduling decisions, it collects more information (i.e., context of the decision and its reward) and adapts its policy accordingly. The scheduling policy is expressed as a function $Q(S_t, A_t)$, that outputs a *Q-value* for taking an action A_t (i.e., a query to execute next) on a buffer state S_t . Given a state S_t and an action A_t , the Q-value $Q(S_t, A_t)$ is calculated by adding the maximum reward attainable from future buffer states to the reward for achieving its current buffer state, effectively influencing the current scheduling decision

by the potential future reward. This potential reward is a weighted sum of the expected buffer hit ratios of all future scheduling decisions starting from the current buffer state. Formally, after each action A_t on a state S_t the agent learns a new policy $Q^{new}(S_t, A_t)$ defined as:

$$Q(S_t, A_t) + \alpha[R_t + \gamma \max_a (Q(S_{t+1}, a) - Q(S_t, A_t))] \quad (2)$$

The parameter γ is the discount factor which weighs the contribution of short-term vs. long-term rewards. Adjusting the value of γ will diminish (e.g., favor choosing queries that will make use of the current buffer state) or increase (e.g., favor choosing queries that will allow long-term increased usage of the buffer) the contribution of future rewards. The parameter α is the learning rate or step size. This simply determines to what extent newly acquired information overrides old information: a low learning rate implies that new information should be treated skeptically, and may be appropriate when a workload is mostly stable but contains some outliers. A high learning rate implies that new information is more fully trusted, and may be appropriate when query workloads smoothly change over time. Since the above is a recursive equation, it starts with making arbitrary assumptions for all Q -values (and hence arbitrary initial scheduling decisions). However, as more experience is collected through the execution of incoming queries, the network likely converges to the optimal policy [27].

3. PRELIMINARY RESULTS

Here, we present preliminary experiments demonstrating that SmartQueue can generate query ordering that increase the buffer hit ratio and improve query execution times compared with alternative non-learned schedulers.

Experimental Setup. Our experimental study used workloads generated using the 99 query templates of the TPC-DS benchmark [29]. We deployed a database with a size of 49GB on single node server with 4 cores, 32GB of RAM. For our experiments, we generated 1,000 random query instances out of these 99 templates and placed them in a random order in the execution queue. The benchmark includes 165 tables and indexes, and the number of blocks for each of these ranged between 100 and 130,000. However, after downsizing both the query vector and buffer state bitmaps, our representation vectors have a size of $165 \times 1,000$, including index tables. We run our experiments on PostgreSQL [1] with a shared buffer pool size of 2GB.¹ For each query, we collect its query plan without executing the query by using the EXPLAIN command.

SmartQueue uses a fully-connected neural network. Our DRL agent was implemented with Keras [12] and uses 2 hidden layers with 128 neurons each. We also use an adaptive learning rate optimization algorithm (Adam [13]) and our loss function is the mean squared error.

In our study, we compare SmartQueue with two alternative scheduling approaches. *First-Come-First-Served (FCFS)* simply executes queries in the order they appear in the queue. *Greedy* employs a simple heuristic to identify the query with the best expected hit ratio given the current contents of the buffer pool. Specifically, for each queued query

it calculates the dot product of the buffer state bitmap with the data requests bitmap, estimating essentially the probability of buffer hits for each data block request. We then order all queries based on the sum of these probabilities over all blocks and execute the query with the highest sum value. Following the execution, the new buffer state is calculated and the heuristic is applied again until the queue is empty. This greedy approach focuses on short-terms buffer hits improvements.

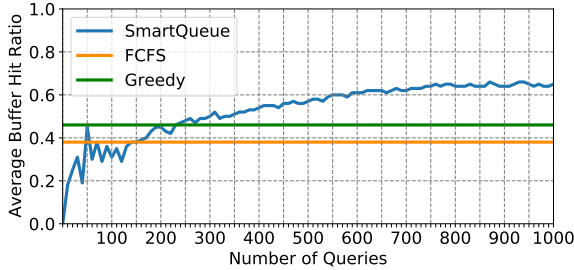
Effectiveness. First, we demonstrate that SmartQueue can improve its effectiveness as it collects more experience. In this set of experiments, we placed all 1,000 queries in the queue and we start scheduling them using SmartQueue. In the beginning our agent will make arbitrary scheduling decisions, but as it schedules more queries, SmartQueue collects more experience from its past actions and starts improving its policy. To demonstrate that, we evaluated the learned model at different stages of its training. Figure 2a and Figure 2b shows how the model performs as we increase the number of training queries. In Figure 2a, we measure the average buffer hit ratio when scheduling our 1,000 queries and we compare it with the buffer hit ratio of FCFS and Greedy (which is not affected by the number of training queries). We observe that the DRL agent is able to improve the buffer hit ratio as it schedules more queries. It outperforms the buffer hit of the other two heuristics eventually converging into a ration that is 65% higher than FCFS and 35% higher than Greedy.

In addition, Figure 2b shows the number of executed queries over time. The results demonstrate that DRL-guided scheduling of SmartQueue allows our approach to execute the workload of 1,000 queries around 42% faster than Greedy and 55% faster than FCFS. This indicates that SmartQueue can effectively capture the relationship between buffer pool state and data access patterns, and leverage that to better utilize the buffer pool and improve its query scheduling decisions.

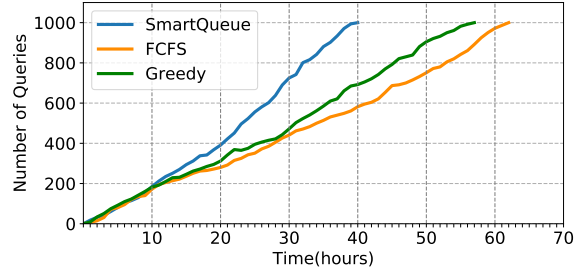
Adaptability to new queries. Next we studies SmartQueue’s ability to adapt to unseen queries. For these experiments, we trained SmartQueue by first scheduling 950 random queries out of 79 TPC-DS templates. We then test the model over 50 random queries out 20 unseen before TPC-DS templates. Figure 3a demonstrates how average buffer hit ratio of the testing queries is affected as SmartQueue collects experience increases from scheduling more training queries. The graph shows that the average buffer hit ratio of the testing queries is increased from 0.2 (when the SmartQueue is untrained) to 0.64 (when SmartQueue has schedule all 950 queries). Furthermore, SmartQueue outperforms FCFS and Greedy after having scheduled less than 500 queries.

Finally, Figure 3b, shows that the query latency of our testing queries keeps decreasing (and eventually outperforms FCFS and Greedy) as SmartQueue is trained on more queries. Our approach enables unseen queries to be eventually executed 11% faster than FCFS and 22% than Greedy. These results indicate that query scheduling policy can adapt to new query templates leading to significant performance and resource sharing improvements.

¹We configured PostgreSQL to bypass the OS filesystem cache. In future work, multiple levels of caching should be considered.

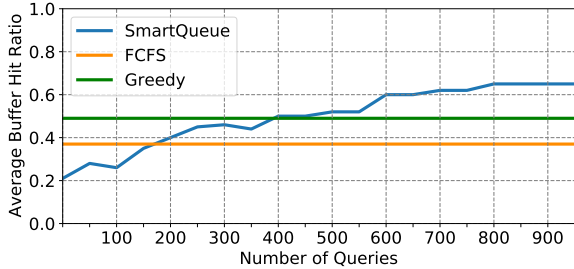


(a) Average buffer hit ratio

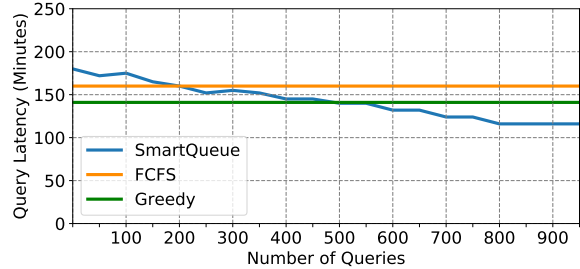


(b) Query execution time

Figure 2: SmartQueue’s effectiveness (buffer hit ratio and query completion rate) with increasing training sets.



(a) Average buffer hit ratio



(b) Query execution time

Figure 3: Buffer hit ratio and latency improvement on unseen query templates and increasing training queries.

Overhead. We also measured the training and inference time. Our proof-of-concept prototype needed 240 mins to incorporate 950 queries in our agent (so in average the training overhead is 3.95 mins per query). This time does not include the execution time of the query. This training overhead can potentially be optimized by offloading it into another thread, introducing early stopping, or re-using previous network weights to get a good “starting point.” There is no training overhead for FCFS and Greedy. The inference time of SmartQueue is 3.12seconds while the inference time for Greedy is 2.52 seconds and 0.0012 seconds for FCFS.

4. RELATED WORK

Prior work on query scheduling have focused on query parallelism [40], elastic cloud databases [4, 9, 18, 19, 22, 24, 30], meeting SLAs [6, 7, 15, 25, 33, 34, 43, 44], or cluster scheduling [20, 36, 39]. In terms of buffer pools and caching, most prior work has focused on smart cache management [2, 10] (i.e., assuming the query order is fixed and choose which blocks to evict or replace), or on (memory) cache-aware algorithms [46]. Here, we take a flipped approach, in which we assume the buffer management policy is fixed and the query order may be modified (e.g., batch processing). More broadly, work on learned indexes follows recent trends in integrating machine learning components into systems [11], especially database systems. Machine learning techniques have also been applied to query optimization [23, 32, 41], cardinality estimation [14, 31, 45], cost modeling [37], data integration [8, 28], tuning [42], and security [35].

5. CONCLUSION AND FUTURE WORK

We have presented SmartQueue, a deep reinforcement learning query scheduler that seeks to maximize buffer hit rates in database management systems. While simple, SmartQueue was able to provide substantial improvements over naive and simple heuristics, suggesting that cache-aware deep learning powered query schedulers are a promising research direction. SmartQueue is only an early prototype, and in the future we plan to conduct a full experimental study of SmartQueue. In general, we believe the following areas of future work are promising.

Neural network architecture. While effective in our initial experiments, a fully connected neural network is likely not the correct inductive bias [26] for this problem. A fully connected neural network is not likely to innately carry much useful information for query scheduling [21], nor is there much of an intuitive connection between a fully-connected architecture and the query scheduling problem [5]. The first layer of our network learns one linear combination per neuron of the entire input. These linear combinations would have to be extremely sparse to learn features like “the query reads this block, which is cached.” Other network architectures – like locally connected neural networks [16] – may provide significant benefit.

SLAs. Improving raw workload latency is helpful, but often applications have much more complex performance requirements (e.g., some queries are more important than others). Integrating query priorities and customizable Service Level Agreements (SLAs) into SmartQueue by modifying the reward signal could result in an buffer-aware and SLA-compliant scheduler.

Query optimization. Different query plans may perform differently with different buffer states. Integrating SmartQueue into the query optimizer – so that query plans can be selected to maximize buffer usage – may provide significant performance gains.

Buffer management. SmartQueue only considers query ordering, and assumes that the buffer management policy is opaque. A larger system could consider both query ordering and buffer management, choosing to evict or hold buffered blocks based on future queries. Such a system could represent an end-to-end query scheduling and buffer management policy.

6. REFERENCES

- [1] PostgreSQL database, <http://www.postgresql.org/>.
- [2] M. Altinel et al. Cache Tables: Paving the Way for an Adaptive Database Cache. In *VLDB '03*.
- [3] K. Arulkumaran et al. A Brief Survey of Deep Reinforcement Learning. *IEEE Signal Processing '17*.
- [4] Y. Azar et al. Cloud scheduling with setup cost. In *SPAA '13*.
- [5] P. W. Battaglia et al. Relational inductive biases, deep learning, and graph networks. *arXiv '18*.
- [6] Y. Chi et al. iCBS: Incremental Cost-based Scheduling Under Piecewise Linear SLAs. *VLDB '11*.
- [7] Y. Chi et al. SLA-tree: A Framework for Efficiently Supporting SLA-based Decisions in Cloud Computing. In *EDBT '11*.
- [8] R. C. Fernandez et al. Termite: A System for Tunneling Through Heterogeneous Data. In *aiDM '19*.
- [9] S. Genaud et al. Cost-wait trade-offs in client-side resource provisioning with elastic clouds. In *CLOUD '11*.
- [10] S. Ghandeharizadeh et al. Cache augmented database management systems. In *DBSocial '13*.
- [11] J. Gottschlich et al. The three pillars of machine programming. In *MAPL 2018*.
- [12] Keras: The Python Deep Learning library. <https://keras.io/>.
- [13] D. P. Kingma et al. Adam: A Method for Stochastic Optimization. In *ICLR '15*.
- [14] A. Kipf et al. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In *CIDR '19*.
- [15] W. Lang et al. Towards Multi-Tenant Performance SLOs. In *ICDE '14*.
- [16] Y. Lecun. Generalization and network design strategies. In *Connectionism '89*.
- [17] Y. LeCun et al. Deep learning. *Nature '15*.
- [18] P. Leitner et al. Cost-Efficient and Application SLA-Aware Client Side Request Scheduling in an Infrastructure-as-a-Service Cloud. In *CLOUD '12*.
- [19] Z. Liu et al. PMAx: Tenant Placement in Multitenant Databases for Profit Maximization. In *EDBT '13*.
- [20] H. Mao et al. Learning Scheduling Algorithms for Data Processing Clusters. *arXiv '18*.
- [21] G. Marcus. Innateness, AlphaZero, and Artificial Intelligence. *arXiv '18*.
- [22] R. Marcus et al. NashDB: An End-to-End Economic Method for Elastic Database Fragmentation, Replication, and Provisioning. In *SIGMOD '18*.
- [23] R. Marcus et al. Neo: A Learned Query Optimizer. *VLDB '19*.
- [24] R. Marcus et al. Releasing Cloud Databases from the Chains of Performance Prediction Models. In *CIDR '17*.
- [25] R. Marcus et al. WiSeDB: A Learning-based Workload Management Advisor for Cloud Databases. *VLDB '16*.
- [26] T. M. Mitchell. *The Need for Biases in Learning Generalizations*. PhD thesis.
- [27] V. Mnih et al. Human-level control through deep reinforcement learning. *Nature '15*.
- [28] S. Mudgal et al. Deep Learning for Entity Matching: A Design Space Exploration. In *SIGMOD '18*.
- [29] R. O. Nambiar et al. The Making of TPC-DS. In *VLDB '06*.
- [30] V. Narasayya et al. SQLVM: Performance Isolation in Multi-Tenant Relational Database-as-a-Service. In *CIDR '13*.
- [31] P. Negi et al. Cost-Guided Cardinality Estimation: Focus Where it Matters. In *SMDB @ ICDE '20*.
- [32] J. Ortiz et al. Learning State Representations for Query Optimization with Deep Reinforcement Learning. In *DEEM '18*.
- [33] J. Ortiz et al. PerfEnforce Demonstration: Data Analytics with Performance Guarantees. In *SIGMOD '16*.
- [34] J. Ortiz et al. SLAOrchestrator: Reducing the Cost of Performance SLAs for Cloud Data Analytics. In *USENIX '18*.
- [35] Shrainik Jain et al. Database-Agnostic Workload Management. In *CIDR '19*.
- [36] B. Sotomayor et al. Virtual Infrastructure Management in Private and Hybrid Clouds. *IEEE IC '09*.
- [37] J. Sun et al. An end-to-end learning-based cost estimator. *VLDB '19*.
- [38] R. S. Sutton et al. *Introduction to Reinforcement Learning*. 1998.
- [39] R. Taft et al. STeP: Scalable Tenant Placement for Managing Database-as-a-Service Deployments. In *SoCC '16*.
- [40] S. Tozer et al. Q-Cop: Avoiding bad query mixes to minimize client timeouts under heavy loads. In *ICDE '10*.
- [41] I. Trummer et al. SkinnerDB: Regret-bounded Query Evaluation via Reinforcement Learning. *VLDB '18*.
- [42] D. Van Aken et al. Automatic Database Management System Tuning Through Large-scale Machine Learning. In *SIGMOD '17*.
- [43] P. Xiong et al. ActiveSLA: A Profit-oriented Admission Control Framework for Database-as-a-Service Providers. In *SoCC '11*.
- [44] P. Xiong et al. Intelligent management of virtualized resources for database systems in cloud environment. In *ICDE '11*.
- [45] Z. Yang et al. Deep unsupervised cardinality estimation. *VLDB '19*.
- [46] J. Zhou et al. Buffering database operations for enhanced instruction cache performance. In *SIGMOD '04*.