Digital Twins for Storage Systems and RAID Pools: Enhancing Data Management in High Energy Physics

Abdalaziz R. Al-Maeeni, Aziz Temirkhanov, Artem Ryzhikov, Dmitry Maevskiy, Yuri Pakhomov, Fedor Ratnikov, and Mikhail Hushchyn

HSE University, Moscow, Russia

LAMBDA • HSE

NATIONAL RESEARCH UNIVERSITY

Motivation and Objectives

High energy physics (HEP) experiments generate vast amounts of data, requiring efficient storage and management. **Digital twin** concept involves creating digital replicas of physical systems, improving **efficiency** of data storage in HEP experiments. Digital twins of storage systems can help **monitor**, **analyze**, and **optimize** various workloads.

- Develop a digital twin of a data storage system (DSS).
- Predict performance of the DSS for a given configuration of the system and data load parameters.

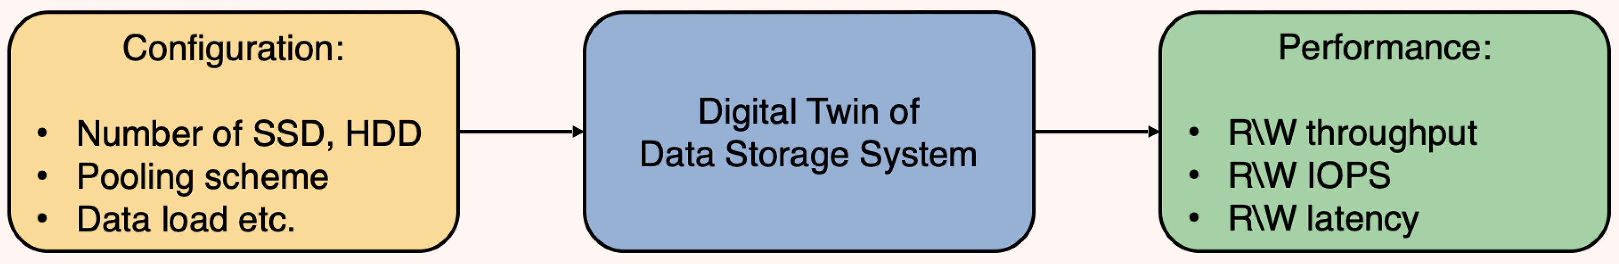

Introduction

In this study, we consider simulation of key components of data storage systems, namely, Solid-State Drives (SSD) storage pools. The goal of this work is to predict the performance of these components for a given configuration and data load parameters, describing the performance by the number of input and output operations per second (IOPS) and their average latencies. The data load parameters are described in table 1. Two load types are considered in this work, random and sequential, each consisting of a mixture of read and write operations. The SSD storage pools have configuration parameters, including the total number of disks and the RAID scheme, described by the number of data and parity blocks.

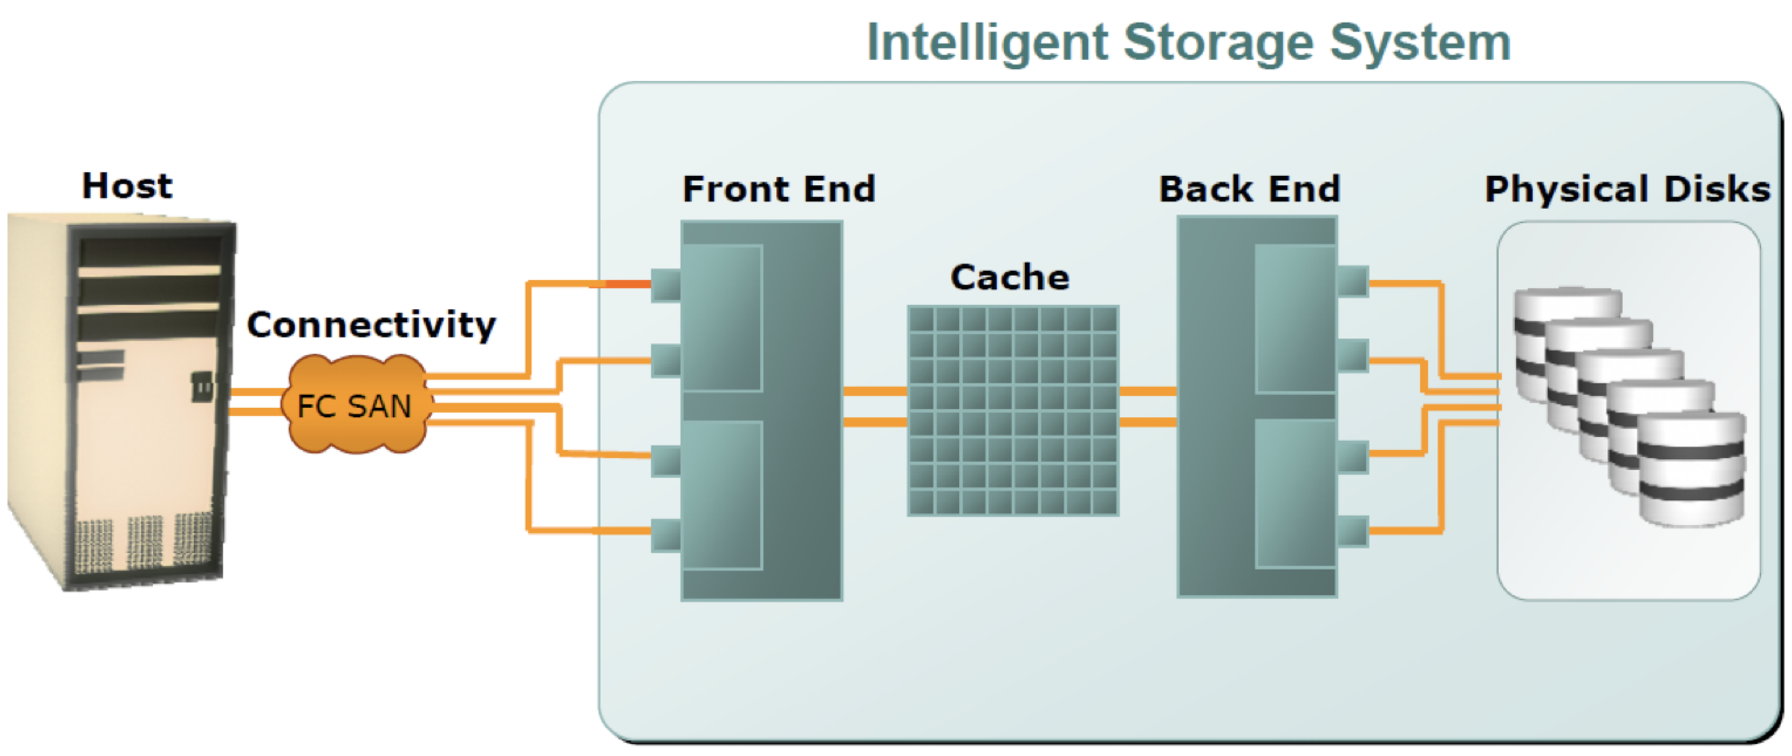

Figure 1. Development of a digital twin of a data storage system (DSS) to predict its performance for a given configuration and data load parameters.

Methods

CatBoost regression model is used as a parametric generative model. Relations between IOPS and latencies within the same data loads are defined by Little's law . All measurements of IOPS and latencies are stochastic. We approximate distribution of their logarithm values by conditional 2D normal distributions $\hat{z}_i = \log \hat{y}_i,\ \hat{z}_i \sim \mathcal{N}(\hat{\mu}(x_i), \hat{\Sigma}(x_i))$ where $\hat{y}_i$ is a vector of predictions for IOPS and latency; the mean $\hat{\mu}(x_i)$ and the covariance matrix $\hat{\Sigma}(x_i)$ depend on input vector x_i of data load and configuration parameters, and are predicted by the CatBoost regression model. We calculate the mean vectors μ_j and the covariance matrices Σ_j for each of these data loads. In addition, we use Cholesky decomposition for the matrices $\Sigma_j^{-1} = L_j L_j^T$. to ensure the positive semi-definiteness of predicted covariance matrices. We fit the CatBoost regression model with the MutliRMSE loss function defined as

$$L^2 = \frac{1}{2m}\sum_{j=1}^{2m}(\hat{\mu}(x_j) - \mu_j)^2 + \frac{1}{2m}\sum_{j=1}^{2m}(\hat{L}(x_j) - L_j)^2$$

, where $\hat{\Sigma}^{-1}(x_i) = \hat{L}(x_j)\hat{L}(x_j)^T$. The model consists of 5000 decision trees. The optimal hyperparameters values are estimated using grid search for each sample in our study.

Quality metric

Mean/Median Absolute Percentage Error (MEAPE):

$$\text{MEAPE} = \left|\frac{\hat{\mu} - \mu}{\mu}\right| \times 100\%$$

$$\mu = \begin{cases} \frac{1}{k}\sum_{i=1}^{k} y_i \\ \text{median}(y) \end{cases}$$

Conclusions

The proposed machine learning-based digital twins for storage systems can be effective in:

- Storage design and performance optimization
- Inefficiencies identification and cost reduction
- Ensuring reliability and scalability by facilitating informed decisions

Data

The datasets has been collected for the SSD pool under random and sequential data loads following the parameteres below:

Parameter	Random	Sequential
Block size	4, 8, 16, 32, 64 KB	128, 256, 512, 1024 KB
Read fraction	0 - 100%	0%
Number of jobs	1 - 32	1 - 20
Queue depth	1 - 32	1 - 32
RAID (K+M)	1+1, 2+1, 2+2, 4+1, 4+2, 8+2	
Number of disks	K+2M, 24, +3 values in between	

Table 1. Data load parameters and their value ranges for the storage SSD pool data set. For sequential and random data loads. We generate 512 different data loads using Sobol sequence, each load run for 120 seconds each, during which we measured IOPS and average latency for read and write operations.

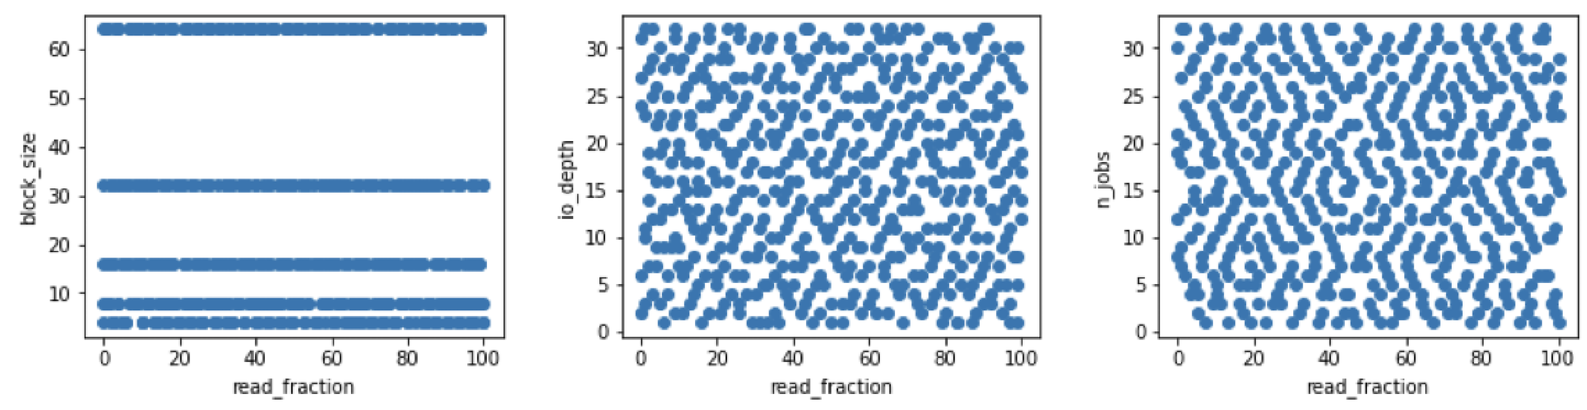

Figure 2. Efficient data collection using Sobol sampling of the data load parameters for SSD disks. Note the even coverage of the parameters space

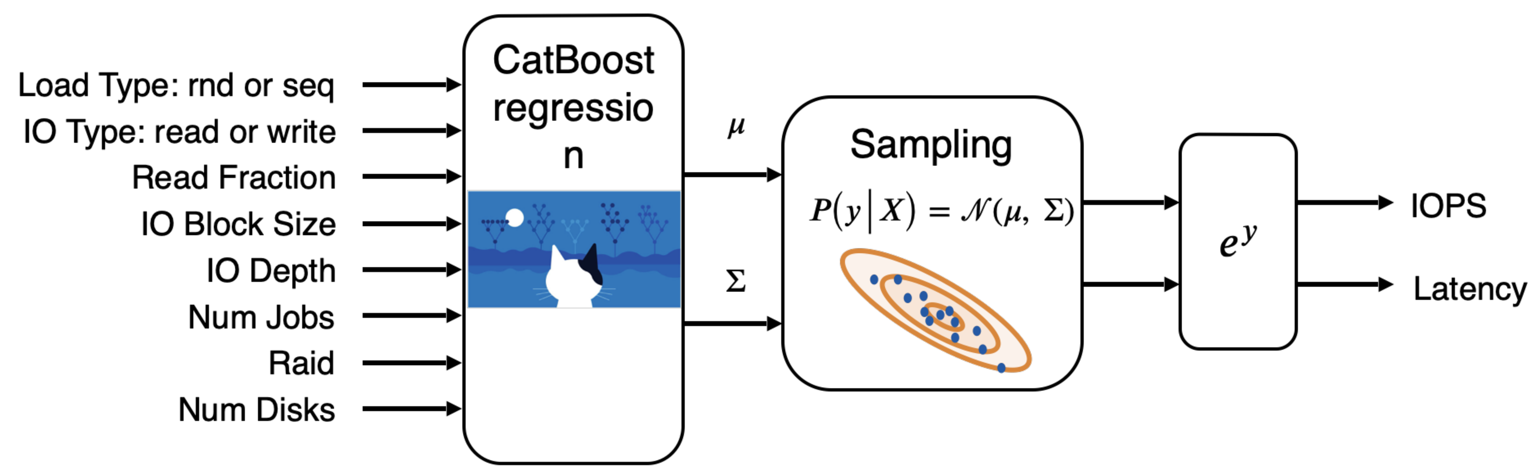

Figure 3. CatBoost model for performance predictions of the storage pools and cache for the given values of data load and configuration parameters.

Results

		Mean MEAPE %	Median MEAP %
SSD Pool, seq	IOPS	5.6	5.3
	Latency	12.2	6.9
SSD Pool, rand	IOPS	6.7	4.3
	Latency	8.3	4.1

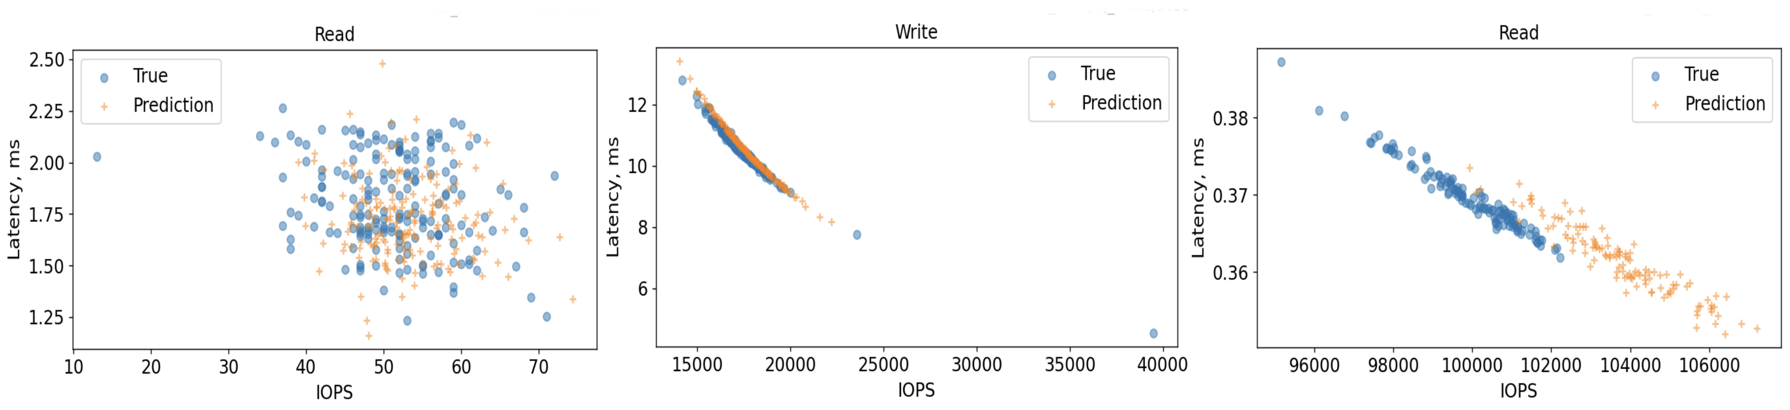

(a) (b) (c)

Scenario	Block size	RAID	num disks	num jobs	IOdepth	Read fraction	IO type	load type
(a)	512	2+2	15	6	1	1%	read	sequential
(b)	256	1+1	8	15	12	0%	write	sequential
(c)	16	4+1	20	2	25	77%	read	random

Figure 4. Samples of simulated IOPS and latencies for for different scenarios.