

# Performance Models of Flash-based Solid-State Drives for Real Workloads

---

Simona Boboila  
Peter Desnoyers

Northeastern University





# Issue

---

- Performance of flash devices on **real workloads** is hard to characterize and predict:
  - workload-dependent
    - real workloads are not well-behaved
  - depends on the Flash Translation Layer (FTL) implementation
    - this implementation is often not known



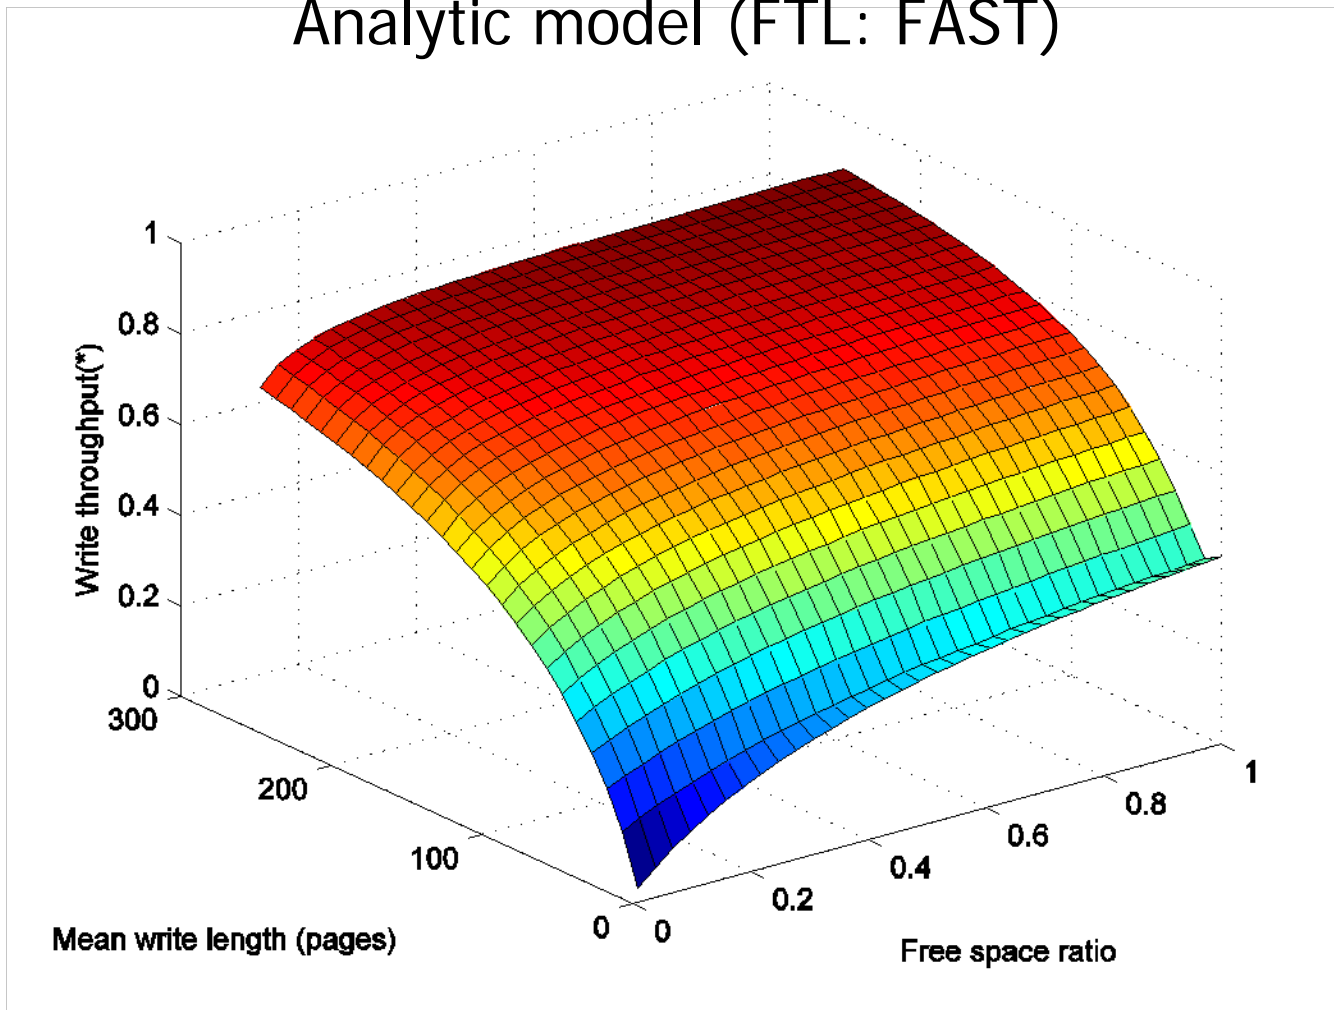
# Comparison with previous work

---

- Benchmarking (e.g. uFLIP) :
  - Performance measures for **specific IO patterns** from synthetic workloads
- Trace-driven simulations:
  - **Run each trace entirely** and get exact performance values per trace
- Our work:
  - Performance estimations for **real workloads** with real IO patterns
  - Black-box model (2 phases):
    - Training:
      - One-time preparation phase per SSD, consisting in running multiple synthetic workloads
    - Application:
      - Given a real workload, compute two statistics and use them to estimate performance, **without running the trace** on the SSD;
  - Analytic model for specific FTLs (probabilistic model for FAST, BAST)


# Model

Analytic model (FTL: FAST)



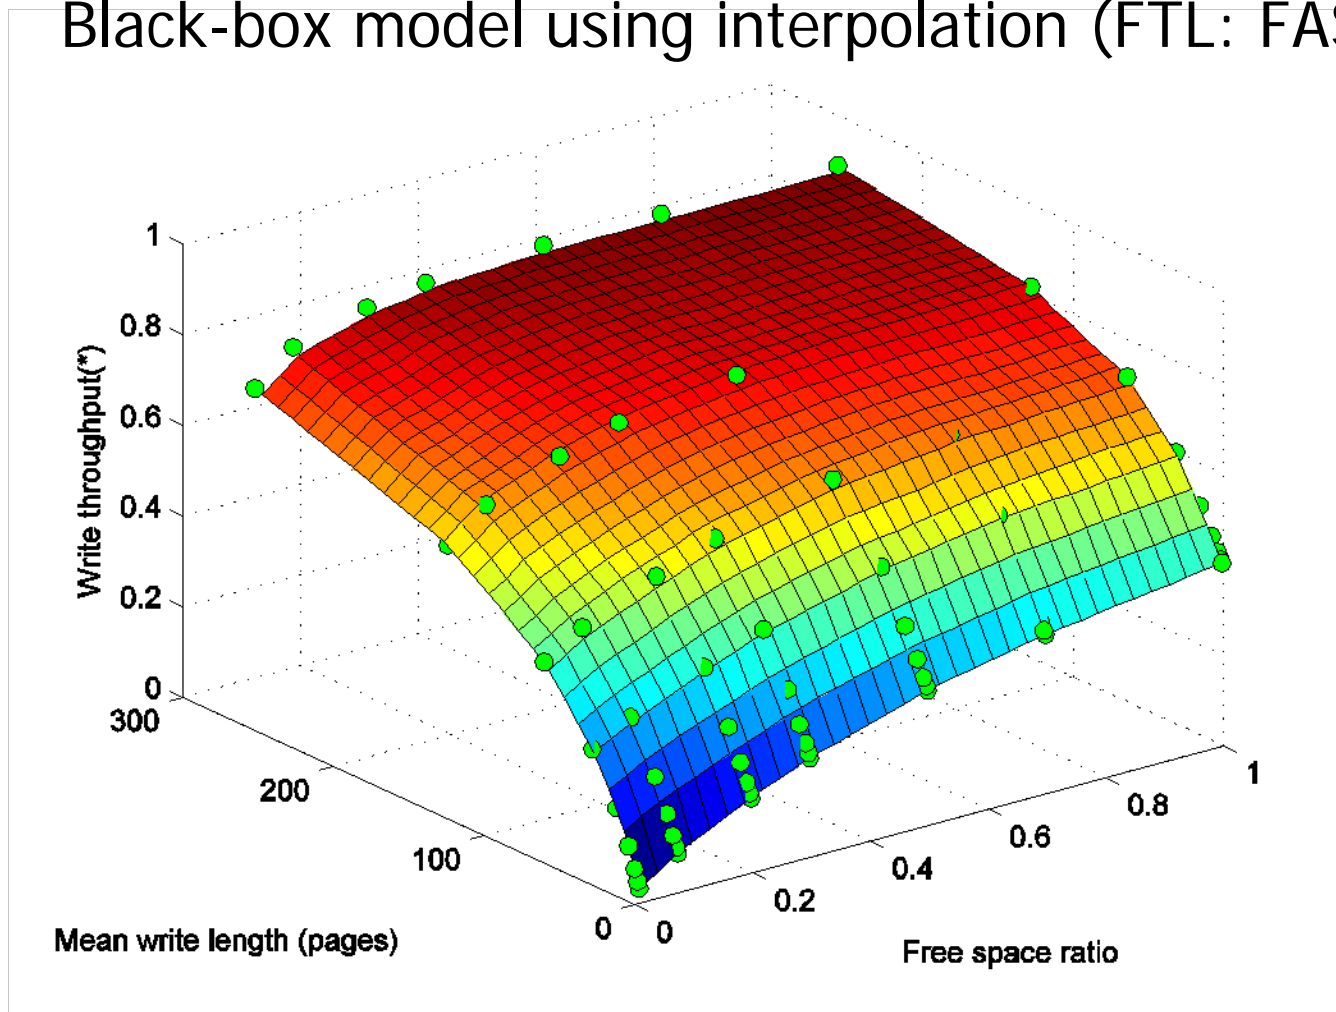
# Model

Black-box model using interpolation (FTL: FAST)




# Model

Black-box model using interpolation (FTL: FAST)



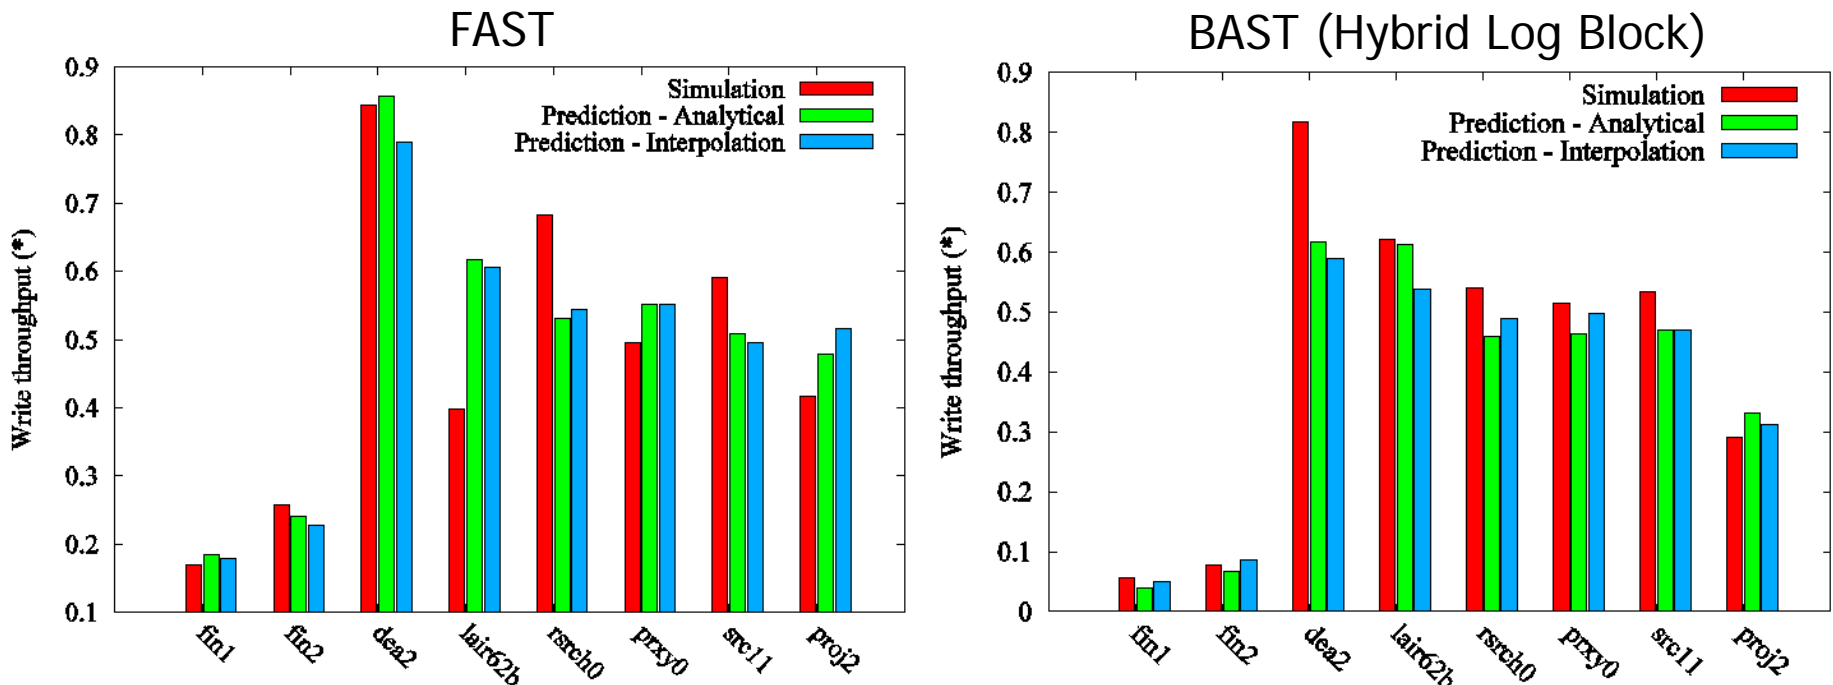
# Method

- For each workload segment:
  - **assume well-behaved traffic**: uniformly distributed requests; exponentially distributed write lengths
  - **compute 2 statistics** which give the parameters of the model (traffic sequentiality, free space utilization across workload segments)
  - **estimate performance**: using an interpolation-based **black-box model** constructed from multiple simulated (FlashSim) or measured (SSDs) data points, or **analytically** (for FAST, BAST)



# Simulation results

Analytical and interpolation-based performance prediction for real workloads




(\*) "Write throughput" is erasure cost – based: Writes (in blocks) / erasures.



# SSD results

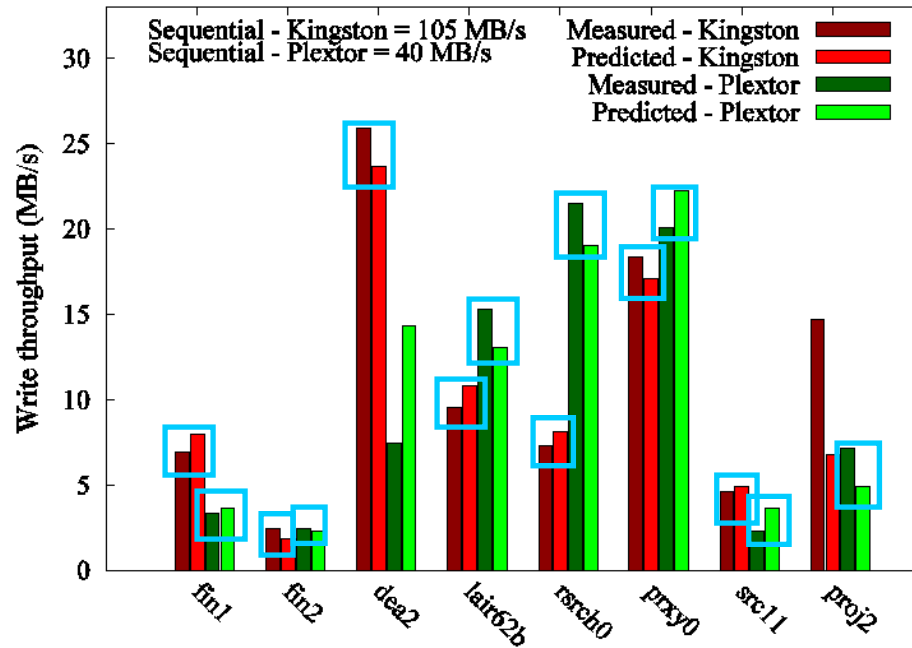
## SSD performance predictions on real workloads



In all but two cases (dea2 for Plextor, proj2 for Kingston) good correspondence between predicted and measured throughput was observed.

# SSD results

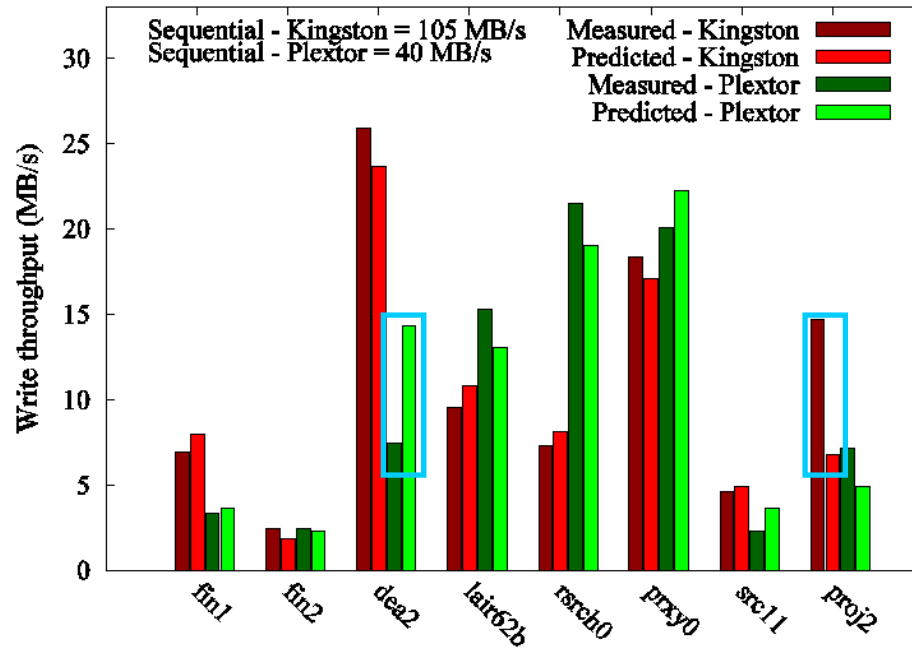
## SSD performance predictions on real workloads



In all but two cases (dea2 for Plextor, proj2 for Kingston) good correspondence between predicted and measured throughput was observed.

# SSD results

## SSD performance predictions on real workloads



In all but two cases (dea2 for Plextor, proj2 for Kingston) good correspondence between predicted and measured throughput was observed.