


V. Tarasov<sup>★</sup>, S. Kumar<sup>★</sup>, J. Ma<sup>†</sup>, D. Hildenbrand<sup>×</sup>, A. Povzner<sup>×</sup>, G. Kuenning<sup>†</sup>, E. Zadok<sup>★</sup>

**Production I/O traces** are valuable as they embody real-world workloads. Unfortunately, traces are cumbersome to replay. **Benchmarks**, on the other hand, are easy to setup and use, but lack the realism of production workloads. In this project we bridge the gap between benchmarks and traces by developing a methodology and a set of tools for **automatic conversion** of large I/O traces into concise workload models that can later be used as benchmarks.


## 1 Traces



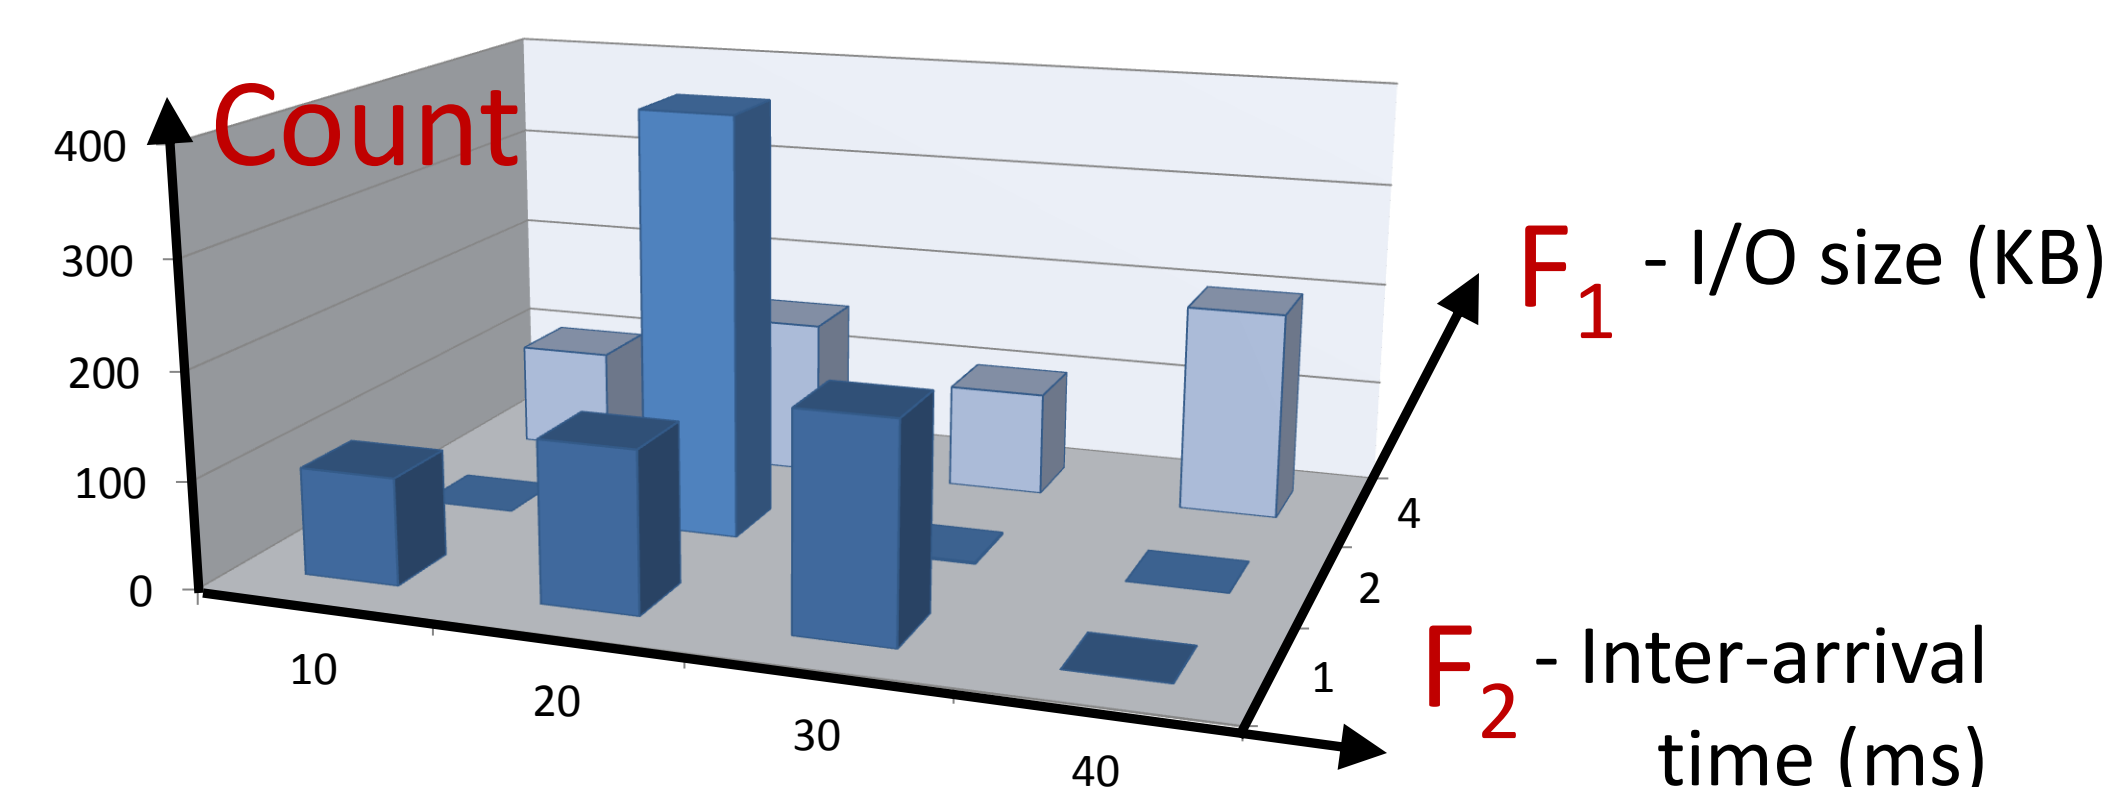
- Represent real-world workloads
- **Credible**: often used in the research  
35 out of 68 papers in FAST, ATC, OSDI, HotStorage, and MSST in 2010
- **Quite a few are collected**  
SNIA's trace repository growing, dozens currently available
- **Massive in size**
- **Cumbersome to replay**
- **Hard to scale**
- **Difficult to adjust even to small workload changes**




## 4 Modeling a chunk



- The number of **workload parameters** is **infinite**
- Select workload parameters:
  - Parameters **used** in the studies **before**
    - ✓ Read/Write ratio, Inter-arrival time, I/O size, ...
  - Make the tool **extensible**
- Use **multi-dimensional** histogram of **feature functions**



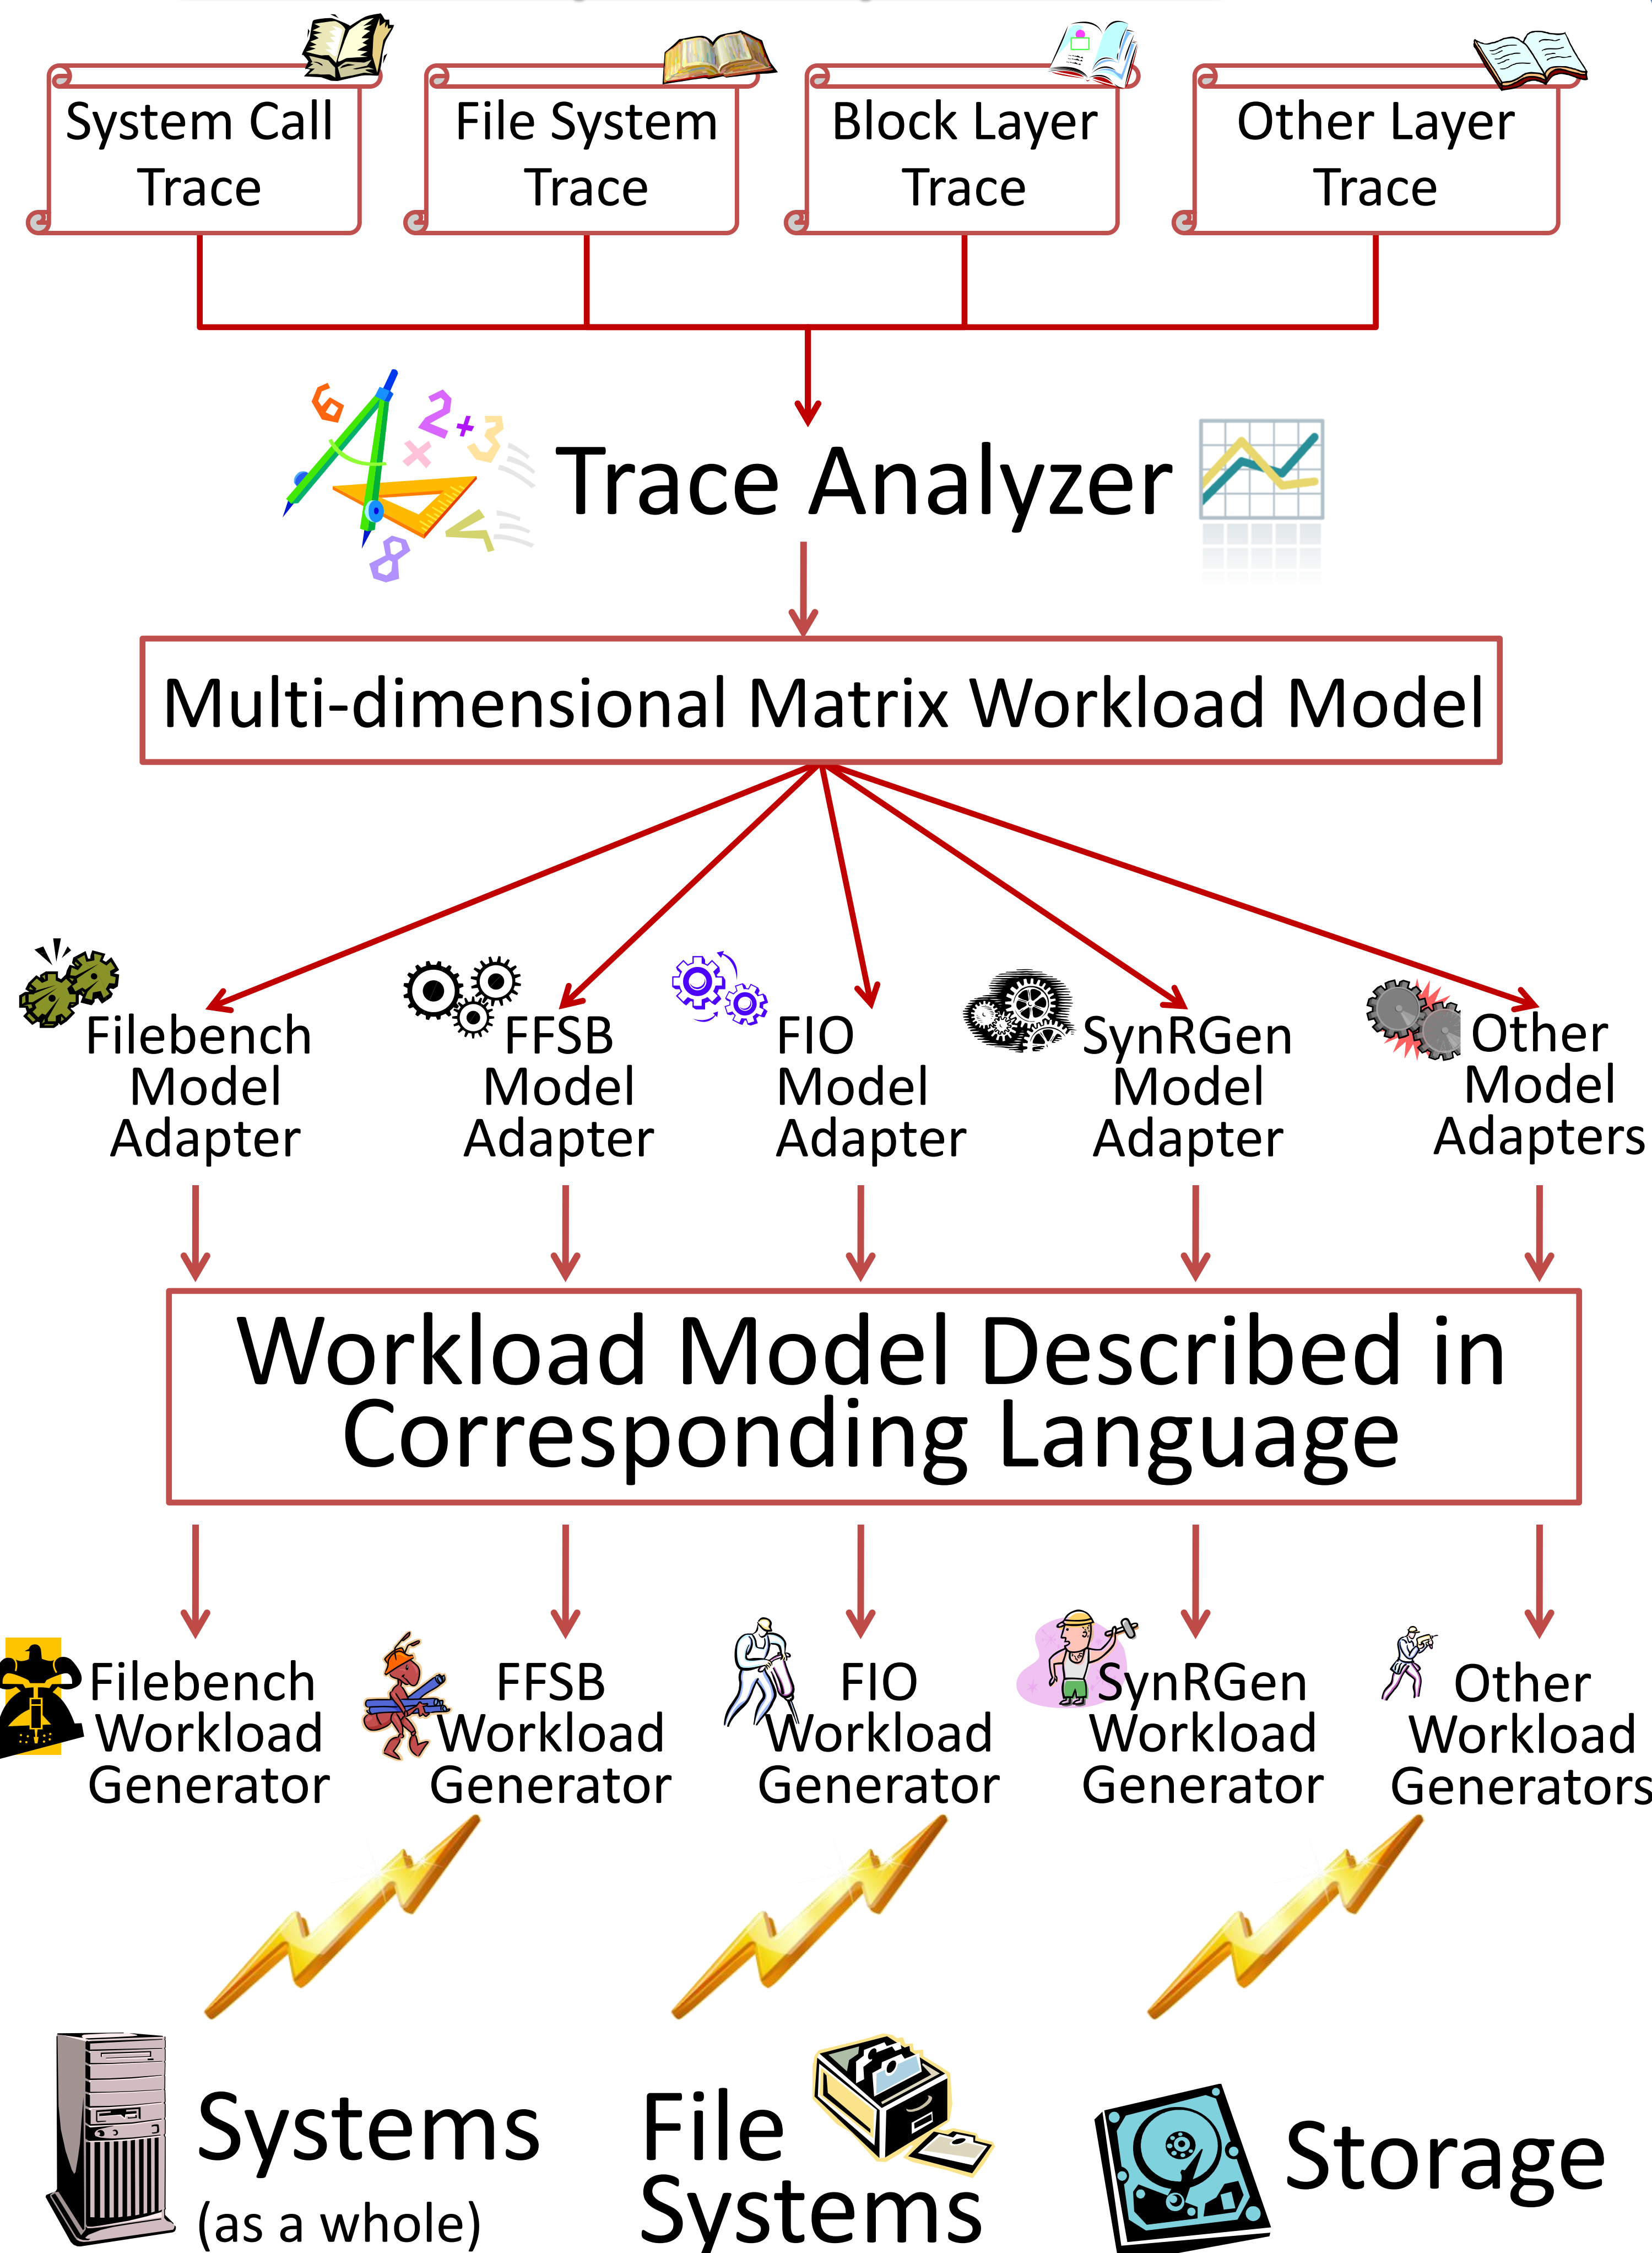
## 2 Benchmarks



- Great for micro-benchmarking
- Easy to set up and run
- Usually scale well
- Workloads are easily modifiable
- **Workloads generated might be not relevant to the real world**



## 5 Complete picture




Can we **combine** the best of two worlds?

## 3 Trace chunking



Average behavior over **entire** trace loses time variance → **Divide** trace into self-consistent chunks



- **Fixed** chunking as a first step
- Perform chunk **deduplication** to save trace space:
  - Identify chunks, where workload is similar
  - Store only **single instance** of similar chunks
  - Maintain a **mapping** between trace time and unique chunk instances