


# File Attribute-Based Predictions and Optimizations

Daniel Ellard<sup>1</sup>, Eno Thereska<sup>2</sup>, Michael Mesnier<sup>2</sup>,  
Gregory R. Ganger<sup>2</sup>, Margo Seltzer<sup>1</sup>


<sup>1</sup>Harvard University

<sup>2</sup>Carnegie Mellon University




## Background

- There is a strong association between the create-time attributes of most files and the operations performed on those files.
- Given the name, permissions, owner, and group we can predict lifespan, size, access patterns, etc.
- Example predictions:
  - “.lock” files are usually zero-length and short-lived.
  - “.tgz” files are usually large and live a long time.
  - Our system finds more complex associations.




## Current Status

- Our system automatically finds associations.
- Our system automatically builds models that make accurate predictions about future files.
  - The models are small, simple and efficient.
  - It is practical to use these models in the kernel to make low-level decisions.




## Using the Predictions

- File layout policies
  - Try to get it right the first time instead of relying on per-file adaptation.
- Caching policies
  - Are we likely to read this file again soon?
- Replication policies
  - Replicate read-mostly data for quick access




## Related Ideas

- Tune global system parameters based on the expected workload:
  - Example: if the workload is dominated by specific types of files, globally optimize for those file types instead of making file-by-file decisions.
- Identify common workloads idioms:
  - Example: protecting file access with a lock. Once a process tries to acquire the lock, we can often anticipate what will happen next.



## Conclusions

- Attribute-based predictions are accurate.
- We can use these predictions at run time to modify system behavior, such as file layout.
- It is a bit awkward to integrate these predictions into current file systems.
- How much do these predictions help? When we benchmarks for adaptive systems running complex workloads, I'll tell you.



## The Obligatory URL Slide

`www.eecs.harvard.edu/sos`

`www.pdl.cmu.edu`