

Caching less for better performance: Balancing cache size and update cost of flash memory cache in hybrid storage systems



Yongseok Oh, Jongmoo Choi, Donghee Lee, Sam H. Noh
 University of Seoul, Dankook University, Hongik University (in Seoul, Korea)

Hybrid Storage System	Over-Provisioned Space (OPS)	Our Goal: Find Optimal OPS Size
<ul style="list-style-type: none"> Combine SSDs and HDDs SSD-like performance for HDD-like price SSDs used as Flash Cache Issues in Flash Cache <ul style="list-style-type: none"> Performance <ul style="list-style-type: none"> Garbage collection (GC) Lifetime <ul style="list-style-type: none"> Erase count 	<ul style="list-style-type: none"> Reserved space for GC in Flash Cache Greatly influence GC cost and hit rate In typical SSDs OPS size is FIXED to an undisclosed size <ul style="list-style-type: none"> Cannot adapt to workload changes & GC cost 	
<p>Conventional Hybrid Storage</p>		

OP-FCL (Optimal Partitioning-Flash Cache Layer): Workload Dependent Optimal Partitioning

- Flash Cache is divided based on u and r

- Hybrid Cost Model: $C_{HY}(u, r)$** represents expected I/O cost
 - Involves
 - Storage Cost Model
 - Flash access cost
 - HDD access cost
 - Workload Pattern
 - Hit rate
 - I/O rate
- See the paper for derivation

- Periodically Execute Optimal Partitioning Algorithm

```

procedure OPTIMAL_PARTITIONING
  step ← segment_size/total_cache_size
  INIT_PARAMS(op_cost, op_u, op_r)
  for u ← step; u < 1.0; u ← u + step do
    for r ← 0.0; r ≤ 1.0; r ← r + step do
      cur_cost ← C_HY(u, r)
      if cur_cost < op_cost then
        op_cost ← cur_cost
        op_u ← u, op_r ← r
      end if
    end for
  end for
  ADJUST_CACHE_SIZE(op_u, op_r)
end procedure
    
```

- Find u and r resulting in Optimal I/O Cost

- Adjust Flash Cache partition

Performance Evaluation

Hybrid Storage Simulator

- CMU DiskSim 4.0+MSR SSD extension
 - 16GB Flash Cache+10K RPM HDDs

Flash Cache Layers

- FP-FCL (Fixed Partitioning)
 - Conventional Hybrid Storage
- RW-FCL (Read Write Partitioning)
 - Fixed OPS size with Read/Write Partitioning
- OP-FCL (Optimal Partitioning) that we propose**

Workload

- Exchange Server
- See the paper for more results!

Conclusion

- Trade-off exists
 - Caching Benefit vs. Update Cost
- OP-FCL balancing caching space and OPS sizes
 - Provides near optimal performance
 - Improves lifetime of Flash Cache

