

Edsger W. Dijkstra

The Structure of the "THE" Multiprogramming System


Communications of the ACM 11(5), May 1968.

①

- 1) a model for multiprogramming OS ~~OS~~ ✓✓✓
 - 2) to prove delaying process does not hurt, ✓✓
 - 3) provable system. ✓✓✓✓
- ~~to provide~~ ✓

- 2) (1) Hierarchical abstractions ~~abstractions~~ ✓
- 2) Semaphores as synchronizing primitives & for mutual exclusion ✓

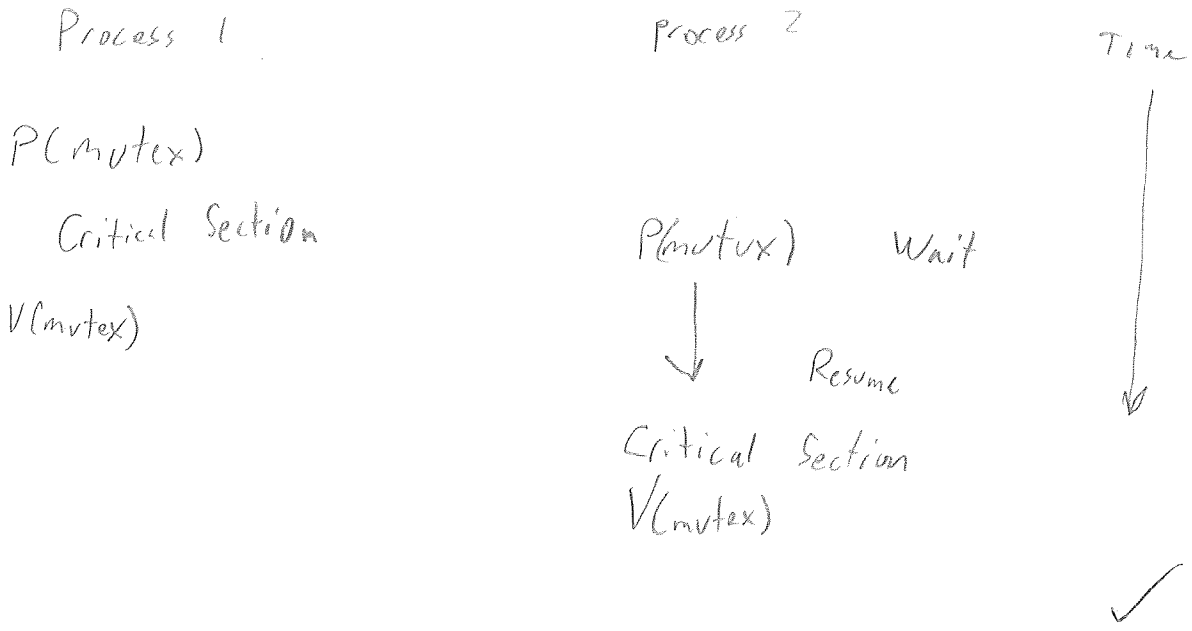
③ Hierarchy

- Level 0 - abstract the processor with a process scheduler
 - Level 1 - segment controller, abstracts secondary storage into "segments"
 - Level 2 - message interpreter, support conversation b/t user and their high-level processes
 - Level 3 - buffering I/O
 - Level 4 - user programs
 - Level 5 - operator (not implemented by THE )
- ✓✓✓

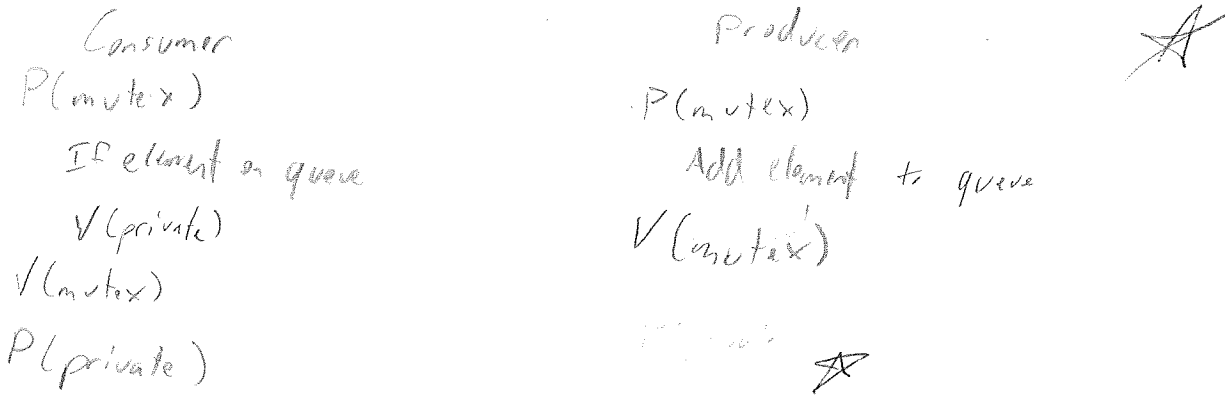
4.) Virtual address space - separates data from where it is stored. Can map virtual address to physical address. Advantage is that [virtual address space can be much bigger than physical address space]. Also don't need contiguous memory, meaning be more efficient memory usage. ✓✓

Disadvantage is takes extra space (page table) and lookups take extra time. ✓

5 Semaphore for mutual exclusion



Semaphore for Scheduling



1. Motivation:
✓ ✓ ✓
Create an extensible OS, which lets the applications decide their own policies and even behaves as OS themselves. They wanted to do this because current OS don't allow enough flexibility to applications.

2. High level point

✓ ✓
The nucleus is used to implement the simulation of processes, comm. b/w them, creation, deletion & control of the processes.

3. Inter-process communication:

✓
Get away with semaphore driven Synchronization for better efficiency & security. Use message passing through buffers. This technique is an asynchronous way of communication & scales well. Basic message calls:
send, wait, } message & send, wait } response
Awesome Writing!

4. Process Hierarchy

✓
The memory space originally belongs to the initiating process, which can allocate a subset of its memory to a process it creates. This results in a parent-child protection scheme for memory access. As an extensible / generic process management scheme, this also allows arbitrary OS's to exist/run without interference, and needing no specialized management.

5. 2 ^{levels of} processes: internal processes, external processes (I/O devices)

- treat internal and external processes uniformly to create extensible system. The difference between them is merely a matter of processing capability.

Ext

Creation: assignment of name to a peripheral device.

Real time system

Review :

1. Introduced process names.
2. ~~Pro~~ Document is also a process (external)

Internal

assignment of name to a contiguous storage area selected by the parent

Ritchie, D.M. and Thompson, K.
The UNIX Time-Sharing System
Communications of the ACM, Vol. 17, No. 7, July 1974, pp. 365-375.

1) * Wanted to make a general multi-user operating system that was simple, easy to use. (Authors thought filesystem was the most important role of an OS) ^{??} so they focused mostly on that. Also wanted an easy to use shell and command execution.

2) Conceptual contributions - I/O equivalent to file system access
- fork/exec/pipe/wait/exit - style process management

Practical contributions - the Unix system, including the file system & shell
- particularly that a multiprocessing system can work on relatively inexpensive hardware

3) → File system with directory tree, inodes links
→ special files for I/O devices - /dev
→ linking of files → protection bits → r/w/x
→ setuid.

adv
→ for general purpose use.
not for any predefined objectives

Disadv
- Separate inode from the data/file → lot of disk seeks.
- small block size → wastes the disk b/w for large files

4) Process - execution of an image (code, registers, open files, current directory, etc.)

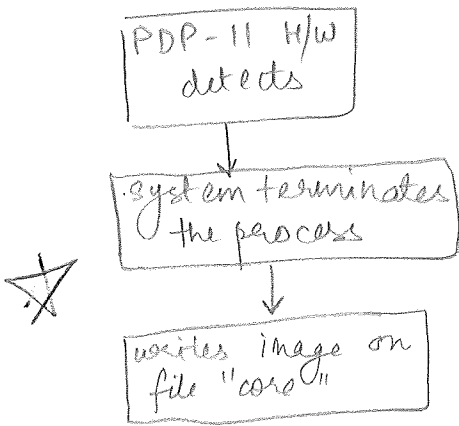
- Exists in memory
- Created by fork() in parent process: copy of memory, shared files, etc.
- IPC with pipes
- Execute new program with exec(ute)()

Adv.
fork() + exec() clean + easy

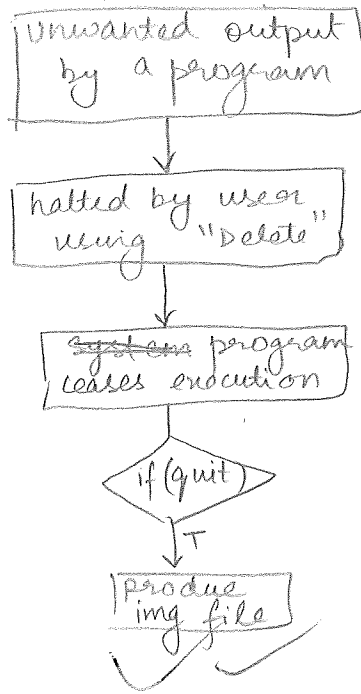
P.T.O.

5) TRAPS → How they work?

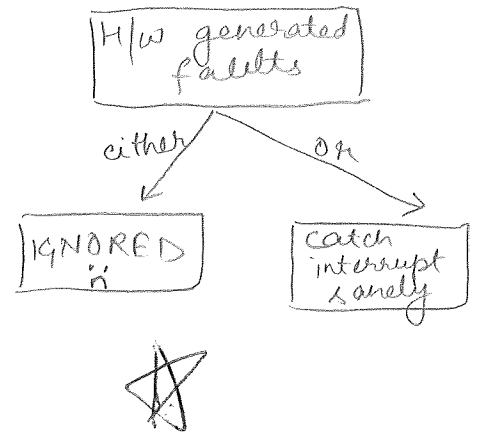
Scenario 1



Scenario 2




Scenario 3



- 1) The motivation of this paper is that researchers use disk models to improve overall I/O performance in disks by running simulations on these models. However the problem is that little work has been done to create accurate disk drive models. This severely limits the useful information gathered, because simulations are run on models that do not reflect reality. Therefore, the author aims to build an extremely accurate disk model.
2. The authors discuss the physical characteristics of disk drives and explain how they affect r/w behavior. They present an accurate model of the disk drive. They created a simulator based on their this understanding which can accurately predict disk emulate disk I/O behavior.
3. The primary technique was to model specific physical elements of the disk's operation. The simulator tried a model with linear seek times, which was not that accurate, and then a model with seek times that accurately reflected the acceleration, max speed movement, deceleration, and head settling time of a seek, resulting in much more accurate modelling.
4. Another major component they modeled was the caching behavior of typical disks. They were able to show that a model that does not provide a cache simulation is widely inaccurate since a large percentage of I/O is handled at the cache. After implementation their model was much more in line with a real disk.

5) Track Skewing: offset block data on each track so that ~~least~~ data can be read continuously with having to seek back on the disk. When head or track switches occur.



McKusick, M.K., Joy, W.N., Leffler, S.J., and Fabry, R.S.

A Fast File System for UNIX

ACM Transactions on Computer Systems, Vol. 2, No. 3, August 1984, pp. 181-197.

1) the problem was that the original UNIX FS was incredibly slow, getting a blazing 2% of max disk transfer. ✓

there were 3 problems: ✓

1. small block sizes - more transfers, w/ fixed cost per transfer
2. poor freelist organization - consecutive blocks not close together?
3. No locality - inodes far from data blocks

2) Contribution.

- Better locality is proposed. - using cylinder group. ✓
- larger block size.
- fragments for small files ✓


3) - Organize data which is accessed together in close spatial locality. Like inodes and data blocks.

- Use 4KB blocks to get better bandwidth, but also support smaller fragments to ~~save~~ reduce wastage. ✓
- Place file data in rotationally optimal blocks to increase sequential r/w performance. ✓

4) - Layout policy: Files in a folder are put in the same cylinder group if they fit and are smaller than 48kb. Advantage: Less seek time when accessing files in same directory. Seek time is less of overall cost for large files. More files in folder fit together.

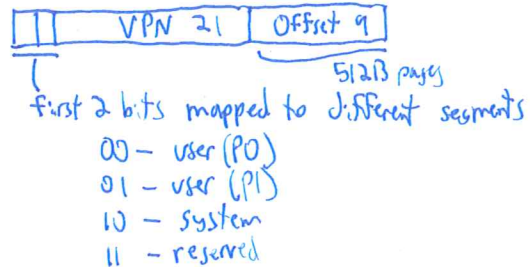
Disadvantage: Large files are split across disk at 1 MB intervals, which increases read time. ✓

5) if the CPU does not have I/O channel, the data ~~has~~ has to be put in a "rotationally optimal" location, so the disk head will be on the right position after the CPU interrupt.



- 1.) Goal was to support a variety of physical mem sizes and demands of timeshared systems, realtime and batch processes. ✓✓ (250k + 8M)
- 2.) Predictable performance can be given by ~~maximizing~~ giving more space in memory (using techniques like swapping and paging out page tables) and reducing I/O (using clustering). ✓? ? ?
- 3) Their implementation of the swapper was a solution to the problem of the high ^{cost} during process cost associated with paging in pages after a process is able to execute again. Without ^{swapper} the swapper keeping a process resident set to load in pages on a process waking each page would have been faulted in. ✓★✓✗
- 4) VAX also implemented a pager service designed to efficiently manage paging in and out. In particular, when a page is removed from a process' resident set, the pager puts it on a free list (if it is unmodified) or a modified list, which serve dual duty as caches for pages that might be read in again later and as sources of free pages for when new pages are read into memory. They could also be used to batch I/O of writing modified pages back out to disk. ★★✗


5) VAX's virtual addresses were as follows:



To ~~convert~~ translate an address into a physical memory location, VAX would:

- ① Look at the first 2 bits, and go to the corresponding segment's base register.
- ② Jump to that segment's page table, and translate VPN into PRN.
- ③ Jump to that physical page, and apply the offset






✓

R. Rashid and A. Tevanian and M. Young and D. Golub and R. Baron and D. Black and W. Bolosky and J. Chew,
Machine-Independent Virtual Memory Management for Paged Uniprocessor and Multiprocessor Architectures
Proceedings of the 2nd International Conference on Architectural Support for Programming Languages and Operating System (ASPLOS), 1987.


- ✓✓✓ 1. UNIX's virtual memory system relied on limited functionality for portability. Mach aimed to create an easily-portable VM system that provided more functionality than UNIX and minimized hardware-dependent code.
- ✓✓✓ 2. The high level contribution of this paper is the abstraction they provide between the machine dependent/independent code. Before this work virtual memory management was almost always machine dependent. However this work demonstrated/implemented the construction of a memory management system that was independent of hardware. This allowed the system to run on any hardware configuration; something commonplace now. Pmap is machine dependent part.
- ✓✓✓ 3. Shadow objects allow multiple processes to use the same memory for the same objects. Instead of allocating memory twice, the second process contains a shadow object which points to the original memory. This uses less memory for shared objects. A disadvantage is that when memory is rewritten to an shadow object, many new smaller memory objects may be created and this can be complicated to manage.
- ✓✓ 4. Inter-procedure Communication (IPC) was ~~not~~ an important part of their design. Each task had "ports", which other tasks were able to use to send/receive messages. Messages could be passed back and forth through the copy-on-write mechanism, which significantly reduced the IPC overhead, relative to Nucleus.
- ✓✓ Pmap acts as a cache of virtual-physical mapping, but kernel mappings need always be kept. Other mappings can be reconstructed by the machine-independent part.

5. ✓ Example of Copy on Write


i) Initial



ii) Fork



iii) C write B to B'



Edouard Bugnion, Scott Devine, Mendel Rosenblum.
 Disco: Running Commodity Operating Systems on Scalable Multiprocessors
 Proceedings of The Sixteenth Symposium on Operating Systems Principles (October 1997)


① They wanted to enable the running of commodity OS's on ccNUMA machines. At the time, new scalable computers with 10s or 100s of processors were being released, but the complexity of implementing software for those machines was too great, and thus, software was lagging behind hardware. ✓

② Disco allowed multiple OS's to run mostly unmodified on one system. It manages resources between all OS's without requiring OS to worry about non-uniform memory. Commodity OS's can be used. ✓ ✓

3) Introduced new layer between hardware and software. New layer supports virtual CPUs, virtual Physical Memory, virtual I/O devices, memory management, & Virtual Network Interface. ✓

4) Disco took advantage of existing virtual memory hardware to virtualize guests' "physical" memory. This allows guests to be moved from processor to processor, share memory, and share disk reads. However, it causes TLB misses to cause traps to Disco; which then has to emulate hardware for guest kernels, increasing the cost of TLB misses. ✓

5) In Disco, system calls to a guest OS look like this:





Nooks

Michael M. Swift, Brian N. Bershad, and Henry M. Levy
Improving the Reliability of Commodity Operating Systems
Symposium on Operating System Principles, SOSP'03, October 19-22, 2003, Bolton
Landing, New York, USA

1) The prevalence of operating system extensions — code executing in fully privileged mode yet potentially written by ~~many~~ ~~many~~ relatively inexperienced programmers with little testing — leads to many system crashes.

✓ (✓) (✓) (✓) ✓

"nice"

2 wrappers - kernel & extension

2) Introduced mechanisms to isolate driver code from the kernel by using wrappers and XPC. ✓ (✓) (✓) ✓

"agreed"

3) Basically, Nooks introduces indirection to implement isolation, interposition, object tracking and recovery.

→ XPC handles extensions, isolating them in the kernel

By using XPC, they put function, arguments, domain ~~into~~ into. And they do deferred calls, so ~~they~~ that they can do batch calls & avoid the overhead of call XPC frequently. ✓ ✓ "mmhmm"

4) Nooks Isolation Manager (NIM) is inserted between the extension and OS kernel. Provides lightweight kernel protection domain by limiting write access of extension to kernel address space. Tracks kernel resources used by extensions & cleans during recovery upon extension crash/misbehaviour. (✓) "well written!"

5) Isolating kernel from extension failures provides 99% recovery from system crashes. ✓
→ high reliability

- Disadv.


- large overhead introduced: ✓

e.g. web server: 60% ↓ in throughput.

point: Was not good at recovering the modules, just good at recovering the system.

Recovery Manager - releases resources used by extension
cleans the system state

- ① Want to support multiprocessor environment and provide sufficient primitives for protection of resources (not security!) ✓✓
Also wanted to avoid a hierarchical system structure, which can be limiting. ✓
- ② - Provided a way to separate mechanisms from policies.
- Abstraction of resources as objects ✓✓✓
- ③
 - i) Caller dependant capabilities & caller independent capabilities
 - ii) Type dependant access rights (for custom objs) & kernel access rights for common operations (like copying) for all objects.
 - iii) Kernel can provide protection without "understanding" the rights bits.
 - iv) Path-names to avoid explosion in no. of capabilities passed between procedure calls. ?
- 4) ✓✓
 - i) capabilities can't be forged. Ensured by HW or OS support
Tag: every word of memory is tagged. User process.
 1: cap
 0: other.
 can't write cap words.
 - ii) procedures → name, data (code & static data), capability list.
 only kernel can modify capabilities → put in segmented memory
- 5) flexibility to obtain extension from both representation of objects as both data & capabilities. ?



Dawson R. Engler, M. Frans Kaashoek, and James O'Toole Jr.

Exokernel: An Operating System Architecture for Application-level Resource Management

In the Proceedings of the 15th ACM Symposium on Operating Systems Principles (SOSP '95), Copper Mountain Resort, Colorado, December 1995, pages 251-266.

1. The motivation of exokernel is that ^{separate protection from management} the lower the level of a primitive, the more efficiently it can be implemented and more flexibility to implementors of higher level abstractions. High level abstractions hurt application performance, hide information & limit its functionality and hence the need for library OS. ✓

- 2) 1. securely expose hardware → ^{run} privileged instructions, access resources
2. expose allocations → request physical resources
3. expose physical names ???
4. expose revocations.
- Multiplex resources, don't abstract.

- 3) i) Secure Bindings → to securely provide applications ~~to the~~ the low level resources of the systems
- Can be implemented in hardware or software
- (ii) Downloading code into kernel to extend the capabilities of kernel.
- Packet filtering by downloading ~~code~~ code which by dynamically generated. ✓

- (4) Visible resource revocation - eg. CPU is explicitly revoked from libos. Libos saves the processor state.
- Abort Protocol - If a libos misbehaves with a resource / is disobedient then Exokernel takes that resource away from libos forcibly. ✓

⑤ ASHs - Application Specific safe handler - An implementation of downloading code into kernel. They do message processing for packets arriving over network. It runs on packet receptions and it can also initiate a message to be sent. This avoids the application to be scheduled to ~~process~~ send a response to the received ~~of~~ n/w packet. ✓

Tyler Harter, Chris Dragg, Michael Vaughn, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau
A file is not a file: understanding the I/O behavior of Apple desktop applications
SOSP '11 Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles Pages 71-83

MOTIVATION:

- Not enough analysis done in the past about home ~~workloads~~ ^{applications}. ✓
- To understand the I/O behavior and how users perceive system latency and app. performance. ✓ ✓ ✓ ✓

Contribution:

In depth case study and analysis of the I/O behaviour of modern home-user applications. The analysis has several interesting revelations. (file is not a file, seq access is not seq, writes are forced, multiple threads perform I/O, frameworks influence I/O etc). ✓ ✓ + renaming is popular + auxiliary files dominate


Solution / Technique advan./disadv.

The results have a strong ramifications for the design of next generation local and cloud-based storage systems. Constructing the iBench task suite tw-fold: 1. representative of tasks performed by home users. ??
← copied from abstract

2. Sol/Tech:

- 4) The authors study four categories of the behavior of the iBench suite,
- 1) Nature of files
 - 2) Access patterns
 - 3) Transactional properties ✓
 - 4) Multi-threading behavior

5) iTunes attempted to block efforts to trace it, but the authors found a way to circumvent this by using gdb to force a system call into a NO-OP. ★



1. The motivation for LFS is to match the ^{write} performance (speed) of the disk with that of the CPU.

2. The use of log to club small random writes into large sequential writes to achieve ^{very good} utilization of the disk bandwidth.

3. They combined ~~threading~~ ^{Interleaving} and compaction into a hybrid approach where they grouped blocks into segments and threaded the segments. They periodically ran a cleaner to find underutilized segments and write their data at the front of the log.

4. i) Identifying live blocks using Segment summary block which contains for every datablock the file inode and offset.

- Some more optimizations like maintaining ⁱⁿ versions in inode map to avoid costly identification of live blocks (eg. truncation of files).

ii) Using 2 timestamps for each of two CR regions and writing to alternate CRs.
 Adv - No inconsistent CR

DisAdv - The CR commit interval is 30s which affects freshness of data ^{segments}.

⑤ ^{As} write cost $\approx \frac{2}{1-u}$ ^{with} ~~making~~ cleaning underutilized ~~blocks~~ ^{segments} would give high performance. ~~The above~~ ^{As} ~~cost~~ ^{As} ~~cost~~ ^{As} having more underutilized blocks, increases cost per byte storage approach to have segments into ~~segments~~.

⑥ A cost benefit segment cleaning policy is proposed.
 * The cost benefit policy involves bimodal distribution of segments,

$$\frac{\text{benefit}}{\text{cost}} = \frac{(1-u) \times age}{(1+u)}$$

✓ To use cost-benefit policy, a data structure which records #live bytes and recent modified time called segment usage table. is used

Segment Usage Table

✓

Vijayan Prabhakaran, Lakshmi N. Bairavasundaram, Nitin Agrawal, Haryadi S. Gunawi, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

IRON File Systems

Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP '05), Brighton, United Kingdom, October, 2005.

Motivation:

Present disk failure model for commodity system is faulty - It assume that disk either work or fail.

✓ In reality, disks exhibit complex disk failure pattern.

* Present paper does have following contributions

2. The paper presents a more nuanced failure model for disks that provides partial failure support and develops a framework to determine how a file system handles failures. Analyzing commodity FSes show that their handling of failures is buggy and inconsistent. Finally, a more robust file system, "ixt3", is demonstrated to show that error-checking techniques such as checksumming provide substantial reliability improvements at a low cost.

3. To determine how the FS handled disk failures, the authors defined the IRON Taxonomy for various levels of detection and tabulated the behaviour for various workloads.

Adv: The tabulation (complex picture) gives a very detailed picture of the behaviour ✓

4. Build a prototype version of an IRON filesystem (ext3)

✓ that is robust to various kinds of partial disk failures. Evaluation of ext3. ✓
↳ + checksumming, metadata replication, parity

5. Works! :- Transactional checksums for ordering journal data. This improves the performance of journal data I/O.

✓ (checksum the transaction blocks + write the checksum with the transaction commit)

1. Motivation :

To understand modern file systems' (ZFS as a representative FS) reliability mechanisms for dealing with disk & memory corruptions. ✓ ✱

2.) ZFS while does checksum verification for on disk blocks, it doesn't use these same checksums in memory to protect against in memory corruption. ✓ ✓

3. Solution/technique: ZFS uses checksums for all blocks on disk. This checksum is stored in the parent block for a given block; by doing this rather than storing it alongside the checksummed block, it becomes possible to detect errors such as misplaced writes. Maintaining and checking these checksums does cause some overhead, and updating a checksum necessitates propagating new checksums up the file system, though the atomicity of this is guaranteed through copy-on-write semantics. ✓ ✓

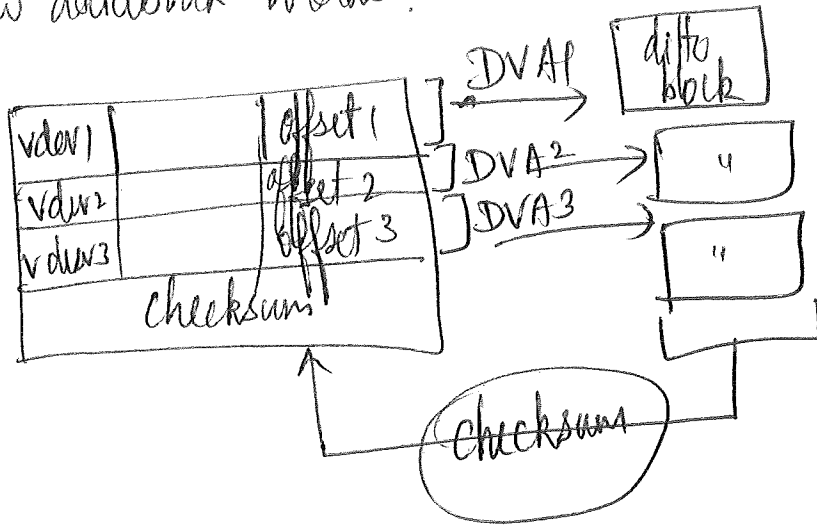
4.) ZFS in addition to checksum, use replication to maintain reliability

→ metadata - 2 ditto blocks
global metadata - 3 ditto blocks - ✱
provide

ZFS uses tiered software based RAID system by having virtual # storage pools → 2 pool and logical).
zpool → it consists of virtual devices (physical and logical).

→ RAID Z (similar to RAID 5)
distr → variable stripes. ✓

5. How datablock works?



Motivation - Crashes are rare but possible in FS. Pessimistic approaches anticipate crashes and do the careful logging to the journal before writing to (FS) disk. These ordering points make the FS perf slow. Main motive is to avoid ordering as much as possible and also decouple ordering from durability. ★

High Level Point: Probability of inconsistency because of crash is high for workloads with a high number of random write I/Os. for probabilistic crash inconsistency. ★

Maintain consistency to same extent as pessimistic but nearly same performance with probabilistic consistency. ★

Solution:
① Add ^{data} checksums to avoid write ordering
② Use asynchronous durability notification ★ to delay checkpointing until durable commitment has happened.
③ Ordered Journaling. (transactions)

Advantage 1. disk itself determines when to flush, don't need to listen to higher-level FS. ★

Disadvantage: - Recovery takes a long time X (not much longer than normal journaling).
- greater code / memory overheads, + stale data

Essential operations of async:

When OptFS calls on async, it writes, [?] ~~in order~~, the data blocks, journal meta-data, and journal commit block. Checksums are added to the commit block so these can become durable in any order. Then, ^{only} after receiving notification (or a simulation of this) that the blocks are durable, will it 'checkpoint'; ~~it~~ ~~write~~ write the metadata to disk, which occurs in background.

1. ~~For~~ DRAM \rightarrow more imp, need full performance of DRAM \rightarrow ✓
Fast Recovery on crash \leftarrow durable + available ✓

2. ~~Technical~~ \rightarrow make use of scale + try to provide * ✓
low latency \rightarrow * ✓
3) ~~single recovery master~~ \rightarrow impact: enable new class of * applications * ✓

2. i) log structured storage (both in back-up & in-memory) ✓ ✓
ii) Extremely low-latency storage ✓ ✓
iii) Randomization techniques to make decisions in scalable fashion ✓
+ Pick best of randomly chosen ✓

3) Disadv.

~~cannot do~~
- Very hard to achieve such short switching times
and very hard to optimizations like kernel bypass &
faster NICs. ✓ but badly worded

Adv

- Harness the power of several nodes to retrieve data & perform recovery. ✓

4) Used DRAM ✓ to support low latency ✓
dis-adv. Volatile so backup efforts made to avoid data loss ✓
o Can potentially be used for massive r.t. computations ✓

5.) Recovery Process


1. When DRAM crashes, controller asks all nodes what pieces of the missing data they have ✓

2). assigns a subset of recovery masters, tells them which segments to incorporate into their own memory ✓

→ recovery is spread out among many disks ✓

Done, serve requests from restored data ✓

+ Servers keep 'wills' indicating how data should be grouped ✓



Patterson, D., Gibson, G., and Katz, R.,
 A Case for Redundant Arrays of Inexpensive Disks (RAID)
 Proceedings of the 1988 ACM SIGMOD Conference on Management of Data, Chicago
 IL, June 1988.

CAN WRITES BE PARALLEL IN RAID 1?

- 1) * Single disk become bottleneck (because seek latency & bandwidth are not improving at a good rate compared to processor speed).
 ✓ Thus, how can we use ~~multiple disks~~ multiple disks/array of inexpensive disks & collectively use them to get greater bandwidth for I/Os

- 2) a) ~~categorize~~ model & categorize ~~a~~ different techniques
 b) take advantage of parallelism
 ✓ c) 12 times higher bandwidth than state-of-art device at that time, ?
 ✓ d) good reliability ~~evaluation~~ of each case. ✓

- 3) low reliability can be dealt with using replication. The technique is simple - buy another disk, and copy the first disk over entirely. An expansion of this is data redundancy, used in all levels. *
 ✓ The advantage is it can improve reads and reliability, while the downside is it requires extra disk space (whether replicating or checksumming, etc.) ✓

- 4) Parallel access improves aggregate bandwidth when → ?
 ✓ data is replicated or when partitioned (RAID-0). RAID-0 lacks redundancy, however RAID-2 can benefit from parallel reads. ✓
 ✓ Write performance of RAID-2 however is no worse than a single disk assuming no bus contention.

- 5) From RAID-0 to RAID-5 the models progress by improving disk utilization (although the performance for small writes decrease as we go through the levels) and reliability through parity disks. RAID-5,

improves upon RAID 3 & 4 by distributing the parity.

RAID 0, RAID 1 with RAID 0 & RAID 5 by abstracting from file system. Automatically manages storage

1. X? Motivation: a two-level storage hierarchy implemented inside a single disk-array controller → RAID 0 vs RAID 5 performance vs cost

? HP AutoRAID: fully redundant storage system ?

2. Hig Contribution: ~~A RAID~~ A storage hierarchy is developed. write-active (RAID 0) and write inactive (RAID 5);

3. Active / freq. accessed data is kept in RAID 0 (mirroring) RAID 5

Transfer of data between the diff. RAID levels according to the need of the workloads. Inactive / less " " "

Data stored as: RB, segments, PEXes, PEGs, LUNs.

4. Advantages:

- Automates the management of RAID groups.
- Good job at balancing performance vs disk utilization (active vs inactive) move data between groups)
- Scalable: possible to add new disks ~~at the time~~ on the fly.
- Upgrade: Replace disks without disrupting operations.
- Extremely easy to use

Disadvantages

- Has variable response times. Not suitable for some workloads.

[Faint, illegible handwritten text at the bottom of the page]