



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

Fault-tolerance

6.033 Lecture 14

Frans Kaashoek

With slides from Sam Madden

Where are we in 6.033?

- Strong form of modularity: client/server
 - Limits propagation of effects
 - In a single computer using OS
 - In a network using Internet
- Two limitations:
 - Isolates only benign mistakes (e.g., programming errors)
 - No recovery plan

Extending C/S to handling failures

- Can we do better than returning an error?
 - Keep computing despite failures?
 - Defend against malicious failures (attacks)?
- Rest of semester: handle these “failures”
 - Fault-tolerant computing
 - Computer security

Plan for fault-tolerant computing

- General introduction: today
 - Redundancy/Recovery/Replication
- Transactions: next 4 lectures
 - updating permanent data in the presence of concurrent actions and failures
- Replication state machines: 2 more
 - Keep computing despite failures

Windows

A fatal exception 0E has occurred at 0028:C00068F8 in PPT.EXE<01> + 000059F8. The current application will be terminated.

- * Press any key to terminate the application.
- * Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue

Availability in practice

- Carrier airlines (2002 FAA fact book)
 - 41 accidents, 6.7M departures
 - ✓ 99.9993% availability
- 911 Phone service (1993 NRIC report)
 - 29 minutes per line per year
 - ✓ 99.994%
- Standard phone service (various sources)
 - 53+ minutes per line per year
 - ✓ 99.99+%
- End-to-end Internet Availability
 - ✓ 95% - 99.6%



Barracuda® 7200.10

Experience the industry's proven flagship perpendicular 3.5-inch hard drive

80 GB to 750 GB • SATA 1.5Gb/s or 3Gb/s and PATA 100


Key Advantages

- First 3.5-inch drive to utilize capacity- and reliability-boosting perpendicular recording technology
- First drive to reach 750 GB—a full year ahead of competition—enabling new solutions for data-intensive applications.
- Industry's most proven and established desktop hard drive available today—more than 16 million shipped to date*
- "One-stop shopping" with a broad range of capacity, cache and interface options for all your computing needs
- Best-in-class environmental specifications and reliability features
- Adaptive Fly Height offers consistent read/write performance from the beginning to the end of your computing workload.
- Clean Sweep automatically calibrates your drive.
- Directed Offline Scan runs diagnostics when storage access is not needed.
- RoHS-compliant design assures an environmentally conscious product.
- Enhanced G-Force Protection™ defends against handling damage.
- Seagate® SoftSonic™ motor enables whisper-quiet operation.

Best-Fit Applications

Desktop and High-Performance PCs

- Gamer PCs
- Workstations
- High-end PCs
- Desktop RAID
- Mainstream PCs
- Point-of-sale devices/ATMs
- USB/FireWire/eSATA personal external storage



*16 million Barracuda 7200.10 drives shipped as of 4/16/07

Contact Start-Stops	50,000
Nonrecoverable Read Errors per Bits Read	1 per 10 ¹⁴
Mean Time Between Failures (MTBF, hours)	700,000
Annualized Failure Rate (AFR)	0.34%

Data Sheet

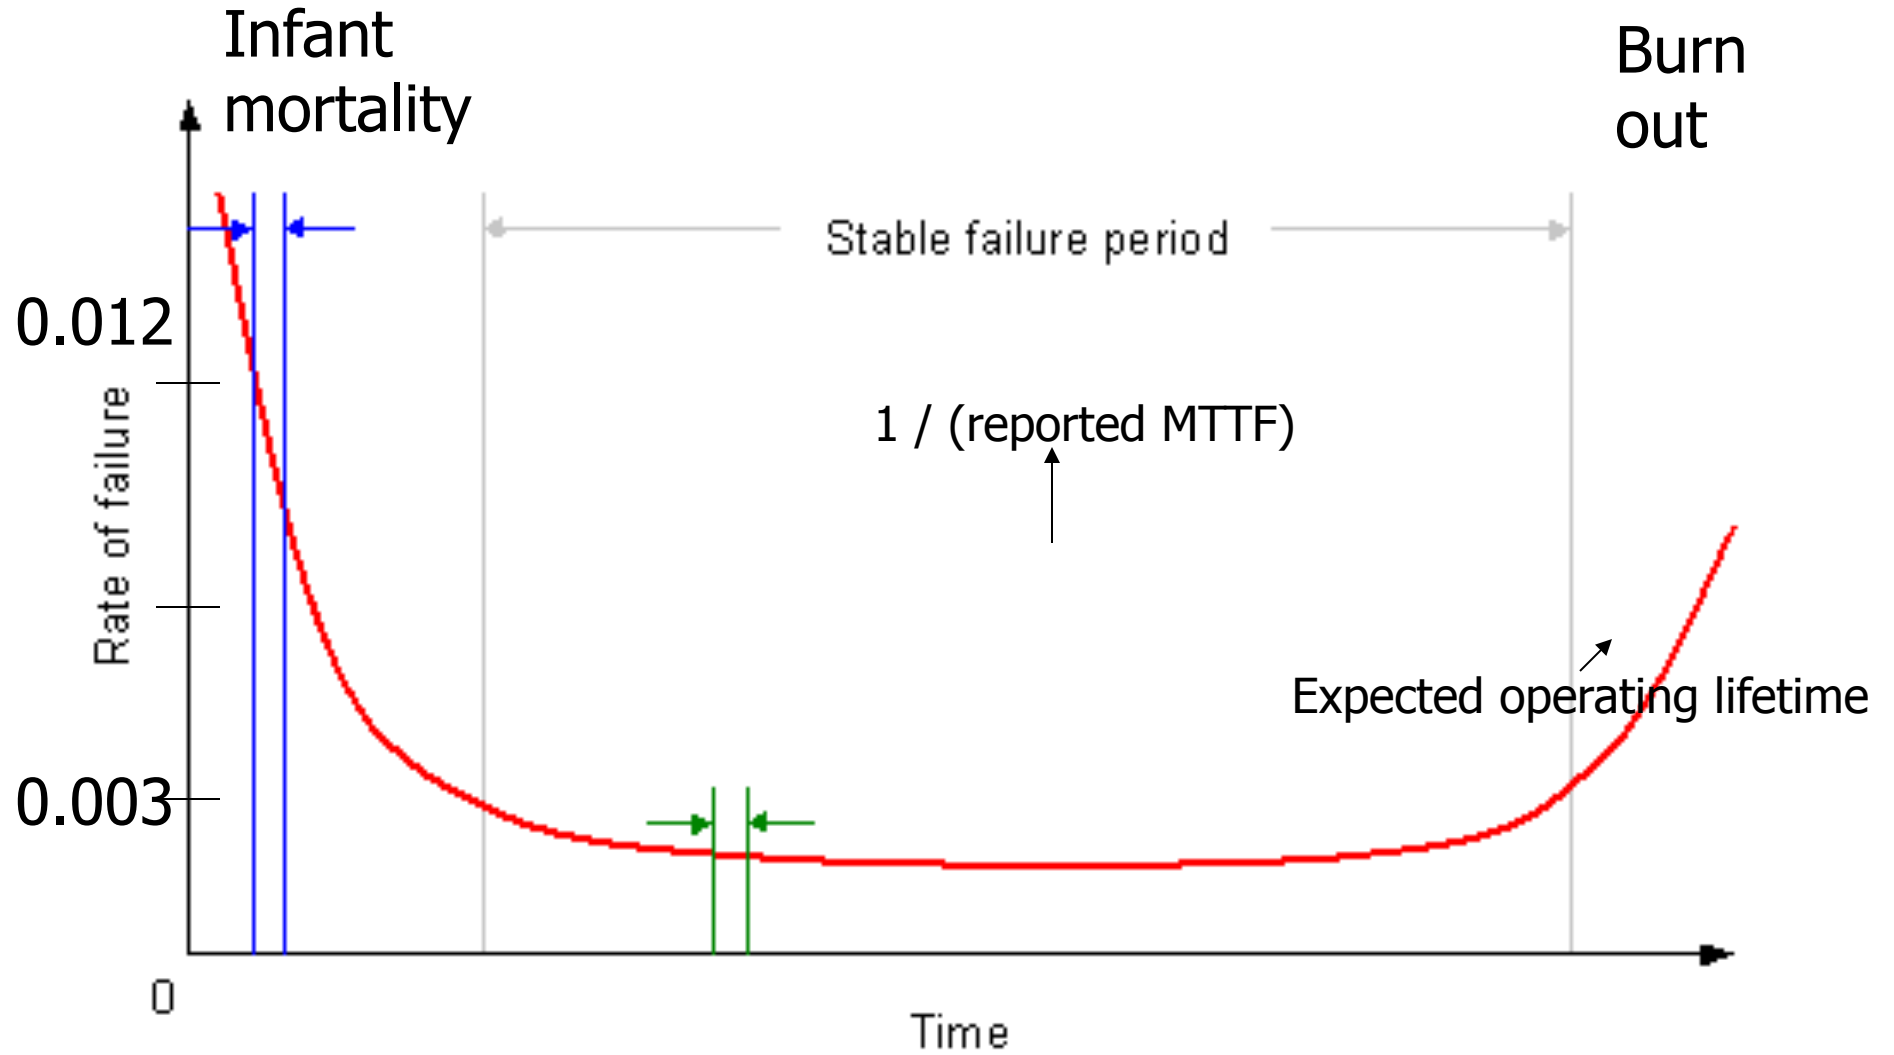
Barracuda® ES.2

High-capacity, business-critical
Tier 2 enterprise drives


1 TB, 750 GB, 500 GB and 250 GB • 7200 RPM •
SATA 3Gb/s, SATA 1.5Gb/s and SAS 3Gb/s

Reliability/Data Integrity	
Mean Time Between Failures (MTBF, hours)	1.2 million
Reliability Rating at Full 24x7 Operation (AFR)	0.73%
Nonrecoverable Read Errors per Bits Read	1 sector per 10E15
Error Control/Correction (ECC)	10 bit
Interface Ports	
SATA	Single
SAS	Dual

Disk failure conditional probability distribution




Bathtub curve



Human Mortality Rates (US, 1999)

From: L. Gavrilov & N. Gavrilova, "Why We Fall Apart," *IEEE Spectrum*, Sep. 2004.
 Data from <http://www.mortality.org>

Disk Age vs. Pr(≥ 1 Reported Read Failure)



Bairavasundaram et al., SIGMETRICS 2007

Relative frequency of hardware replacement

COM1	
Component	%
Power supply	34.8
Memory	20.1
Hard drive	18.1
Case	11.4
Fan	8.0
CPU	2.0
SCSI Board	0.6
NIC Card	1.2
LV Power Board	0.6
CPU heatsink	0.6

10,000
machines

Pr(failure in
1 year) $\sim .3$

Fail-fast disk

```
failfast_get (data, sn) {  
    get (s, sn);  
    if (checksum(s.data) = s.cksum) {  
        data ← s.data;  
        return OK;  
    } else {  
        return BAD;  
    }  
}
```

Careful disk

```
careful_get (data, sn) {  
    r ← 0;  
    while (r < 10) {  
        r ← failfast_get (data, sn);  
        if (r = OK) return OK;  
        r++;  
    }  
    return BAD;  
}
```

Replicated Disks

write (sector, data):

 write(disk1, sector, data)

 write(disk2, sector, data)

read (sector, data):

 data = careful_get(disk1, sector)

 if error

 data = careful_get(disk2, sector)

 if error


 return error

 return data

Technical specifications

Processors	2–16 per node Intel Itanium processor 9100 series processors, 1.6 GHz single core processors
Cache	12 MB L3
RAM standard/maximum	Minimum: 4 GB Maximum: 16 GB (32 GB ²)
RAM type/speed	PC2100 ECC registered DDR266A/B
ServerNet I/O	Minimum: 10 Maximum: 60
I/O adapters supported	Fibre Channel, Gigabit Ethernet
Fibre Channel disk modules	14 disks per module
Disk drives supported	146 GB and 300 GB 15K RPM Fibre Channel internal hard disk drives HP Disk Array family (e.g., XP24000, XP20000, XP12000, and XP10000 disk arrays)
Standard features	N + 1 power supplies N + 1 fans

2 Although 32 GB is available, the total available memory is limited to 16 GB.



How about an error in software?

- Big problem!
- Software for fault tolerant systems must be written with great care
 - Stringent development practices
 - Well-defined stable specification
 - Modeling, simulation, verification, etc.
 - N-version programming is tricky
- Will also be a problem for secure software
- Good design: small fraction is critical