

important concepts from last time:

- capacity
- availability
- performance
- quality attribute
- reliability

fault tolerance is a quality attribute of a system that describes ability to tolerate faults without degradation of system.

fault tolerance is preventing faults from being (system) failures

failure: not meeting its specification

error: system state that may lead to failure

fault: caused an error

faults can originate from internally or externally

faults may have several reasons
accidental
malicious

fault detection can be transient → difficult to identify, since it does not always appear in particular, needs debugging intermittent environment

fault/failure models

fault metric for measuring fault tolerance

at-resilient system contribute up to and including to faulty components

things which are going wrong in interesting processes:

- crash server halts
- omission server fails to respond/receive/send
- timing response outside specified time interval
- response server response (value/control flow) incorrect
- arbitrary arbitrary response at arbitrary time

faulting types:

- fail-stop crash failures, reliably detectable
- fail-stop actually reliably detectable
- fail-silent omission/crash failures: crash cannot tell what was wrong
- fail-safe arbitrary, yet benign failures → can not do anything
- fail-arbitrary arbitrary, with malicious failures

message loss = lack of delivery

fair-loss delivery: if correct process p infinitely often sends message m to correct process q, then q delivers m an ∞ number of times

reliable delivery: if correct process p sends m to correct process q, then q actually delivers m

message duplication

finite duplication: if correct p sends m finitely many times, correct q cannot deliver m infinitely

stable delivery: if correct p sends m to correct q, q delivers m an ∞ number of times

no duplication: no message m is delivered more than once

message creation

no creation: if q delivers m with sender p, then m was previously sent to q by p

authenticity: if q and p are correct, and m is delivered to q with sender p, then p previously sent m to q.

UDP guarantees FLD

TCP guarantees RD

a unicast service is reliable if it satisfies the requirements for a perfect lib:

$$RD + \underbrace{ND + NC}_{\text{integrity}} \downarrow \text{validity}$$

validity: message actually delivered

integrity: message delivered & not message, no duplication

process resilience:

flat group all nodes run same process

hierarchical group coordinator + workers
could be server + clients

consensus: each correct process executes the same commands in the same order

a group communication (multicast) is reliable if it satisfies

validity if correct process multicasts m, it will eventually deliver m

integrity message is delivered at most once, and if delivered, sender belongs to group of m and m was previously sent/multicast by S

agreement if correct process delivers m, all correct processes will eventually deliver m

consensus problem: set of correct processes make the same decision i.e. have agreement

consensus guarantees:

termination every correct process eventually decides some value

validity if process decides v, then v was proposed by some process

integrity no process decides twice

either of agreement no two correct processes decide differently

uniform agreement no two processes decide differently

fault tolerance tactics

fault masking

error detection

acceptance test

exceptions

watchdog timers

ping/echo

heartbeat

error diagnosis

error/fault estimation

error recovery

forward error recovery find new error-free state to continue predictable in terms of time & energy → suitable in real-time systems

backward error recovery restore previous error-free state e.g. checkpointing, suitable for high transaction failures

redundancy

- hot spare updates are synchronous with primary
- warm spare updates are updated periodically
- cold spare off-line, not updated, needs to be brought alive

for failure masking

add extra bits error correction

time repeat failure operations

physical redundancy multiple hardware/software components

active all replicas performing calculation (updates) in parallel

passive replica only activated if previous replica fails

fault prevention:

sound development methodologies coding standards, test procedures, design patterns

temporary removal from services occasional restores

transactions: ACID

atomic commit/abort & rollback, indivisible

consistent maintain system invariants

isolated concurrent states appear serialized - shared states not observable

durable committed transaction is persistent i.e. system failure cannot undo result of transaction

availability + reliability improvements mainly comprise

- selecting uptime
- restricting downtime