



Microservices and DevOps

Scalable Microservices
Stability Patterns

Henrik Bærbak Christensen



Anti --- Antipatterns

- Nygaard lists 14 antipatterns
 - Users, blocked threads, slow response, chain reactions, etc.
- So – how to combat these???
 - The *Stability Patterns*
 - (and the ‘Remember this’ section in each antipattern chapter, which mentions stability patterns that are not mentioned in this chapter, hmmm)

- Note: Figure from *first edition!*
- Missing in 2nd Ed, and a bit different pattern set
- But – nevertheless...

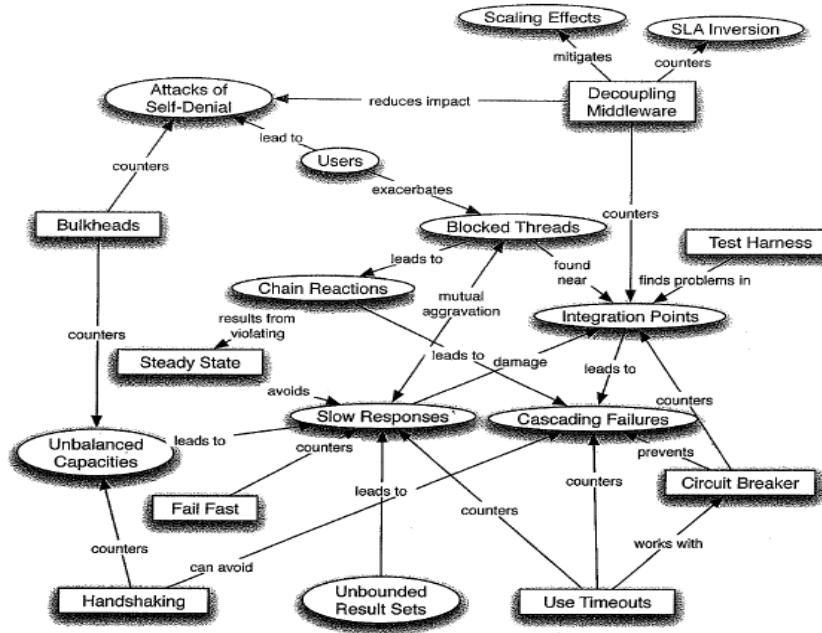


Figure 3.1: INTERACTION OF PATTERNS AND ANTI PATTERNS



Is All This Clutter Necessary?

AARHUS UNIVERSITET

- As stated by Nygard:
- *... handling all the possible timeouts creates undue complexity in your code. It certainly adds complexity.*
- *... Your users may not thank you for it, because nobody notices when a system **doesn't** go down, but you will sleep better at night.*

- *Timeouts*: Guard *any* call to remote units with timeout to avoid waiting forever on an answer that never arrives
 - Well placed time-outs provide fault isolation: your code does not break due to a failure in another subsystem.
 - Beware of vendor supplied APIs, sometimes they have forgotten to add time out parameters ☹
 - Or it is in the API, but not implemented ☹ ☹ ☹

- *Timeouts*: Guard any call to remote units with timeout to avoid waiting forever on an answer that never arrives
 - Apply to Integration Points, Blocked Threads, and Slow Responses
 - Consider *delayed retries*
 - Network issues take time to go away – do not retry again immediately
 - See ‘Retry’ pattern (that I introduce) later...

Timeouts on Locks

Circuit Breaker

- The idea of a *fuse*: Burn before your house does!
- *Circuit breaker*: Da: Maksimalafbryder
 - ... en automatiskafbryder, som bruges til udkobling af overstrømme.
 - Består af bryderdel og relædel.
Relædelens funktion er at måle strømmen og i tilfælde af en overstrøm udkoble afbryderen
 - HFI relæ



Circuit Breaker

- *Circuit Breaker*: Wrap dangerous operations with a component that can circumvent calls when the system is not healthy. Differs from retries, in that circuit breakers exist to *prevent* operations rather than to re-execute them.

Circuit Breaker

- **Closed state:**

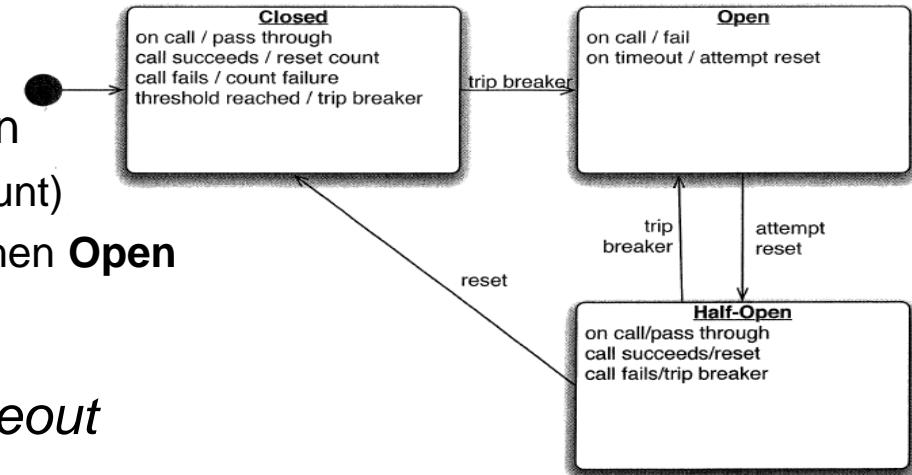
- Execute operation as normal
 - If operation fails (timeout) then
 - Note it (increment failure count)
 - If failure count > threshold then **Open**

- **Open state:**

- *Fail fast, avoid the wait for timeout*
- After a set time, switch to **Half-open**

- **Half-open state:**

- Execute operation again
 - if fail then goto **Open** state immediately
 - if success then goto **Closed** state



Exercise: Why the need of the half-open state?

Circuit Breaker

- Is a way to make "graceful degradation"
 - Degrade functionality when under strain
 - Avoid DogPile, as it gives producer service time to recover
 - Important to involve stakeholders
 - "What to do if in the open state?"
 - What to do if we cannot verify credit card?
- Remember
 - Use together with timeouts
 - Expose, track and report state changes
 - Log any "popped fuses"!

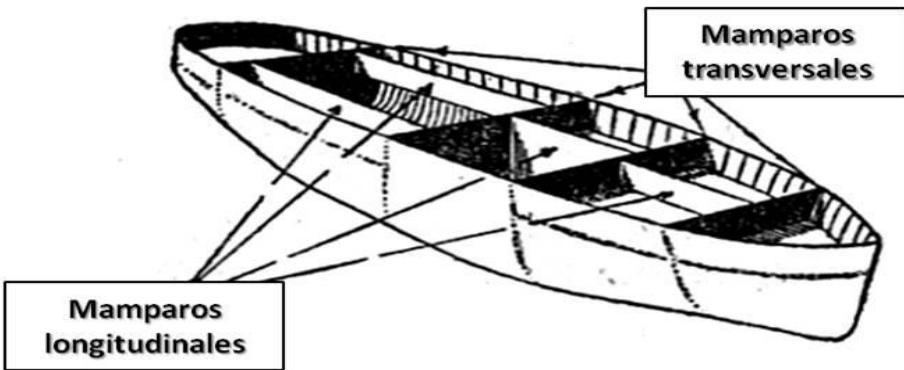
- NetFlix is a large micro service system, based upon circuit breakers.
 - Open sourced their circuit breaker implementation w. dashboard



- ... is
 - A *lightweight, easy-to-use fault tolerance library inspired by Hystrix, but designed for Java 8 and functional programming.*
- You can pick and choose just the piece you want
 - CircuitBreaker: Nygard's pattern in it's *frequency* form
 - Bulkhead: Limit number of concurrent executions
 - RateLimiter: Limit rate of requests (or queue them)
 - Retry: Retry call N times with M mS delay between
 - TimeLimiter: Nygard's *Fail Fast* pattern
 - Cache: You guessed it ☺

- Da: Skot

- *Partitions that can be closed preventing water from moving from one section to the next*
 - *Damage containment*



- *Bulkheads*: Partitioning a system so failures in one part does not lead to system failure
- Simplest (most common) form: **Redundancy**
 - Have two or more servers handling the load

- Mission-critical form:
 - Pool of servers/services reserved for critical use while the rest are available for non-critical use
- Example
 - Servers dedicated to airline check-in (critical)
 - Others serve flight status checking (non-critical)
- Liability (see next slide)
 - You now have two disjoint resource pools that are subject to *unbalanced capacities*
 - I.e. you need more reserved capacity

Bulkheads

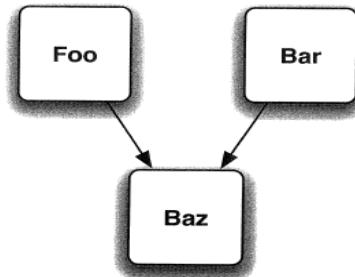


Figure 5.2: HIDDEN LINKAGES

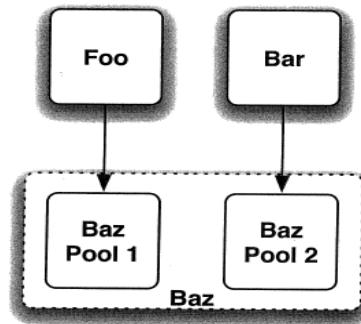
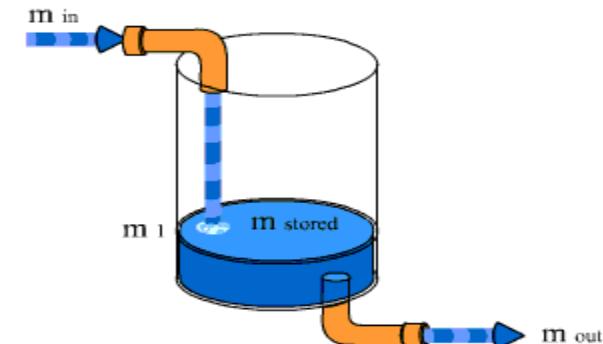


Figure 5.3: PARTITIONED SYSTEM

Steady State

- *Steady state*: For every mechanism that accumulate resources, some other mechanism must recycle that resource.
- If not, accumulated resources outgrow capacity
 - Log files, DB rows, caching, ...
- If capacity is exceeded - bad things happen



Examples

- MongoDB
 - Allocated disk in max 2GB chunks
 - If it cannot it simply stops processing write requests, only processing read requests
 - Will not reenable write requests until after restart
 - (and data purge and compacting is a write request 😊)
- Docker container logs are written to storage
 - So if you never look at them, they just grow

Steady State

- Another potential definition
- *Steady state*: A system should be able to run indefinitely without intervention
- Otherwise, you get used to *fiddling*
 - *Which leads to what Nygard terms "oh-no-second"*
 - The split second when you realize you have hit the wrong key, shut down the wrong server, deleted the wrong DB table...
 - My personal fear was "db.GFKRE003.drop()"
 - 6TB of data from 4 years of data collection

My own Fiddling trick

- A linux shell looks like ... a linux shell
- Make staging machines and production machines look different!
 - Production machine always have awful color choices!
 - I do not mistake one for the other!
 - I hope...

- Data purging
 - Remove old data from the DB
 - Can be pretty tricky in RDB
 - Referential integrity, orphaned rows, ...
 - And perhaps even more so in NoSQL
 - Log files
 - If not purged, you run out of disk space
 - Java.io.IOException!
 - Dump a stacktrace in the face of the user ☺
 - Review the Log4J RollingFileAppender for non Docker use
 - For Docker, you write to StdOut which is stored ☹
 - Rewire to a logging system like ELK or Humio instead
 - » Which then run out of disk space ☺

- *Fail fast*: Check resource availability at the start of a transaction, and fail immediately in case any is not available.
- If not, you waste CPU and human time doing stuff that will eventually have to be redone or undone
- Cook's: *mise en place*
 - Find all ingredients before starting
 - Or the fish will burn while to try to find the chili paste...

- Part of fail fast is also to validate human entered values as best possible before proceeding
 - Typically values entered in web form or similar
 - Avoid connecting the DB and do a query only to find that one of the query parameters were null...
- Fail Fast is a way to combat *slow response*
- Example:
 - CPF system will fail immediately if a key is not found in file.

- *Let it Crash*: Create system-level stability, by sacrificing service-level(*) stability. The cleanest state your program has, is right after startup.
 - Nygard writes ‘component-level’, but guess that is synonym with our ‘services’.
- **Require**
 - Limited granularity: isolate the crash to the service
 - Avoid cascading failures
 - Fast replacement: Only ‘let it crash’ if restarts are quick
 - Supervision: Don’t do local restarts, monitor on higher level
 - Reintegration: Consider how ‘back to work’ is orchestrated

Inspired by the Erlang programming language

- We will return to HEALTHCHECK in Dockerfiles and compose-files.
- Allows
 - Granularity is the container
 - Restart by the swarm (restart time is within the ~seconds)
 - Simple supervision and reintegration is built-in in the swarm's stack
 - Restart-policy

```
version: "3.8"
services:
  redis:
    image: redis:alpine
    deploy:
      restart_policy:
        condition: on-failure
        delay: 5s
        max_attempts: 3
        window: 120s
```

Handshaking

- *Handshaking*: Allows a client to assess whether a server has the capacity to answer a request, essentially providing the server the ability to tell the client to "back off"
- Compare
 - Ping-echo and Heartbeat tactics
- Actually not implemented in most modern protocols, like HTTP, RMI, etc.
 - You will probably have to implement it yourself
 - Watch out for doubling the traffic: 'can I do (1), then I do (2)'
 - Any HTTP call is expensive...

- *Test Harness*: A substitute for the remote end of an integration point that produce out-of-spec failures
 - Like
 - Refuse connection, accept and then die, packet loss
 - Slow responses, garbage responses, protocol errors
 - Send one byte every 30 seconds (DOS attack), send Megabyte long replies instead of Kilobyte long replies
 - Send HTML instead of XML, refuse all authentication credentials, ...

- Da: Seletøj
- *Integrations test*: Testing that the parts are working together to identify integration failures
 - Obviously closely related to integration points and the large set of stability antipatterns
- Nygard: No – it is not strong enough
 - Integration tests test within-spec failure modes...
 - Operations excel in ***out-of-spec failure modes*** !!!
 - You do not get an error code from the remote server, it simply does not reply, or provide *slow response*

Decoupling Middleware

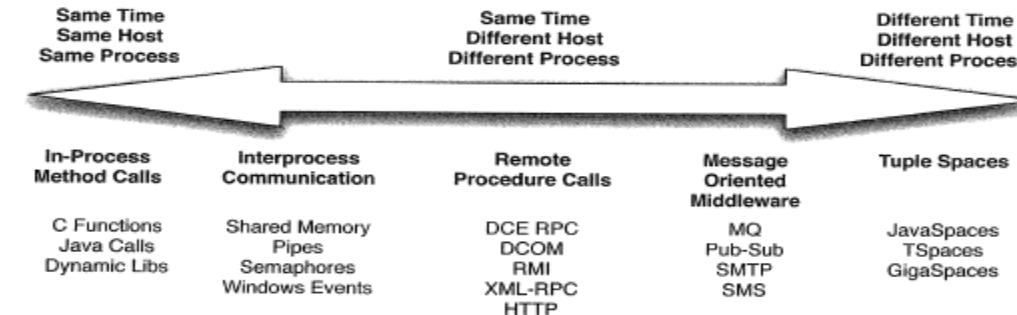


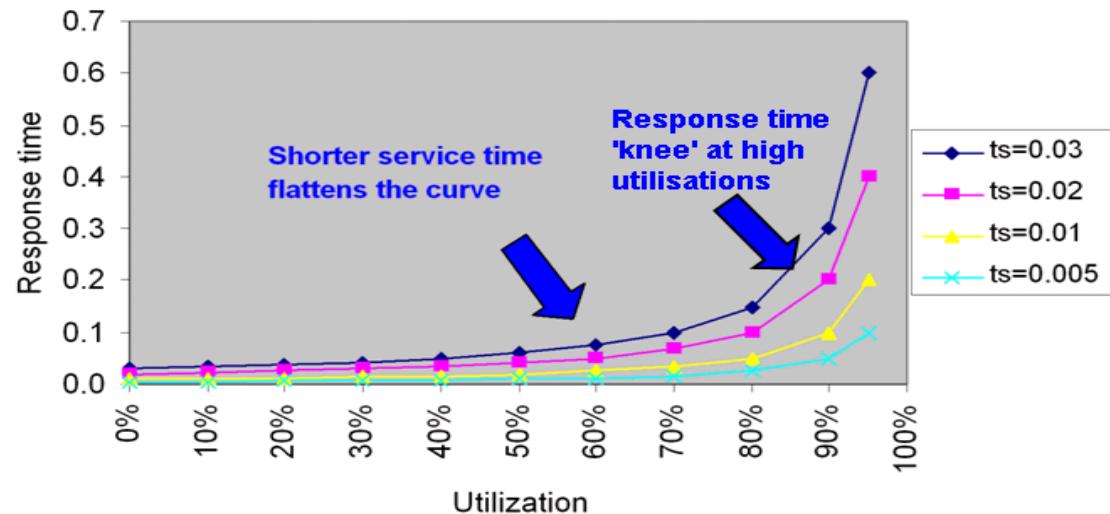
Figure 5.4: COUPLING SPECTRUM OF MIDDLEWARE

- Middleware decisions effect the implementation cost of systems significantly
 - Learn many architectural styles to ensure you pick the right one

- *The world can always create more load than you can handle...*
 - There is no difference between ‘really, really slow’ and ‘down’.
- ***Shed load:*** Refuse new requests, if load gets too high.
 - Similar to ‘fail fast’ but you do not fail on service-level but on request-level
 - Not ‘time out exception’ but ‘503 service unavailable’
 - Guard calls to shed-loading services with circuit breakers or timeouts or ...
 - Meaning of ‘Too high’?
 - Monitor own SLA to determine the tripping point

Shed = skille sig af med

- Back in the Queue Theory stuff...
 - Systems with ‘randomly timed requests’ will follow this distribution as the workload increase – response times increase exponentially ☹
 - Shed load when you hit the upper parts of the ‘knee’



Create Back Pressure

- *Every performance problem starts with a queue backing up somewhere...*
 - Little's law: $L = \lambda \times R$
 - L = number of requests in queue (Think Føtex kasse queue)
 - R = response time of request (Think 'time until I leave')
 - $R = W + S$, wait time + service time (in line + getting served)
 - λ = arrival rate (arrival per sec, of custom.)
 - So
 - λ is constant (influx of requests from the world)
 - If response time gets longer... (S large (dankortterminal i stykker'))
 - ... then more requests in queue
 - ... and then queue eats all memory => crash

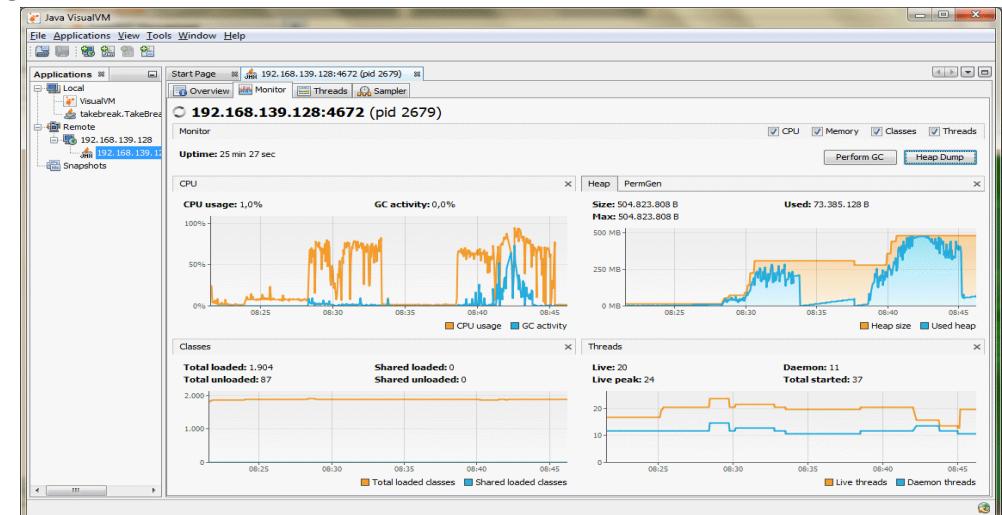
Create Back Pressure

- So
 - ... we do not want *unbounded queues*!
- Bounded queue, what to do with 'out of space'?
 - Pretend to accept new item, but actually drop it
 - Accept new item, drop something else
 - Refuse the item (= shed load)
 - Block request (producer) until there is room in queue
- Dropping options
 - In many real-time systems, only latest reading is interesting
 - Aircraft flight control – who gives a damn about that angle the rudder was in 30 seconds ago – more interesting what is *now*

Create Back Pressure

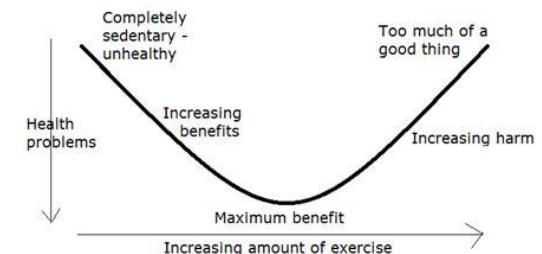
- Block producer option
 - Introduces ‘flow control’, applying ‘back pressure’ upstream
 - (probably) propagates all the way back to the client, who will be throttled down until queue releases...
- *Create Back Pressure*: Use finite queues and block producers if queue overflows, to slow down instead of crashing [Own definition]
 - Alternative to ‘shed load’
 - May lead to ‘blocked threads’ obviously
 - Are we crashed or just extremely slow?
 - Only use ‘within system boundaries’
 - Use ‘shed load’ across system boundaries instead, like open internet

- My Lyon airport load test
 - MongoDB was slow, so node memory exhausted
- Solution
 - I set the ‘prefetch’ to 1!



- Now Java connector does not fetch next until MongoDB has stored it (acknowledge message)
- Thus, the RabbitMQ itself acts as Queue, rather than the consuming node’s heap acting as queue!
- Thus *applying back pressure upstream until I hit a service that can actually handle the pressure...*

- Steam engines: They can run so fast, they will break!
 - Engineering device: Put a speed limit on: the ‘governor’
- Governor: Stateful and time-aware control plane logic that prevents a system from exceeding its safe limits. Actions within safe limit should be fast, outside increasing resistance must be applied
 - *Shutting down instances is unsafe, deleting data is unsafe ...*
 - Antidote for the ‘force multiplier’ antipattern...
 - Safe limits is kind of ‘U-shaped curve’
 - Slow actions down if moving outside the bottom of the U curve...





Bounded Result Sets

- Nygard does not mention a pattern to combat *unbounded result sets*
 - Should be reasonably obvious...
- *Bounded Result Sets*: Return large results sets in *chunks* that can be *iterated*. (Pagination)
 - *Not "get bible", but "get bible, page 7564, page 7565, ..."*
 - *Mongo: use find().skip(n*page).limit(n)*

- Perhaps too obvious, but...
- Why does Nygard not mention ‘Retry’ as a pattern?
 - Only indirectly in ‘time outs’ pattern
- *Retry*: If a request fail/time out, then retry the request some time later a number of times before giving up
 - Only do it, if it makes sense! And who does the retry?
 - Fixed retry intervals
 - Beware: May lead to dogpiles! We herd the calls together in lumps
 - Gaussian retry intervals (?)
 - ‘Exponential backoff’
 - Retry after 1s, 2s, 4s, 8s, 16s, and then give up.

Use Timeouts

- Example

- Exponential Back-off

- Wait 1 second, then 2, then 4, then 8, then 16...

```
    } catch ( ShutdownSignalException sse ) {
        // Happens when rabbitmq shut downs more or less gracefully
        incrementRetryCountAndWaitBeforeProceeding("MQ shutdown signal");
    } catch ( ConnectException connectException ) {
        // Happens in case we cannot connect to ANY of the MQs
        // The origin is the openChannelAndSetRouting method.
        incrementRetryCountAndWaitBeforeProceeding("MQ connection exception");
    } catch ( Exception otherExc ) {
        retryCount++;
        String theTrace = ExceptionUtils.getStackTrace(otherExc);
        log.error(theTrace);
    }
}

private void incrementRetryCountAndWaitBeforeProceeding(String exceptionDescription) {
    retryCount++;
    long delay = this.calculateBackoffDelayInMs();
    log.info( exceptionDescription + ", will backoff for "+
              delay+"ms and retry (Retry #"+retryCount+"')");
    try {
        // wait a bit to if things get better
        Thread.sleep( calculateBackoffDelayInMs() );
    } catch ( InterruptedException interExc ) {
        String theTrace = ExceptionUtils.getStackTrace(interExc);
        log.error(theTrace);
    }
}
```

Who does Retry?

- Exercise:
 - PlayerServant delegate to CaveStorage delegate to MongoDB
 - MongoDB's primary fails and throws exception...
 - We will talk about passive replication technique shortly
 - Exception caught in CaveStorage but
 - Handle it locally in CaveStorage and do a retry?
 - Rethrow as 'ElectionException' and handle in PlayerServant?
 - Pro and Con of each solution?

- In ‘design for failure’ in F2020 I herded students into trying to find graceful degradations
 - Like retries
- General impression
 - Seemed that the code became cumbersome and the results for the users not quite intuitive
 - “dig n My new room”
 - “Could not dig your room, will retry later”
 - ???
- So this year
 - *Fail fast and report it*
 - “Could not dig your room, you have to try again later”

- Phew...
- Plus
 - Shed load
 - Back pressure
 - Govenor
 - Bounded Result sets
 - Retry

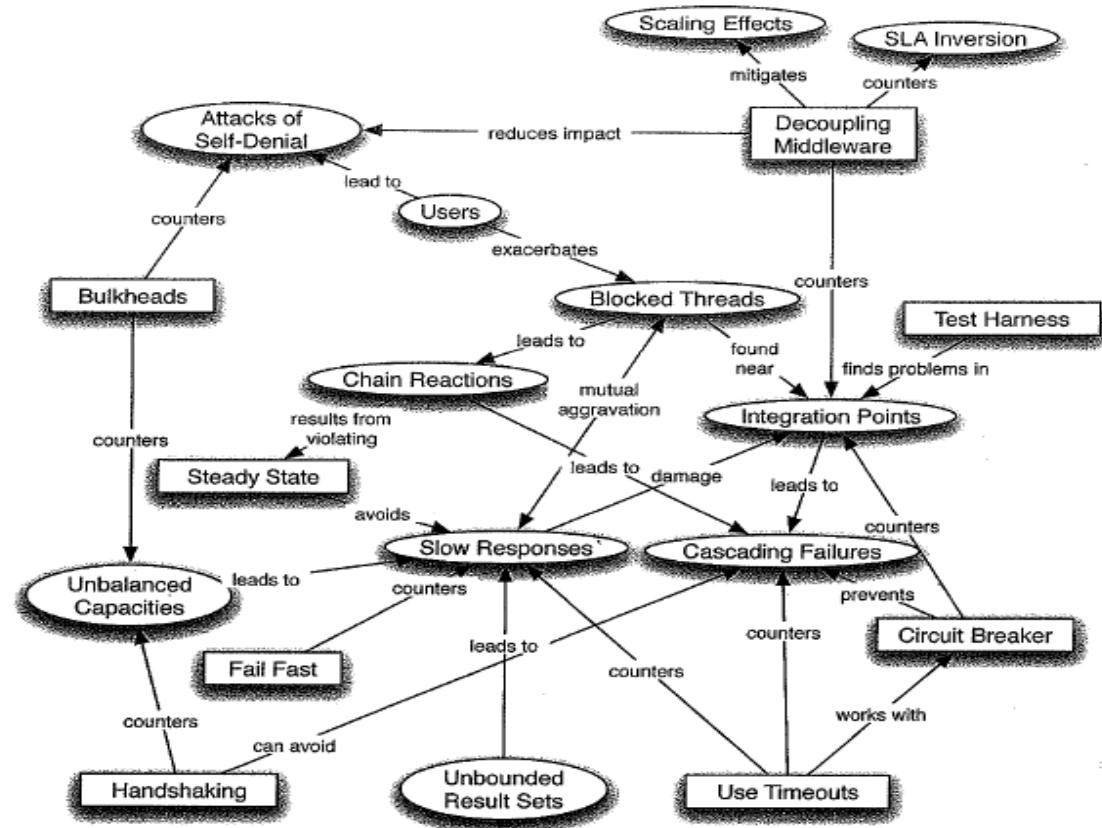


Figure 3.1: INTERACTION OF PATTERNS AND ANTI PATTERNS
