


Software Dependability:





A case study on Software Defined Networks

Carmen Mas Machuca, Petra Vizarreta
Chair of Communication Networks,
Technical University of Munich, Germany



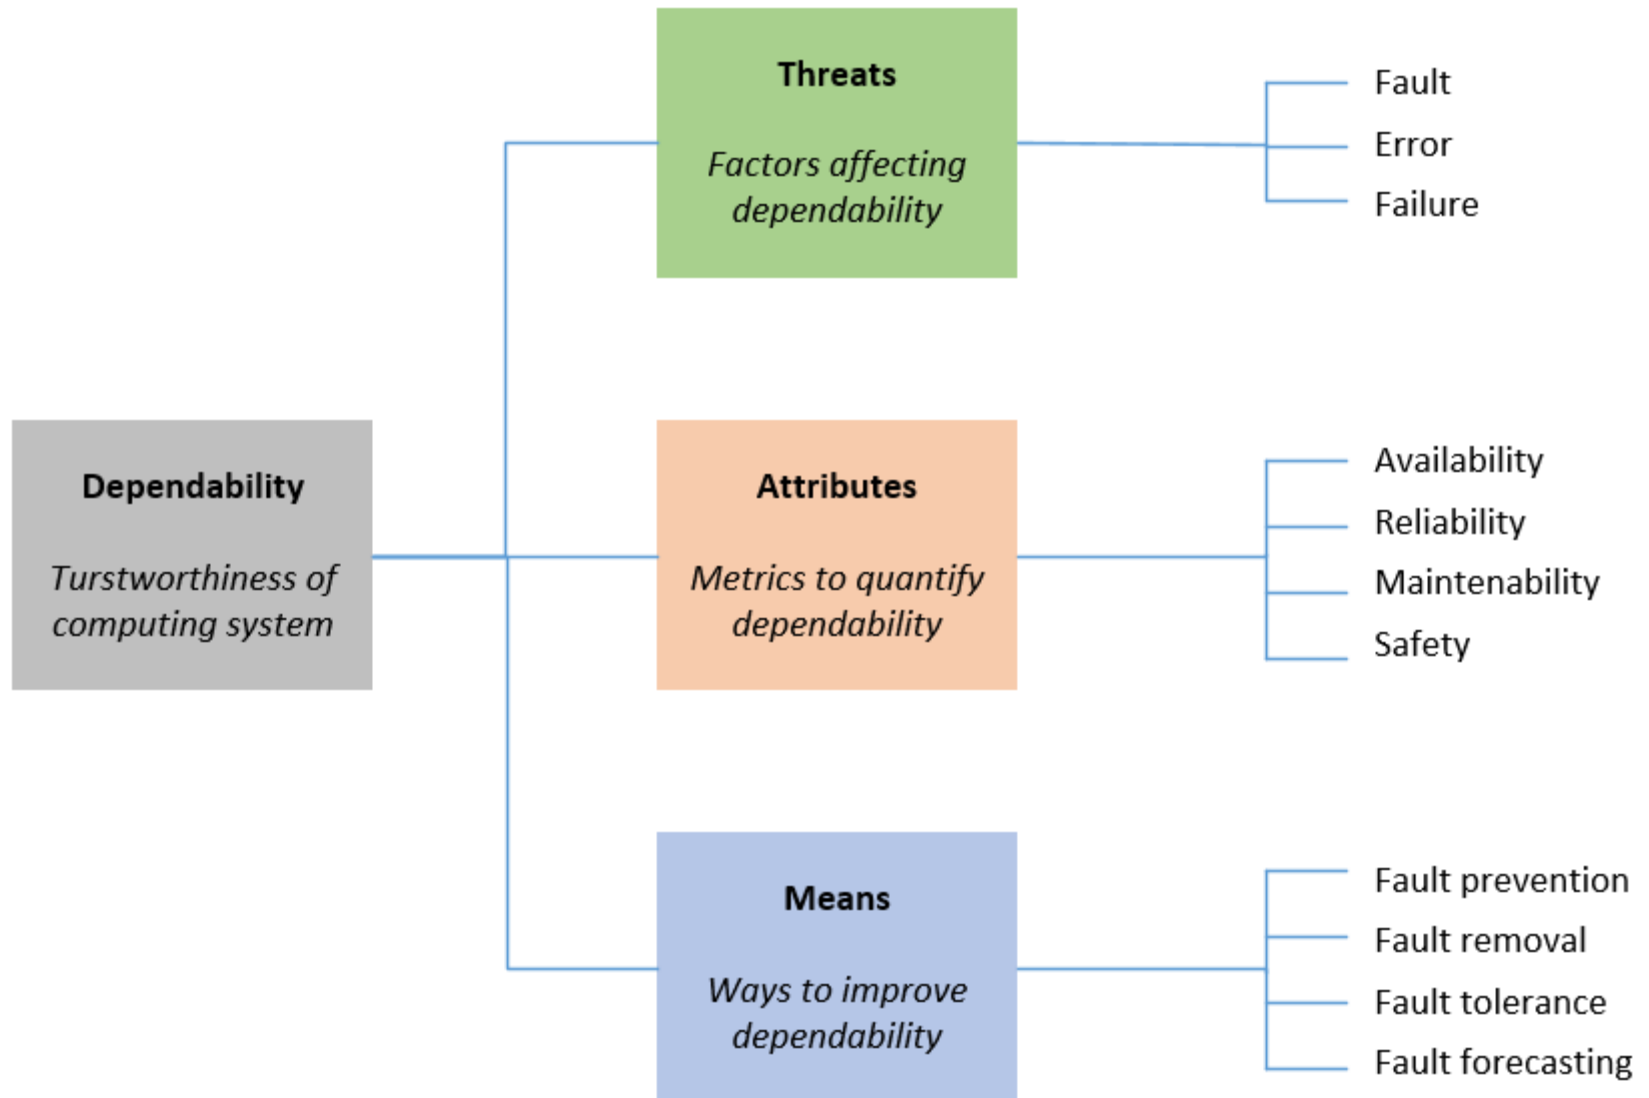
Ubiquity and magnitude of software failures

- Software bugs contribute more than 35% of critical network outages [Google2016]
- Bugs caused more than 33% of customer impacting incidents [Microsoft2017]


 <p>28.02.2017 09:37 (PST) S3 Service Disruption the Northern Virginia (US-EAST-1) Region</p> <p>https://aws.amazon.com/message/41926/</p>	 <p>22.06.19 03:00 to 22.06.20 A widespread BGP routing Internet services and a p</p> <p>https://bgr.cc 2019-google-amazon-rec</p>	 <p>26.03.20 16:14 to 27.03. Cloud IAM experienced across many services (resulting in continued subset of services) for</p> <p>https://sta</p>	 <p>19.05.20 13:30 to 16:30 (UTC). A bug caused high resource utilization in the internal cluster service that is responsible for receiving and executing service management operations in the East US region. The bug was encountered in all the service instances of the region leading to failures and timeouts for management operations</p> <p>https://status.azure.com/en-us/status/history</p>
---	--	---	--

PST: Pacific Standard Time
CET: Central European Time

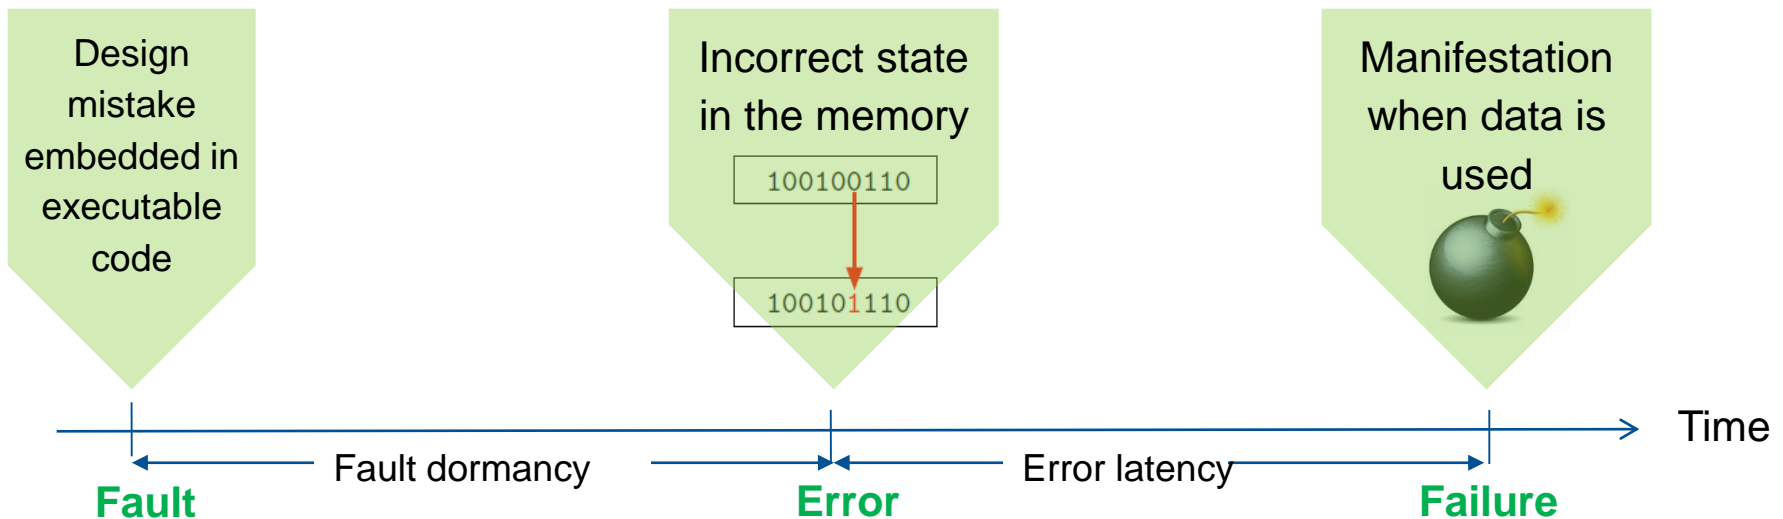
- Terms and Taxonomy
- Software Dependability Problem
- Addressed questions applied to SDN:
 - How reliable a controller is? → Steady-state availability
 - How often does software fail? → Bug forecasting and Software Maturity evaluation
 - What is the impact? → User-perceived service
- Conclusions



Source: IFIP WG10.4 Dependable Computing and Fault Tolerance <https://www.dependability.org/wg10.4/>




- **Fault:** Adjudged or hypothesized cause of an error.
- **Error:** Part of a system state which is liable to lead to failure.
- **Failure:** Deviation of the delivered service according to its specification.



- Active: it produces an error
- Dormant: it has not produced an error
- Detected: it has manifested as failure
- Latent: it has not been detected



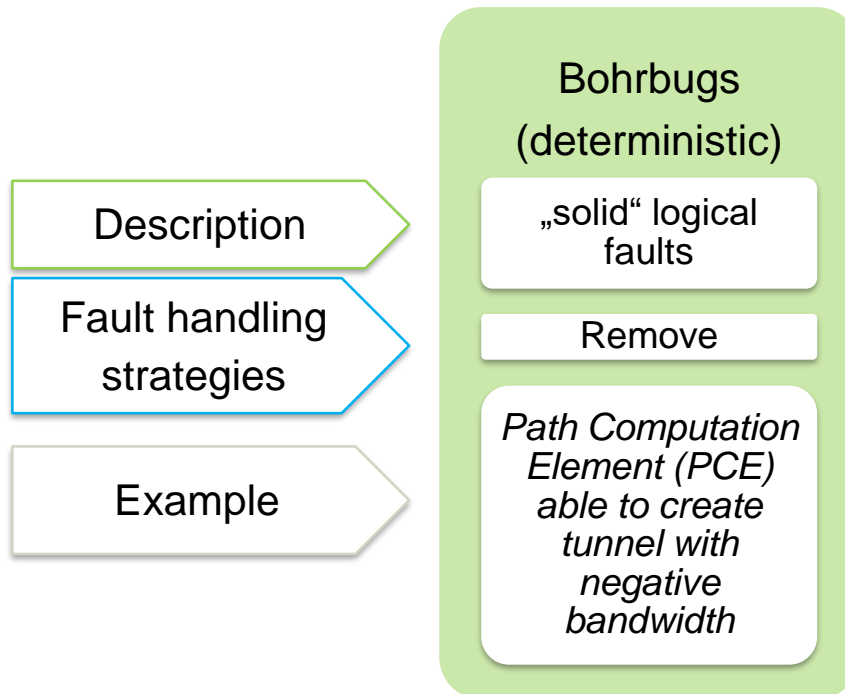
- **Availability:** The ability of an item to perform its required function, under environmental and operational conditions at a stated instant of time.
- **Reliability:** The ability of an item to perform its required function, under environmental and operational conditions, for a stated period of time.
- **Maintenability:** the probability of performing a successful repair and maintenance action within a given time.
- **Safety:** Ability of an item to provide its required function without the occurrence of catastrophic consequences on the user(s) and the environment.



- **Fault prevention** is attained by quality control techniques employed during the design and manufacturing of hardware and software.
- **Fault removal** is performed both during the development phase (verification, diagnosis, and correction), and during the operational life of a system (either corrective or preventive maintenance).
- **Fault tolerance** is intended to preserve the delivery of correct service in the presence of active faults.
- **Fault forecasting** is conducted by performing an evaluation of the system behaviour with respect to fault occurrence or activation: either qualitative (identify, classify, rank the failure modes), or quantitative (probabilities to which some of the attributes are satisfied).

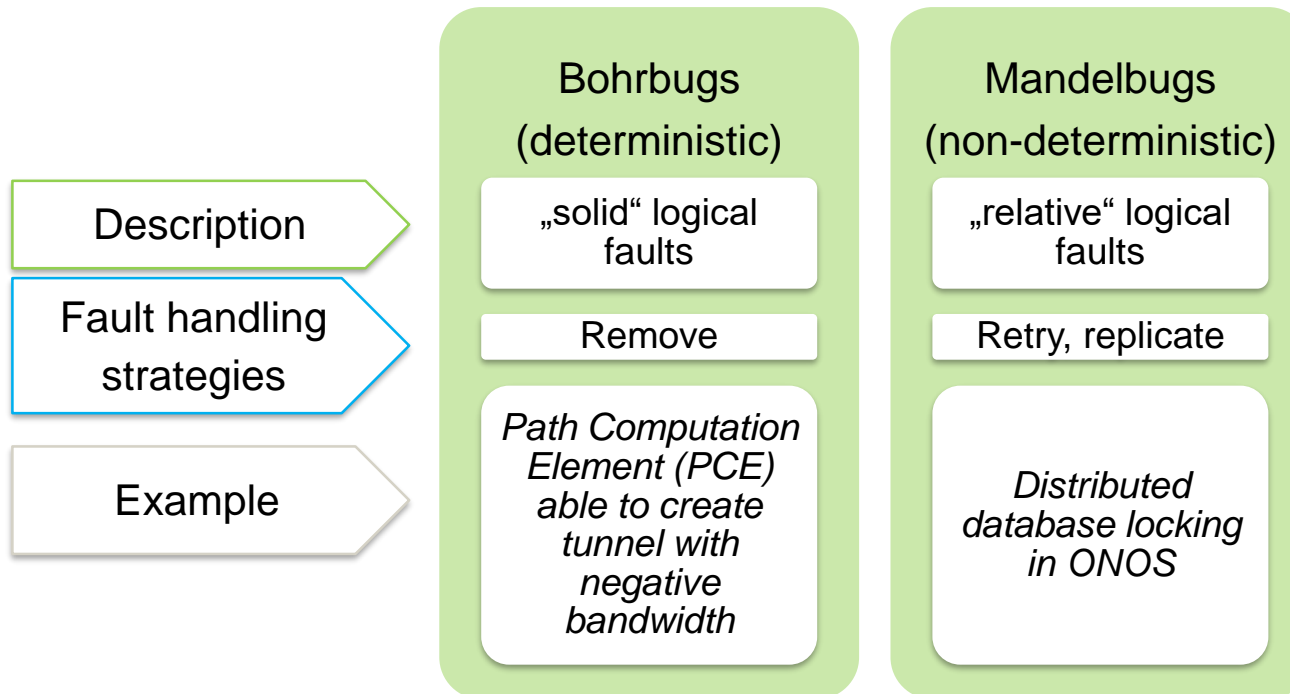


- Software fault = bug
- Types of software faults:



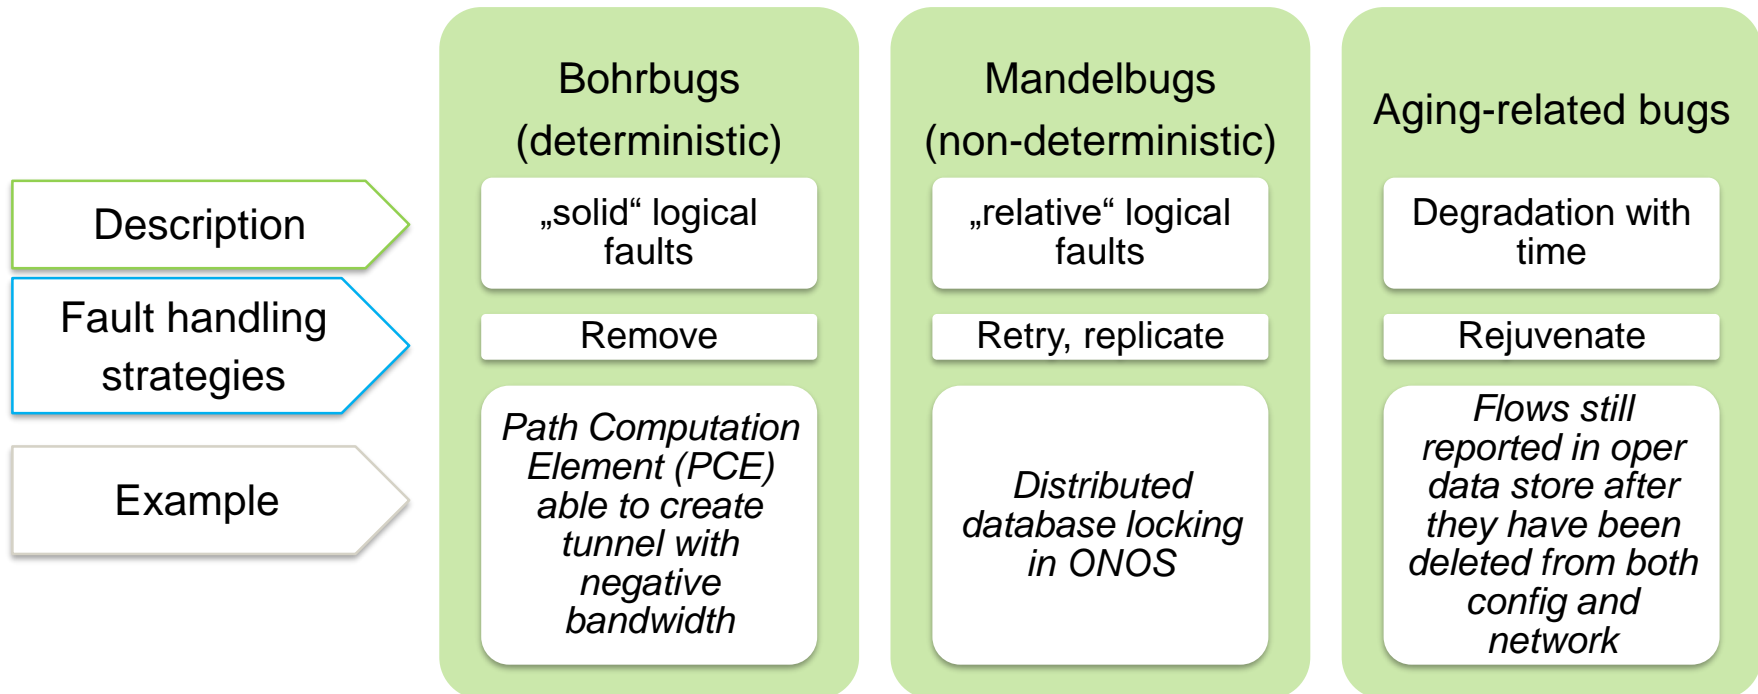


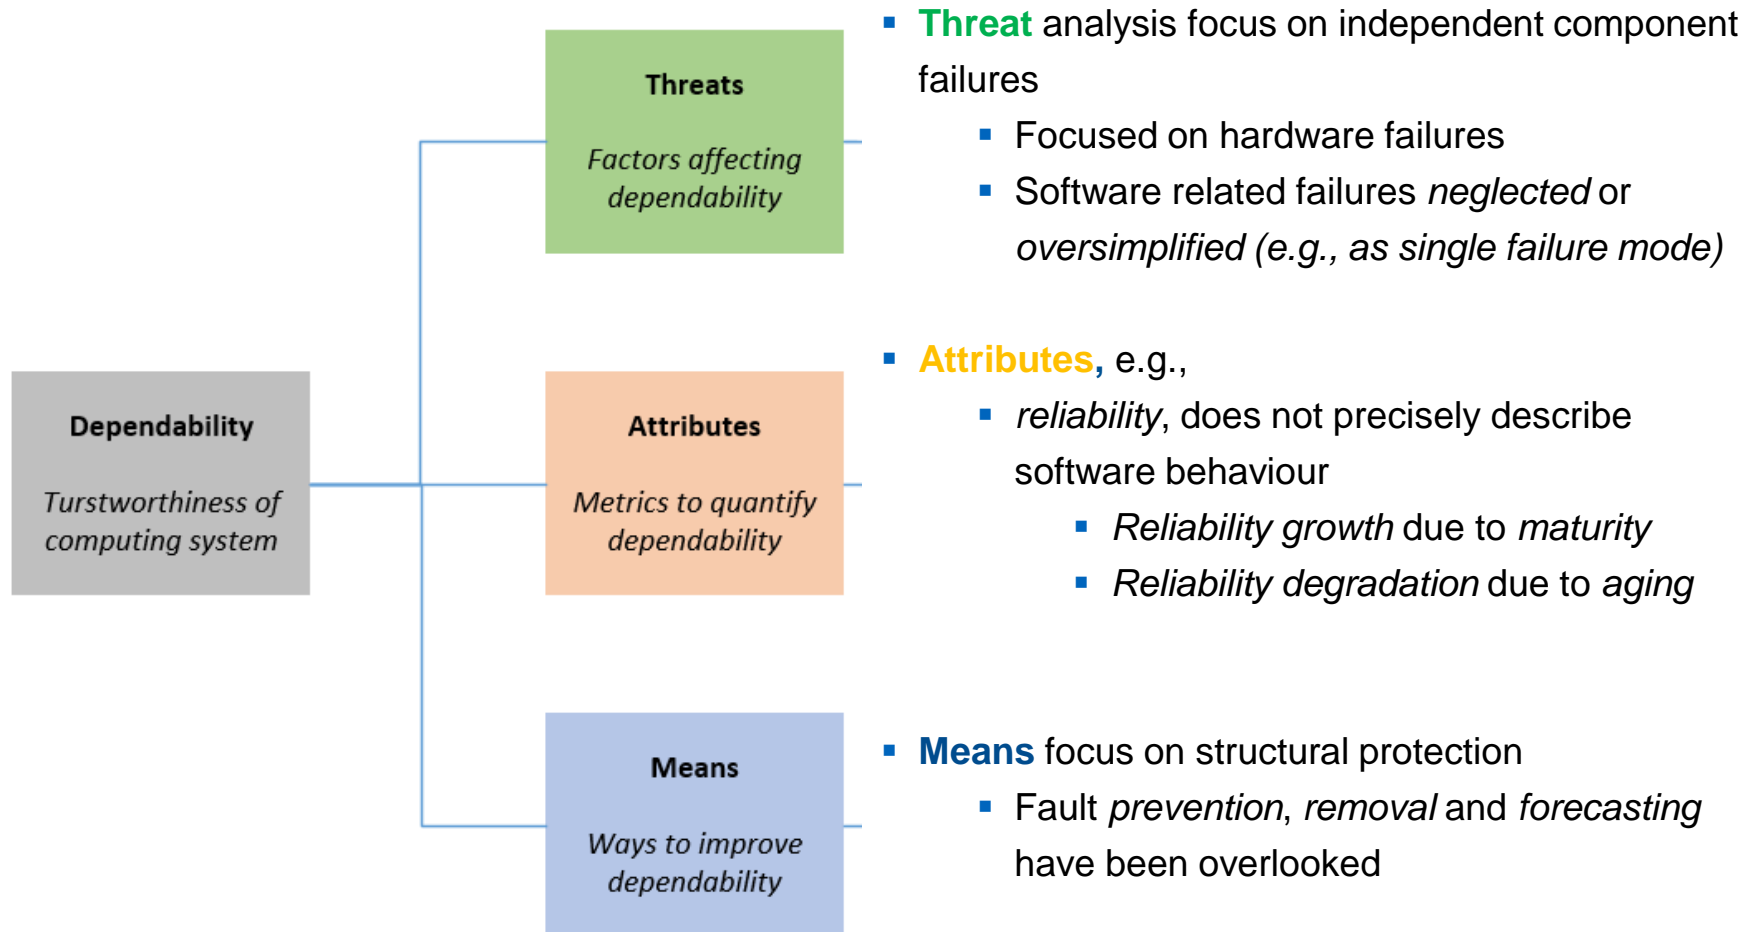
- Software fault = bug
- Types of software faults:





- Software fault = bug
- Types of software faults:





Software Dependability Problem

- Softwarized networks
- Open source code

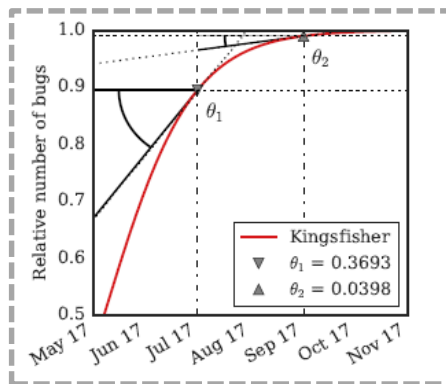


Target: Realistic and practical dependability assurance framework

Proposed methodology based on Statistical inference techniques and stochastic dependability models

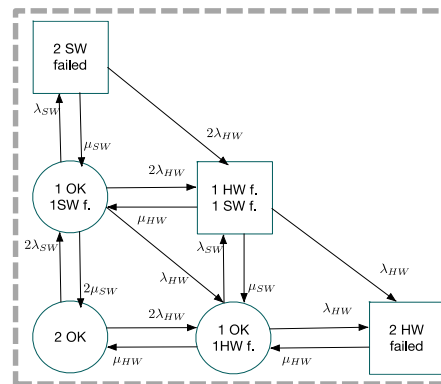
How often does SW fail?

*Failure forecasting and
Software Maturity*



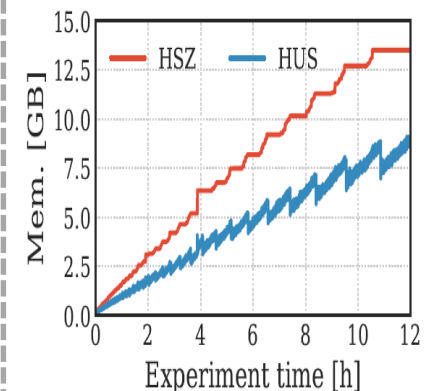
How often is the controller available?

Steady-state availability



Do Softwarized networks age?

Proposed framework



Software Dependability Problem

- Softwarized networks
- Open source code




Target: Realistic and practical dependability assurance framework

Proposed methodology based on Statistical inference techniques and stochastic dependability models


How often does SW fail?

*Failure forecasting and
Software Maturity*



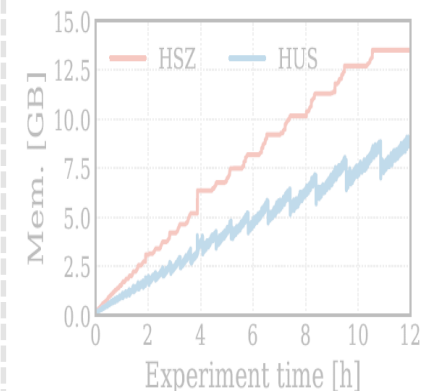
How often is the controller available?

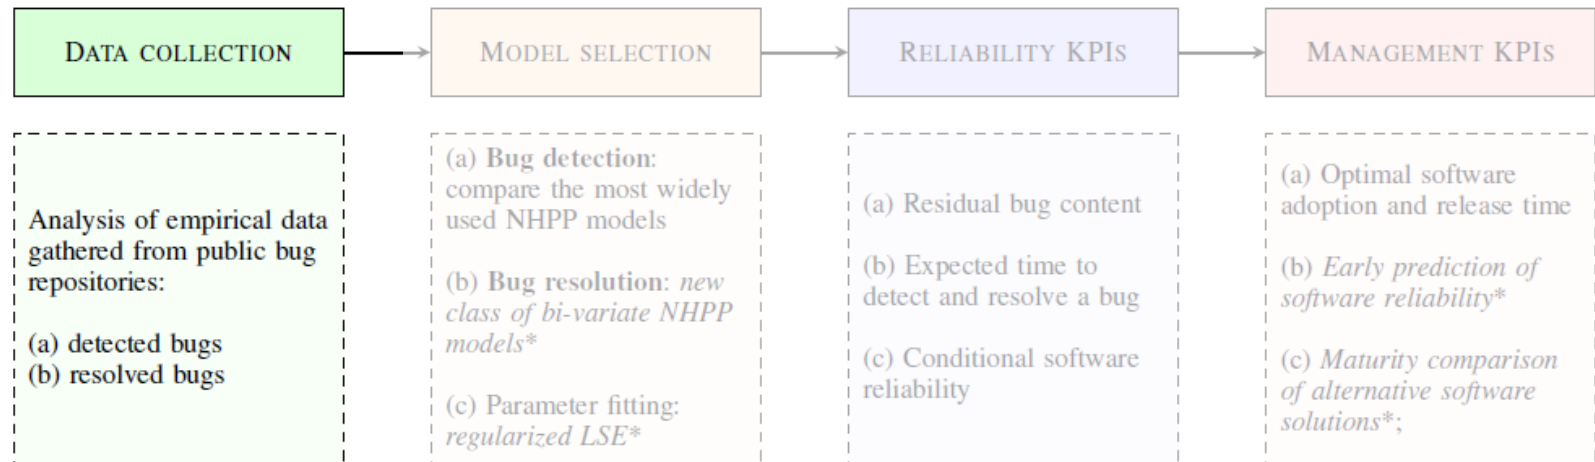
Steady-state availability



Do Softwarized networks age?

Proposed framework





Failure Forecasting and Software Maturity

OpenDaylight Bugzilla - Main P
Home | New | Browse | Search |
Forgot Password

OpenDaylight has migrated to JIRA

Welcome to

File a Bug

Enter a bug # or some
Quick Search
OpenDaylight

Bug 9218 - Oxygen development is limited by SingleFeatureTest running out of heap space

Status: IN_PROGRESS

Product: odparent
Component: General
Version: Nitrogen
Hardware: All All

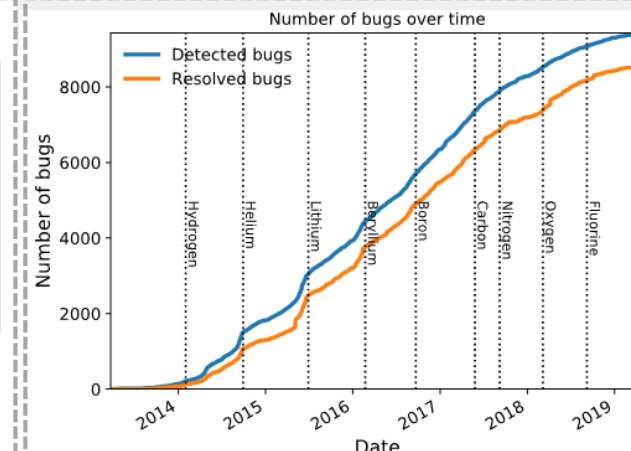


Importance: --- blocker
Assigned To: Robert Varga

URL:
Whiteboard:


Depends on:
Blocks:

Reported: 2017-09-27 08:45 UTC by
Modified: 2017-10-09 10:45 UTC (H
CC List: 1 user ([show](#))

See Also:
ODL SR Target Milestone: Oxygen
Issue Type: Bug
External References:



Bug detection as Non-Homogeneous Poisson Process (NHPP)




- Initial number of bugs N is Poisson random variable with $E[N] = a$
- Probability of detecting a single bug (manifested SW fault) by time t
- Assuming time to discover every bug is i.i.d. we have Bernoulli trials

$$P(N = n) = \frac{a^n}{n!} e^{-a}$$

$$p = F(t)$$


- The cumulative number of detected bugs
- Expected number of detected bugs by time t

$$P(N(t) = k) = \frac{[aF(t)]^k}{k!} e^{-aF(t)}$$



Software Reliability Growth Models: *Model selection*


Bug detection as Non-Homogeneous Poisson Process (NHPP)



The eight most widely used NHPP models for modelling of the bug detection process are:

Model	Shape	Mean value function
Musa-Okumoto logarithmic	Concave	$m_{mo}(t) = a \ln(1 + bt)$
Goel-Okumoto exponential	Concave	$m_{go}(t) = a(1 - e^{-bt})$
Generalized Goel-Okumoto	S-shaped	$m_{ggo}(t) = a(1 - e^{-bt^c})$
Ohba's inflection S-shaped	S-shaped	$m_{iss}(t) = a \frac{1 - e^{-bt}}{1 + \phi e^{-bt}}$
Yamada delayed S-shaped	S-shaped	$m_{dss}(t) = a(1 - (1 + bt)e^{-bt})$
Yamada exponential	Concave	$m_{yex}(t) = a(1 - e^{-r(1 - e^{-bt})})$
Gompertz	S-shaped	$m_{gomp}(t) = ak^{bt}$
Logistic	S-shaped	$m_{logist}(t) = \frac{a}{1 + ke^{-bt}}$

Commonly used Non-Homogeneous Poisson Process (NHPP) [Lyu95]



Bug resolution (R) is a combination of two processes: bug detection (D) and bug correction (C)

$$f_R(t) = \int_0^t f_D(t-x)f_C(x)dx = [f_D * f_C](t)$$

$$m_R(t) = aF_R(t) = a \int_0^t [f_D * f_C](x)dx$$

- Closed form solution exist only in trivial cases

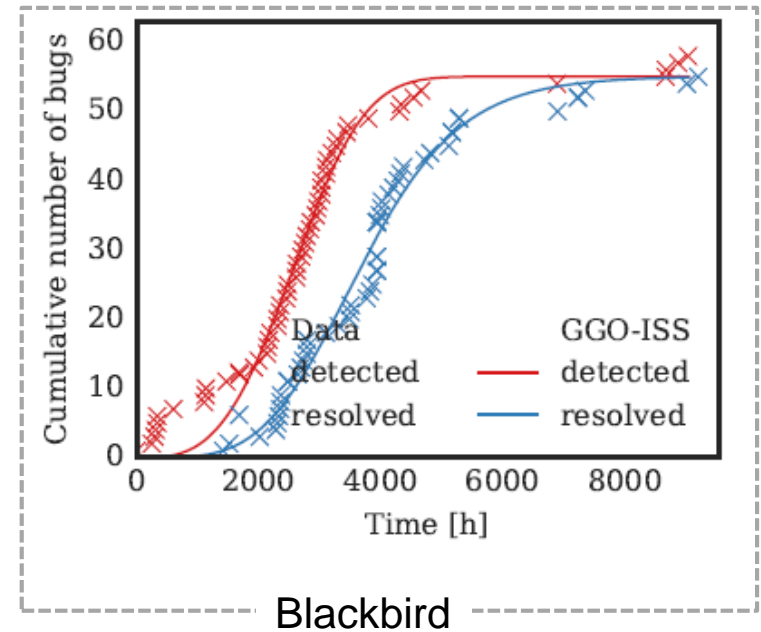
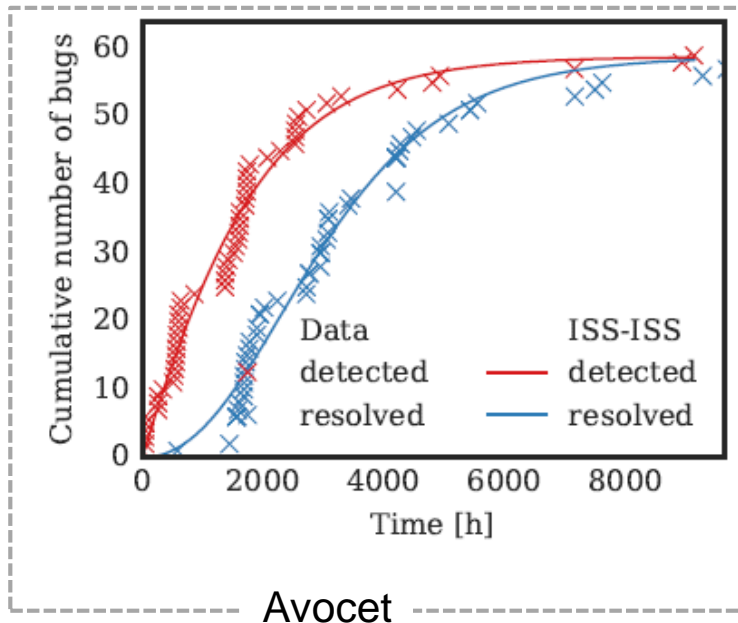
$$m_R^{go-go}(t) = a \left[1 - \frac{b_1 e^{-b_2 t} - b_2 e^{-b_1 t}}{b_1 - b_2} \right]$$

- PCA: Piecewise Constant Approximation is used for fitting instead


$$\tilde{F}_R(t) = \sum_{i=0}^{n=t/\Delta x} [f_D * f_C](i\Delta x)\Delta x$$

$$F_R(t) = \lim_{\Delta x \rightarrow 0} \tilde{F}_R(t)$$



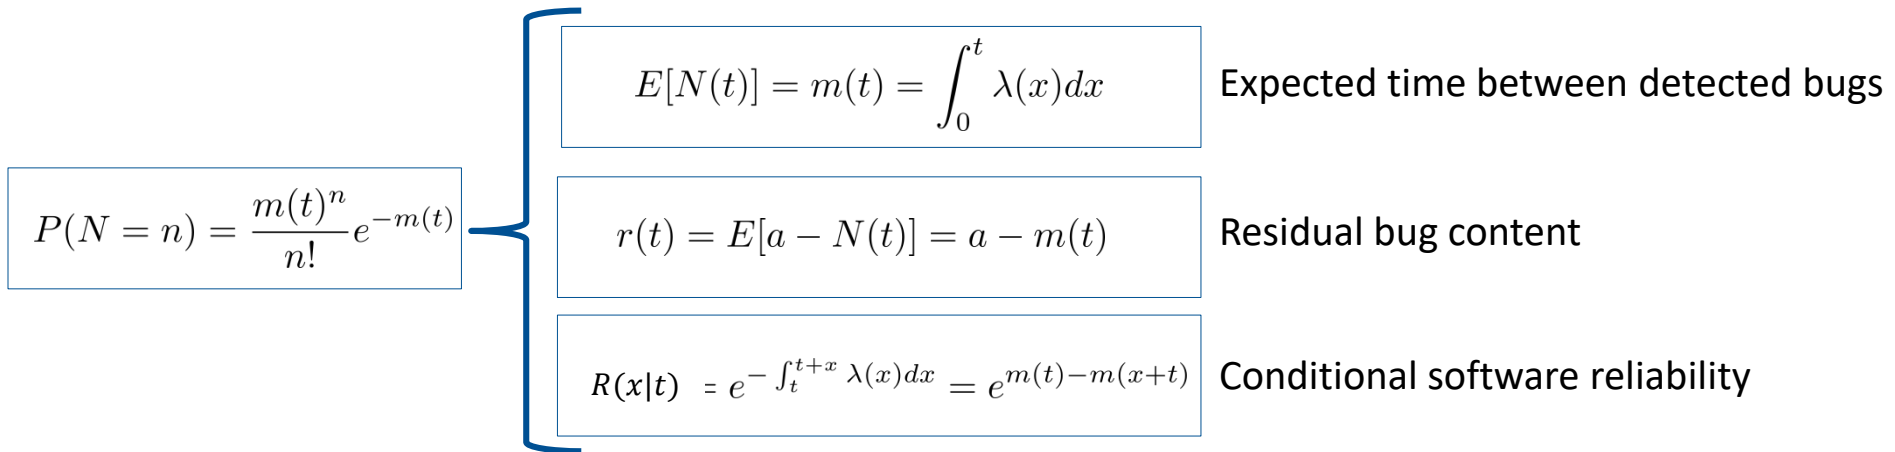


The best fitting models for detected and resolved bugs may be different.



Bug detection

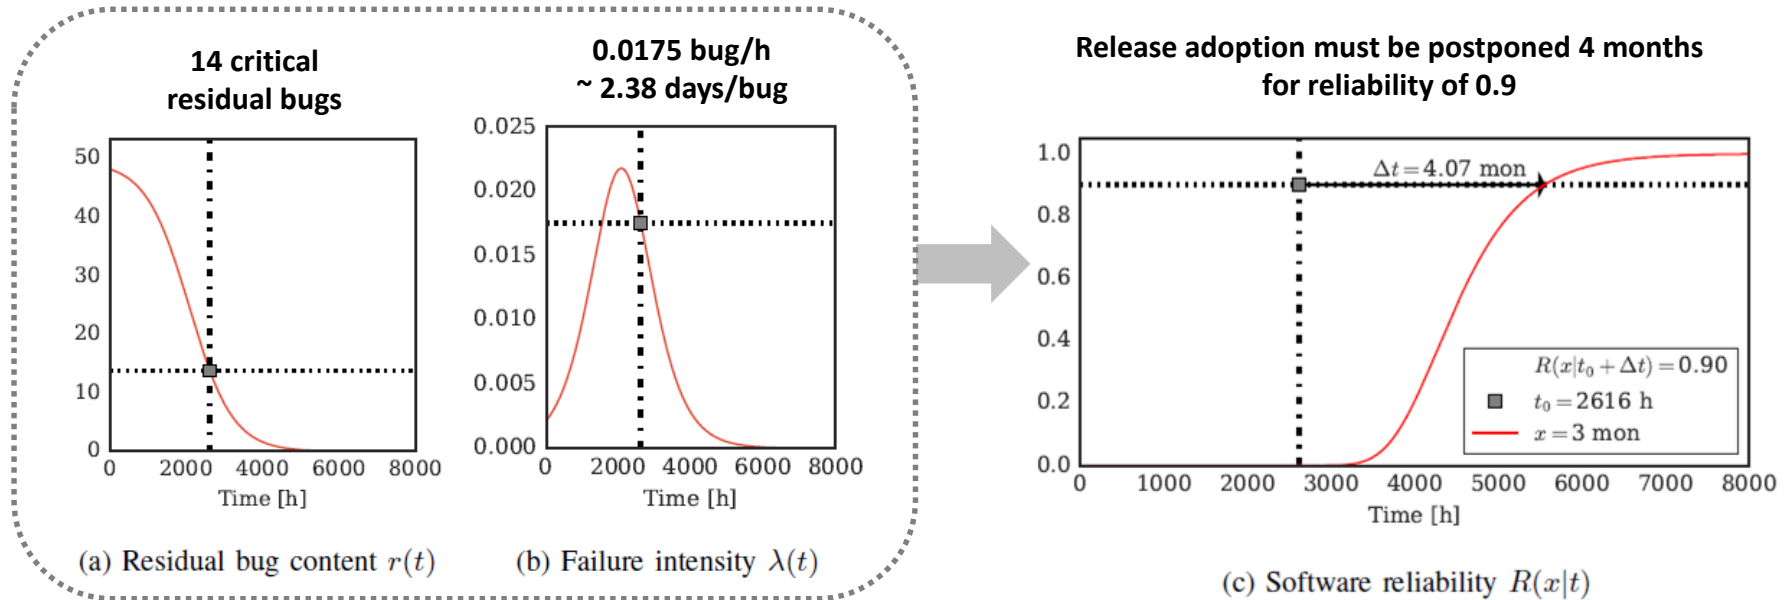
NHPP model is completely described by its mean value function $m(t)$



Similarly for **Bug resolution**



Based on the selected model



ONOS Junco release: 28.02.2017




3 months=2160 hours

ONOS Kingsfisher release: 31.05.2017



Software Maturity Metric

- defined as the scaled gradient of the cumulative number of bugs, i.e., $\frac{\lambda(t)}{m_{max}}$.
- measures how far is the software from the stable region at any given moment.




ONOS Kingsfisher final release (FR): June 2017

ONOS Loon release: September 2017




How often does SW fail?

*Failure forecasting and
Software Maturity*




How often is the controller available?

Steady-state availability



Do Softwarized networks age?

Proposed framework



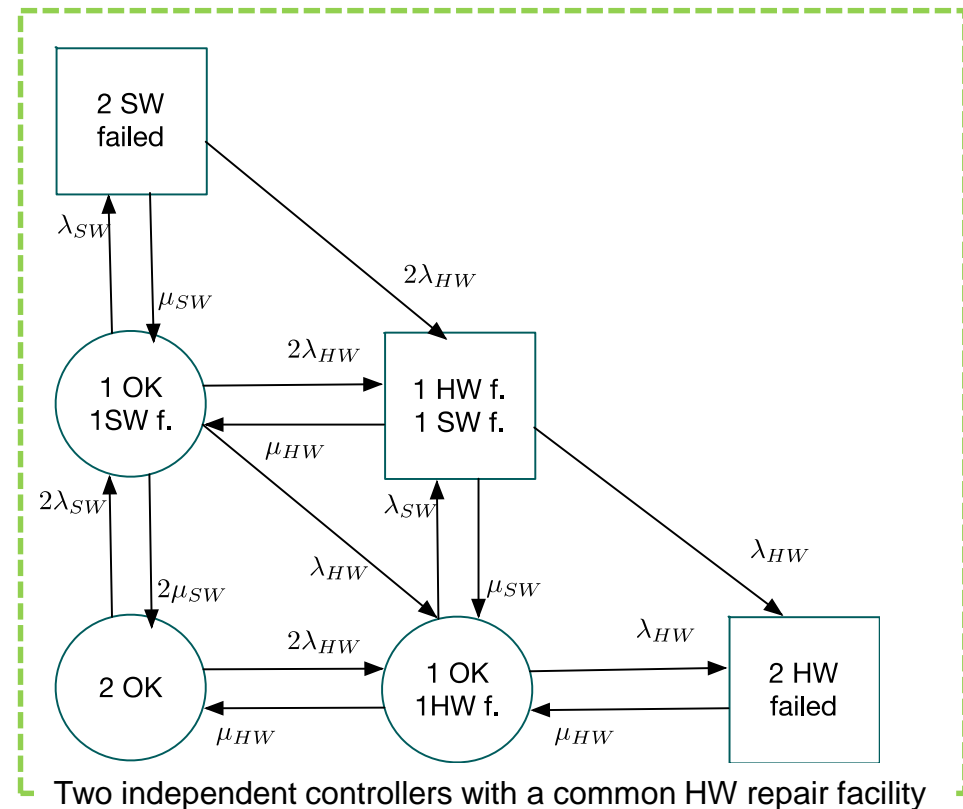
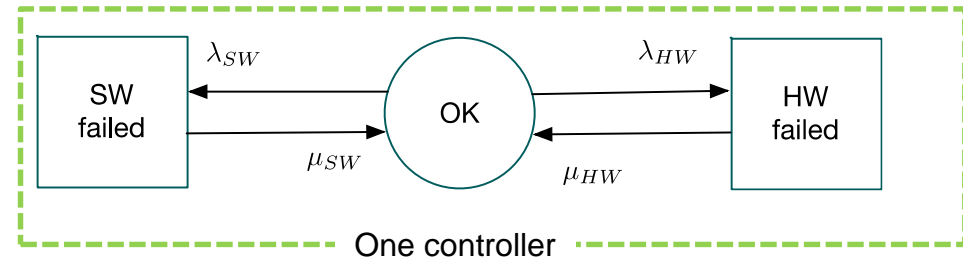
Steady State Availability

Homogeneous Markov Chains

- Single failure modes
- Usual assumptions
 - $\lambda_{HW} < \lambda_{SW}$
 - $\mu_{HW} > \mu_{SW}$
 - $\mu_x \gg \lambda_x$
- Failure shadows

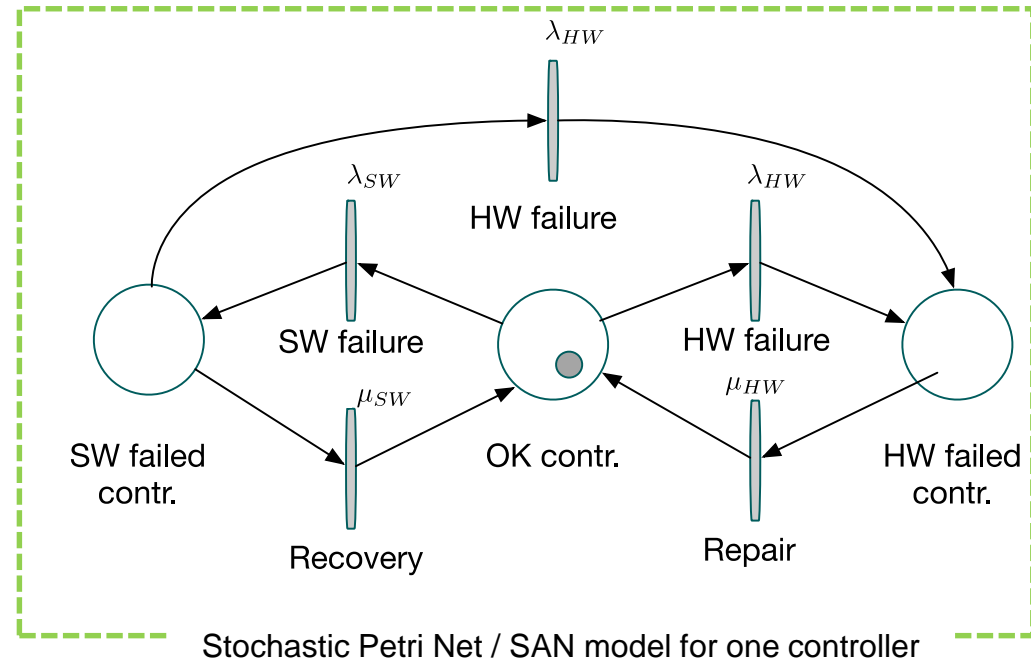
$$\left. \begin{aligned} \sum_{\forall i} \hat{p}_i &= 1 \\ \hat{P} &= \hat{P} \cdot \mathcal{T} \end{aligned} \right\}$$

$$A = \sum_{i \in \Omega_W} \hat{p}_i$$



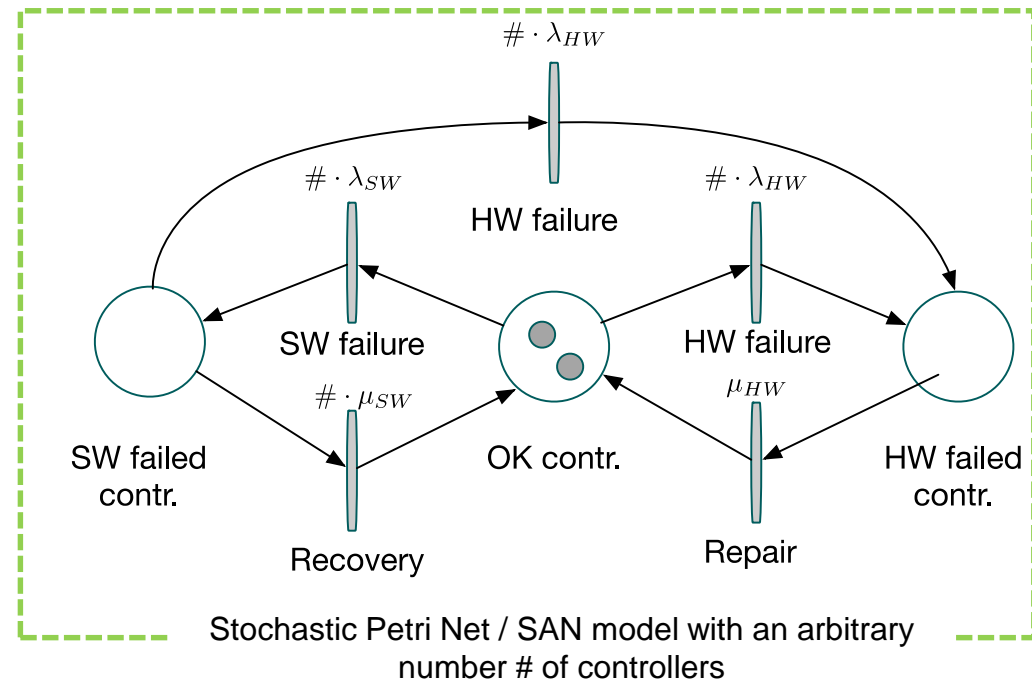
Stochastic Petri Nets/ Stochastic Activity Networks (SANs)

- Single failure modes
- Usual assumptions
 - $\lambda_{HW} < \lambda_{SW}$
 - $\mu_{HW} > \mu_{SW}$

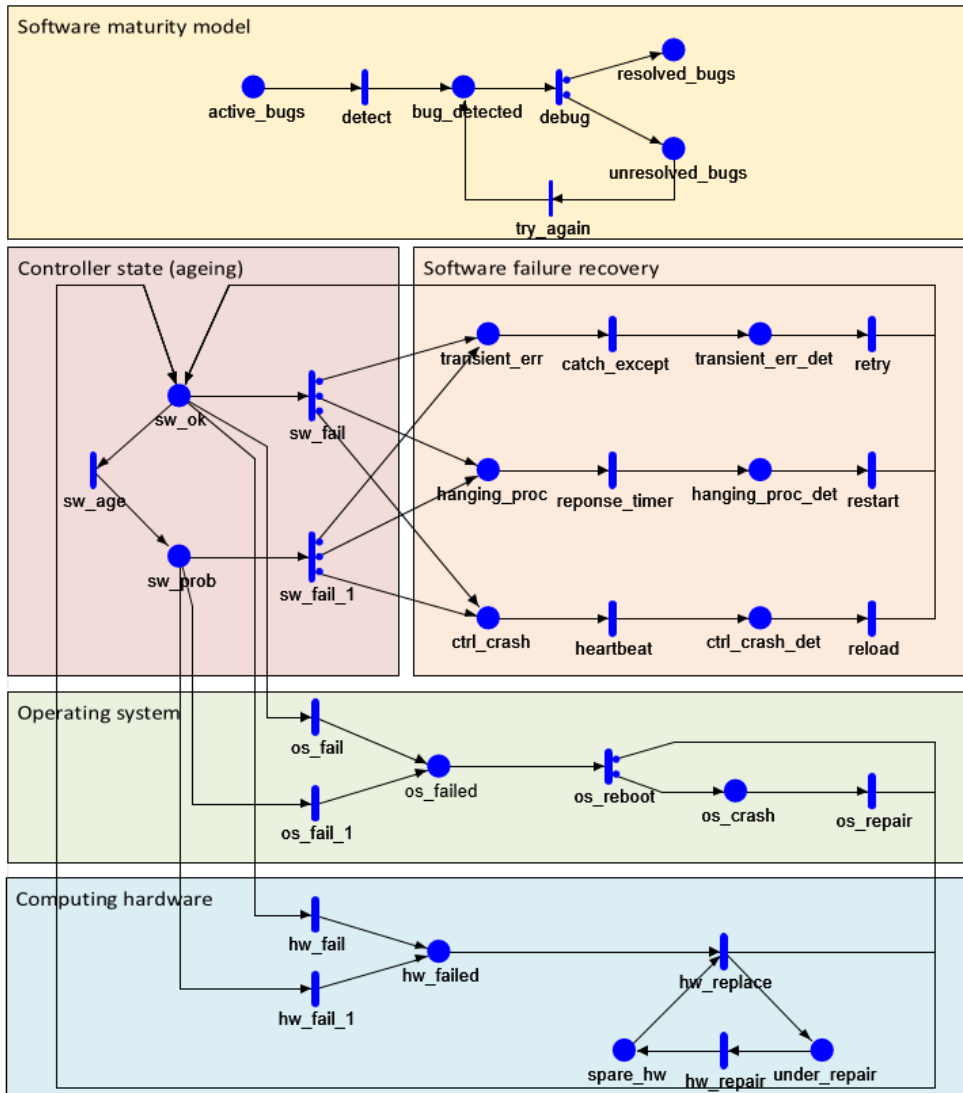


Stochastic Petri Nets/ Stochastic Activity Networks (SANs)

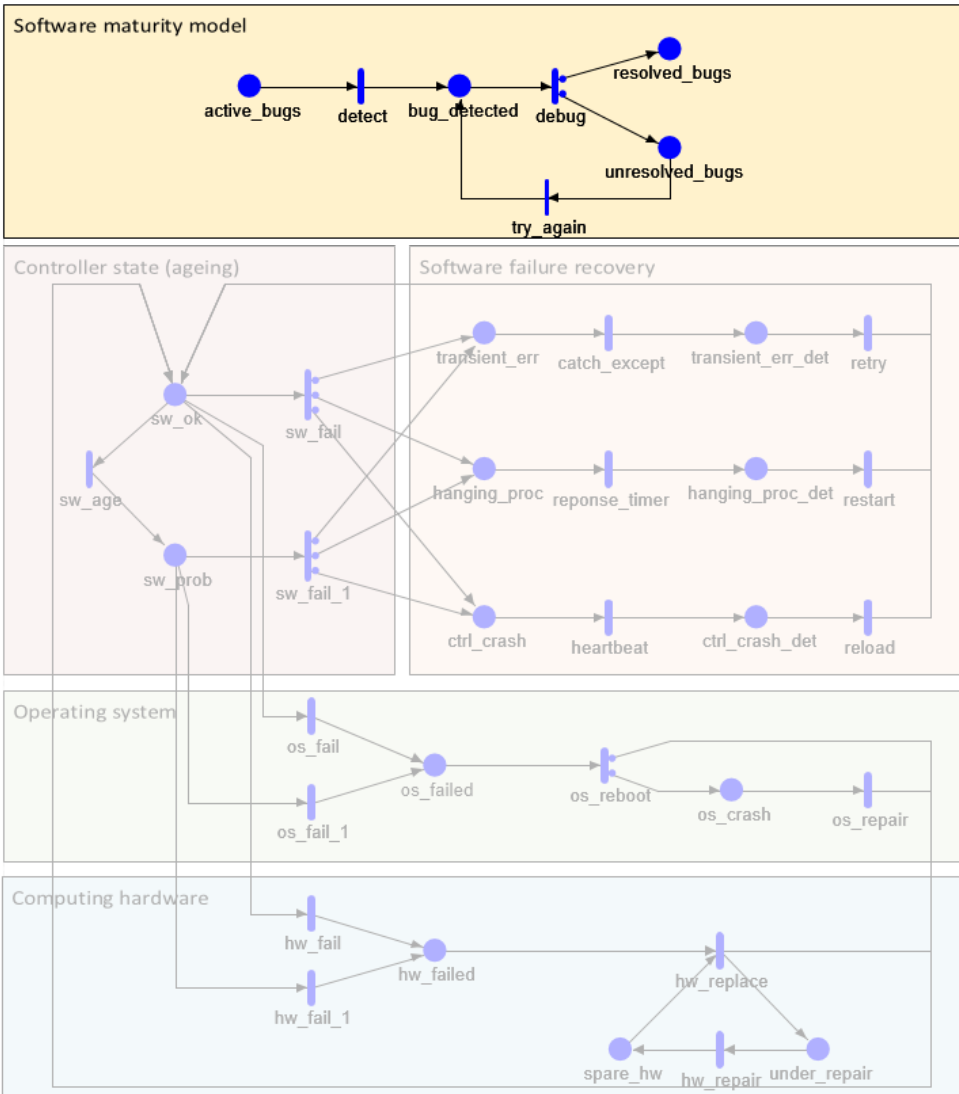
- Single failure modes
- Usual assumptions
 - $\lambda_{HW} < \lambda_{SW}$
 - $\mu_{HW} > \mu_{SW}$



SRN: A Case Study on SDN Controllers



SRN: A Case Study on SDN Controllers




1. Software reliability growth

long term variations of software reliability

Software maturity model

Parameter	Description	Baseline value
N_{bugs}	Initial number of active bugs	60
p_{debug}	Debugging success rate	0.99
$\lambda_{bug_detect}^{-1}$	Bug detection rate	60 days
μ_{debug}^{-1}	Debug rate	60 days

- Model: Jelinski-Moranda with imperfect debugging



SRN: A Case Study on SDN Controllers

2. Software aging


short term variations of software reliability

Parameter	Description	Baseline value
$\varphi_{sw_fail}^{-1}$	Baseline software failure rate	7 days
$\lambda_{sw_age}^{-1}$	Rate of software ageing	1 day
$\lambda_{age_fail}^{-1}$	Ageing failure rate	7 days
$p_{retry}(ok/prob)$	Failures recovered by retry	0.15, 0.15
$p_{restart}(ok/prob)$	Failures recovered by restart	0.15, 0.70
$p_{reload}(ok/prob)$	Failures requiring reload	0.70, 0.15

Failure frequency rate depends on controller state:

- highly robust state **sw_ok**
- vulnerable state **sw_prob**

Aging rate
 λ_{sw_age}



SRN: A Case Study on SDN Controllers

3. Nature of failures

Parameter	Description	Baseline value
μ_{catch}^{-1}	Catch the exception	1 msec
$\mu_{timeout}^{-1}$	Detect hanging process	1 sec
$\mu_{heartbeat}^{-1}$	Detecting controller crash	10 sec
μ_{retry}^{-1}	Retry the operation	0.5 sec
$\mu_{proc.restart}^{-1}$	Process restart	5 min
μ_{reload}^{-1}	Restart controller and reload	30 min

Transient failures

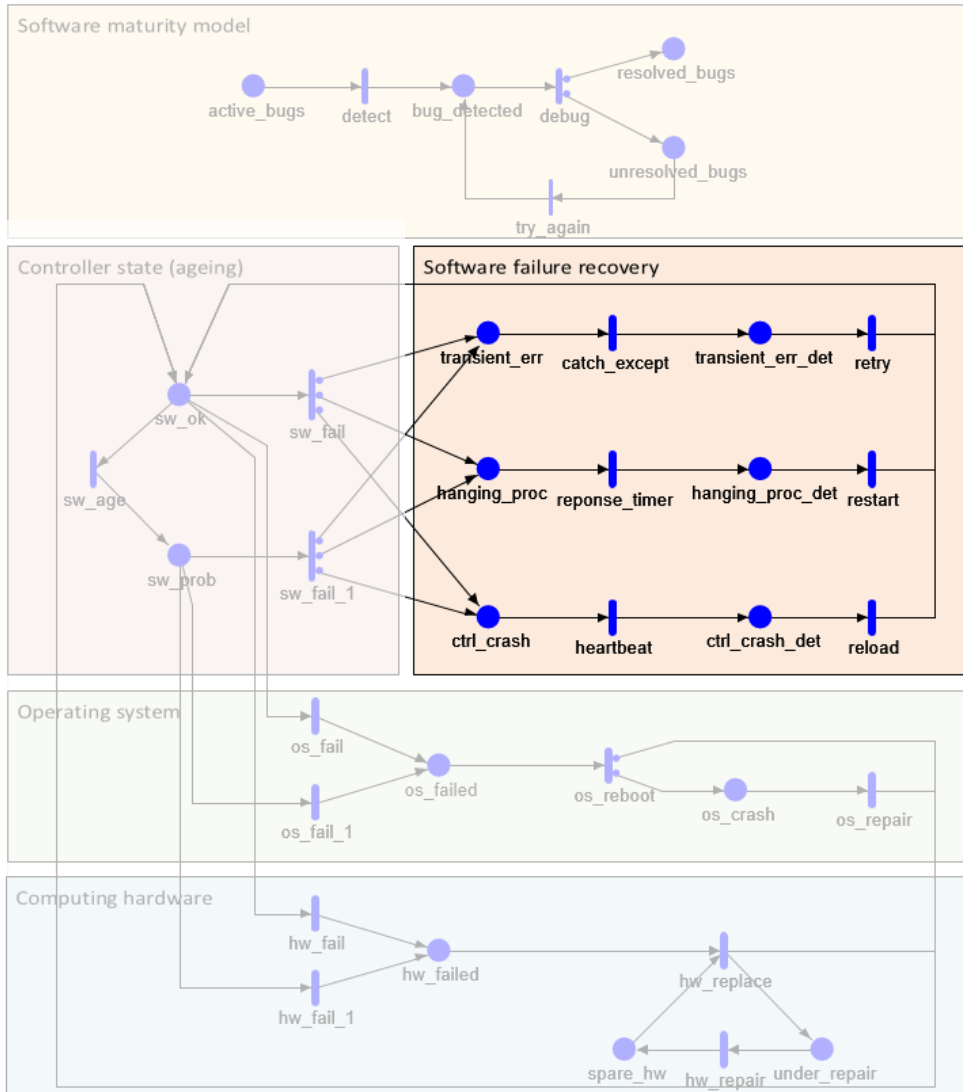
- detected by catch-except routine
- mitigated by retrying the operation

Hanging failures


- detected by response timers
- mitigated by bundle restart

Crash failures

- detected by heartbeat messages
- controller software reloaded from the last checkpointed (saved) state



SRN: A Case Study on SDN Controllers



4. Operating system

Failures of operating system (OS)

Parameter	Description	Baseline value
$\lambda_{os_fail}^{-1}$	Mean time between OS failures	60 days
pos_reboot	Success of OS reboot	0.9
$\mu_{os_reboot}^{-1}$	OS reboot time	10 min
$\mu_{os_repair}^{-1}$	OS repair time	1 h

5. General purpose Hardware

Failures of computing hardware (HW)

Parameter	Description	Baseline value
$\lambda_{hw_fail}^{-1}$	Mean time between HW failures	6 months
$\mu_{hw_replace}^{-1}$	HW replace time	2 hours
$\mu_{hw_repair}^{-1}$	HW repair time	24 hours
N_{spare_hw}	Spare computing hardware	1

Evaluation of SDN controller

Steady state availability

- At least two controllers are needed to achieve “3-nines” availability


Component	Controller	SW	OS	HW
Availability	0.99889	0.99956	0.99981	0.99951

[Ros14] assumed much higher availability of SDN controller

$$A > 0.999975$$

Further study on clustering:
imperfect failover and state
synchronisation

- Identification of the most critical parameters (local sensitivity analysis)




Critical parameters

- External failure rates
(well studied and documented)
- Software aging rate
(uncertain, load dependant)

A comprehensive study on
software aging

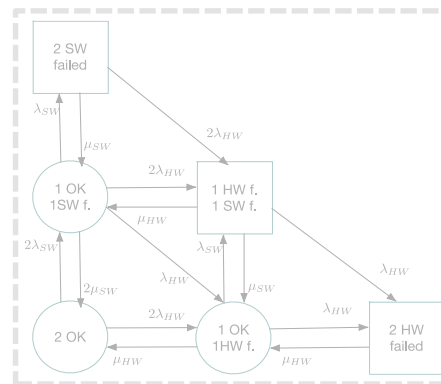
How often does SW fail?

*Failure forecasting and
Software Maturity*




How often is the controller available?

Steady-state availability



Do Softwarized networks age?

Proposed framework



Detection

SOFTWARE AGING DESCRIPTORS

PERFORMANCE ISSUES

Gradual response time increase and throughput regression

RESOURCE LEAKS

Resource leaks, e.g., unreleased file locks, thread leakage

MEMORY LEAKS

Memory leaks leading to out-of-memory crash during runtime

OTHER ISSUES


Other aging symptoms, e.g., data corruption, storage fragmentation

Not all are due to bugs (undesired behaviour of the code that should be corrected), but rather a deliberate design decision

E.g., In ONOS, when flow rules are added and removed, they are not deleted from the controller datastore; Instead, they are replaced with thumbstones (placeholders), to ensure stability of Gossip protocol. This also affects other eventually consistent network state primitives which rely on Gossip


Impact evaluation

Aging observed at **system** level:





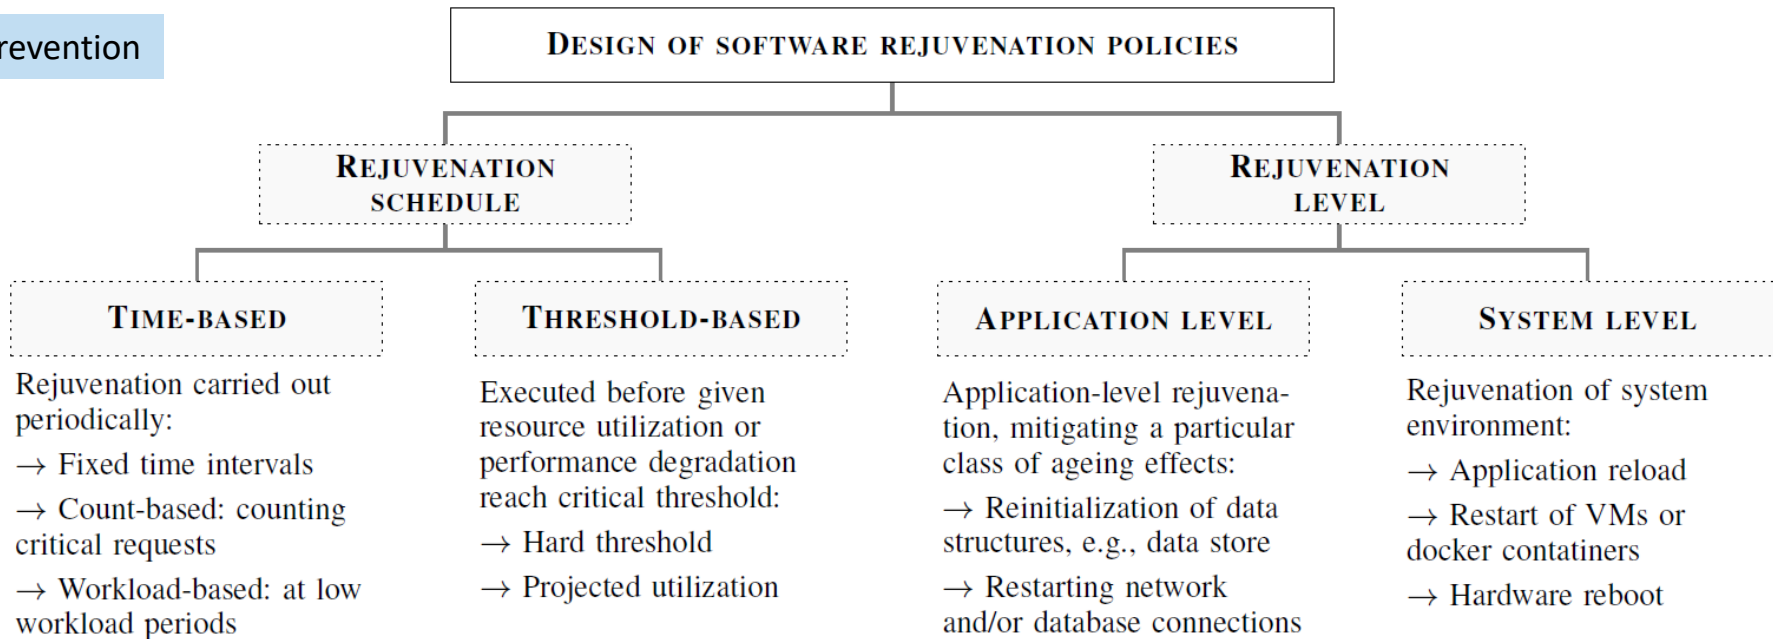
- Allocated heap (HSZ) and used heap memory (HUS) continuously grow
- System crashes after HSZ exhausts all 14 GB of available memory
- **Crash happens after 18h** at 300 intent/s

Aging observed at **application** level:

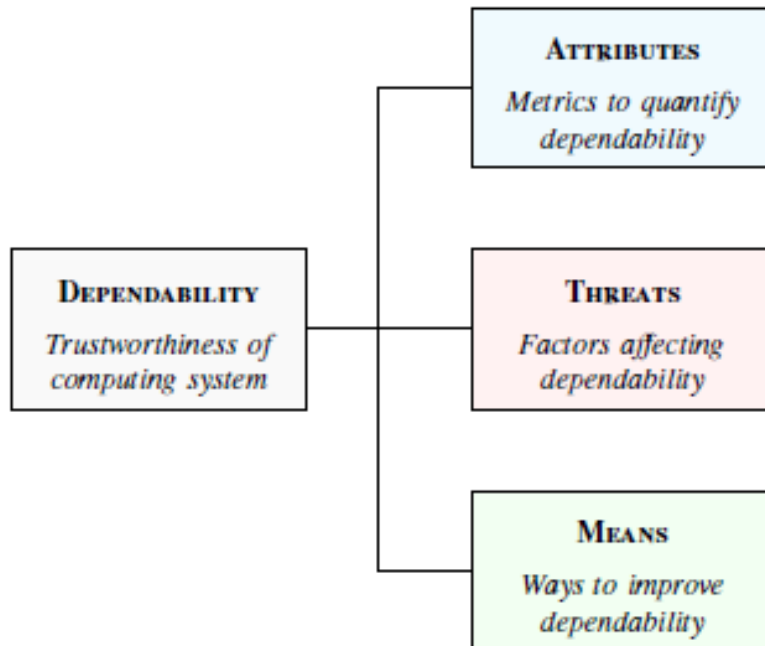


- ONOS response time increases linearly at constant workload
- **Response time increases 50%** for intent installation and withdrawal after the first day of operation

Prevention



Summary



More **metrics** are required to quantify the software dependability:

- Temporal reliability variations due to maturity and aging
- User-perceived service availability

Improved **threat** analysis to identify and classify software threats

Improved **threat** models and characterization

Software-aware **means**:

- (In)efficiency of software redundancy
- Network software rejuvenation



Questions?

Own related Publications



[J1] Vizarrreta, Petra; Trivedi, Kishor; Helvik, Bjarne; Heegaard, Poul; Blenk, Andreas; Kellerer, Wolfgang; Mas Machuca, Carmen, *Assessing the Software Maturity of SDN Controllers Using Software Reliability Growth Models*. Transactions on Network and Service Management (TNSM), June 2018

[J2] Vizarrreta, Petra; Van Bemten, Amaury; Sakic, Ermin; Abbasi, Khawar; Petroulakis, Nikolaos; Kellerer, Wolfgang; Mas Machuca, Carmen *Incentives for softwarization of wind park communication networks*, IEEE Communication Magazine, 2019

[J3] Vizarrreta, Petra; Trivedi, Kishor; Mendiratta, Veena; Kellerer, Wolfgang; Mas Machuca, Carmen, *DASON: Dependability Assurance Framework for Imperfect Distributed SDN Implementations*, Transactions on Network and Service Management (TNSM), Volume: 17, Issue: 2, June 2020

[J4] Vizarrreta, Petra; Sieber, Christian; Blenk, Andreas; Van Bemten, Amaury; Ramachandra, Vinod; Kellerer, Wolfgang; Mas-Machuca, Carmen; Trivedi, Kishor., *ARES: A Framework for Management of Software Ageing and Rejuvenation in SDN* , Transactions on Network and Service Management (TNSM), October 2020

[C1] Stojasavljevic, Petra; Trivedi, Kishor; Helvik, Bjarne; Heegaard, Poul; Kellerer, Wolfgang; Mas Machuca, Carmen, *An Empirical Study of Software Reliability in SDN Controllers*. CNSM, Tokyo, Japan, 2017

[C2] Vizarrreta, Petra; Sakic, Ermin; Kellerer, Wolfgang; Mas Machuca, Carmen *Mining Software Repositories for Predictive Modelling of Defects in SDN Controller*, In Proc. of IFIP/IEEE International Symposium on Integrated Network Management (IM), April 2019

[C3] tojsavljevic, Petra; Heegaard, Poul; Helvik, Bjarne; Kellerer, Wolfgang; Mas Machuca, Carmen, *Characterization of Failure Dynamics in SDN Controllers*. In Proc. of IEEE Int. Workshop on Reliable Networks Design and Modeling, Alghero, Italy, 2017

References (I)

SDN controllers

- [Ros14] F. J. Ros and P. M. Ruiz, “Five nines of southbound reliability in software-defined networks,” in Proceedings of the third workshop on Hot topics in software defined networking. ACM, 2014, pp. 31–36.
- [Onos17] ON.Lab, “ONOS: Open Neetwork Operating System,” <http://onosproject.org/>, 2017.
- [Odl17] Linux Foundation, “Opendaylight.” [Online]. Available: <https://www.opendaylight.org/>

Modelling approach

- [SAN01] W. H. Sanders and J. F. Meyer, “Stochastic activity networks: Formal definitions and concepts,” in Lectures on Formal Methods and PerformanceAnalysis. Springer, 2001, pp. 315–343.
- [Möb00] D. Daly, D. D. Deavours, J. M. Doyle, P. G. Webster, and W. H. Sanders, “Möbius: An extensible tool for performance and dependability modeling,” in International Conference on Modelling Techniques and Tools for Computer Performance Evaluation. Springer, 2000, pp. 332–336.

Reliability growth and ageing

- [JM72] Z. Jelinski and P. B. Moranda, “Software reliability research,” Statistical Computer Performance Evaluation, pp. 465–484, 1972.
- [Huang95] Y. Huang, C. Kintala, N. Kolettis, and N. D. Fulton, “Software rejuvenation: Analysis, module and applications,” in Fault-Tolerant Computing, 1995. FTCS-25. Digest of Papers., Twenty-Fifth International Symposium on. IEEE, 1995, pp. 381–390.

References (II)

Fault mitigation

[Qin05] F. Qin, J. Tucek, J. Sundaresan, and Y. Zhou, “Rx: treating bugs as allergies—a safe method to survive software failures,” in *Acm sigops operating systems review*, vol. 39, no. 5. ACM, 2005, pp. 235–248.

[Gro07] M. Grottke and K. S. Trivedi, “Fighting bugs: Remove, retry, replicate, and rejuvenate,” *Computer*, vol. 40, no. 2, 2007.

[Tri10] K. S. Trivedi, M. Grottke, and E. Andrade, “Software fault mitigation and availability assurance techniques,” *International Journal of System Assurance Engineering and Management*, vol. 1, no. 4, pp. 340–350, 2010.

Model parameters

[Kim09] D. S. Kim, F. Machida, and K. S. Trivedi, “Availability modeling and analysis of a virtualized system,” in *Dependable Computing, 2009. PRDC’09. 15th IEEE Pacific Rim International Symposium on*. IEEE, 2009, pp. 365–371.

[Nec16] G. Nencioni, B. E. Helvik, A. J. Gonzalez, P. E. Heegaard, and A. Kamisinski, “Availability modelling of software-defined backbone networks,” in *Dependable Systems and Networks Workshop, 2016 46th Annual IEEE/IFIP International Conference on*. IEEE, 2016, pp. 105–112.