



# Resilience assessment and dependability benchmarking: challenges of prediction

Lorenzo Strigini

Centre for Software Reliability  
City University, London, U.K.  
strigini@csr.city.ac.uk

DSN 2008 Workshop on Resilience Assessment and Dependability Benchmarking

1

## Outline

- a condensed ensemble view of the topics
- some known difficulties
  - especially with *representativeness*
- evolution of needs
- some ways forward that we could explore

# Empirical assessment of dependability and resilience: why?

various purposes, e.g.:

- direct prediction of dependability of a specific system in use
- same, but with emphasis on *comparing* systems
- coverage assessment of fault tolerance mechanisms, to refine design of a specific system
- same, to steer design of future systems
- dependability or coverage measurements for categories of systems and mechanisms, to assist in predictions about future systems
- ....

3

## Declaration of interests

- much of my work is about safety *assessment* of critical software
- users need strong confidence that a system will be safe enough, *before* seeing any real operation (quantitative predictions)
- in this, I am a *consumer* of measurement data
  - wishing to find in measurement a stronger basis than poorly justified existing consensual processes
  - affected by problems of poor data quality, accessibility, difficulties of measurement...
  - but also by more fundamental issues with prediction

4

# Empirical assessment of dependability and resilience: how?

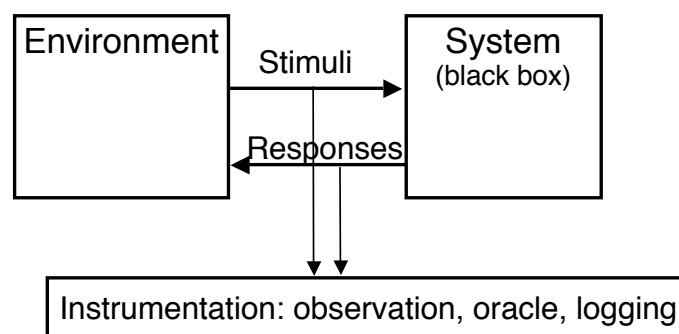
At least two broad ways:

- show that (how often) the system works well despite stress, shocks and flaws
- show that (how often) the defences in the system work well against stress, shocks and flaws
- both are forms of *prediction* of an uncertain future, based on *measurement*, involving *models*
- Prediction is difficult...

5

## The dependability assessment view

- observe the system in operation (in the lab or in real life or somewhere in between)

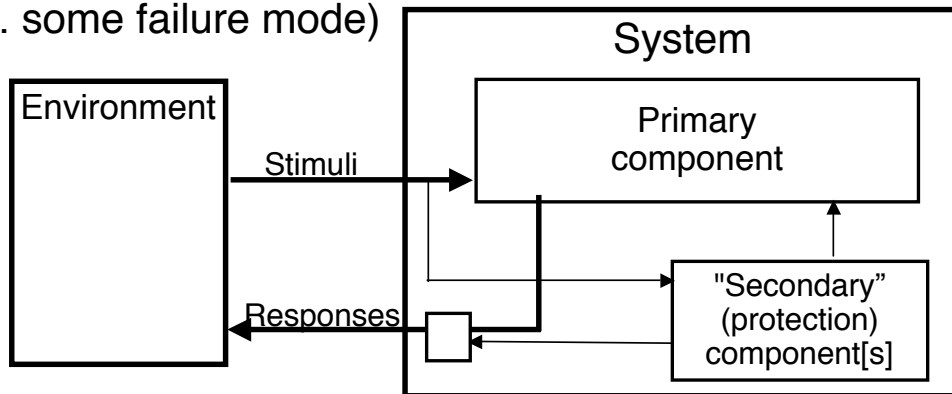


- take direct measures to calculate the dependability measure of interest
  - e.g. probability of failure per demand, MTBF, ...
  - plug into a *prediction system* based on some model
    - + a typical, simple model is "the future will be like the past"
    - + or sometimes: "the future will be a continuation of past trends"

6

# The fault tolerance and coverage view

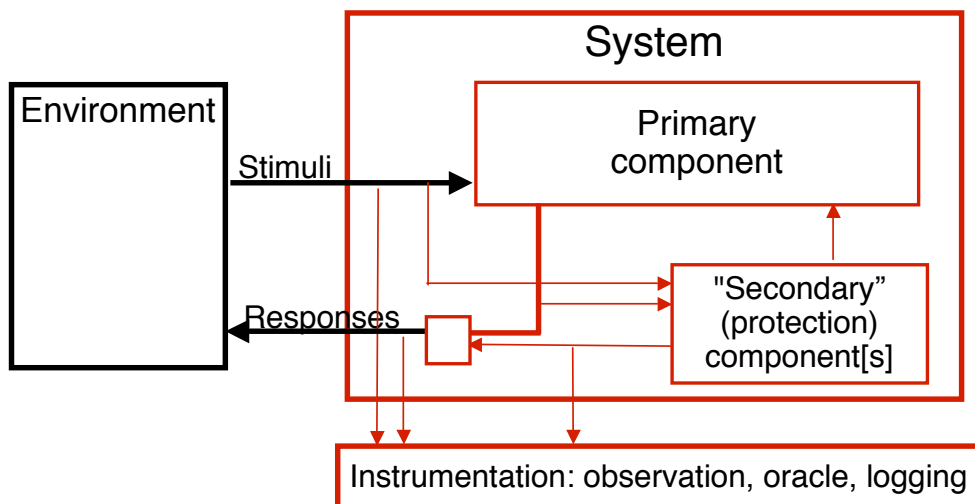
Look at a clear-box model of the system:  
(w.r.t. some failure mode)



then  $P(\text{system failure}) = P(\text{primary fails}) * \underbrace{P(\text{secondary fails} \mid \text{primary fails})}_{1\text{-coverage factor}}$

7

## Coverage assessment



- use the secondary (in the lab or in real life or ...)
- measure primary's errors and successful handling-> coverage
  - you only need a sample of *erroneous* responses from the primary
  - an attractive possibility, for right reasons and wrong reasons
  - a natural application for *fault injection*: save time
  - many sophisticated techniques have been developed ...

8

## Attraction of coverage estimation

- gives insight into why the system behaves well (or badly)
- seems cheap:
  - if I need system probability of failure  $< 10^{-7}$  (expensive to test for!) I can just test for  $10^{-4}$  for primary,  $10^{-3}$  for protection
  - usually wrong: I need the "right" sample of primary's errors
- robust:
  - when system moves to a more stressful environment (when reusing mechanism in a new system), I can re-use my estimate of coverage with the new probability of primary failure
  - wrong in the general case: expect different error distributions
- good at least for comparison:
  - if system A has higher coverage for each fault class than system B, it will be more dependable, given the same environment of use
  - wrong in the general case: e.g. consider tolerance of operator faults, or of internal design faults: B might have fewer to start with

9

## A spectre is haunting measurement ...

... the spectre of unrepresentativeness

(the risk of the measurement conditions not being *statistically* representative of the conditions for which prediction is sought)

10

## Measurement and representativeness

A generic description of system failure probability:

define

- $P(x), P(f)$ : distributions of demands (generalised inputs) and of faults
- the spaces of possible demands and possible responses by the primary
- and Boolean functions
  - $iswrong(x,y)$ :  $y$  is by specification an incorrect response to  $x$
  - $out(x,y,f)$ : the response of primary to demand  $x$  given fault  $f$  is  $y$
  - $gap(x,y)$ : secondary will not "cover" (e.g. detect) erroneous response  $y$  on demand  $x$

(Note: at least  $P(f)$ ,  $out(x,y,f)$  are usually unknown)

then:

$$P(\text{system fails}) = \sum_{x \in \text{demands}} P(x) \sum_{y \in \text{responses}} iswrong(x,y) \sum_{f \in \text{faults}} P(f) out(x,y,f) gap(x,y)$$

11

## Probability of failure, coverage

With the notation defined above

$$P(\text{system fails}) = \sum_{x \in \text{demands}} P(x) \sum_{y \in \text{responses}} iswrong(x,y) \sum_{f \in \text{faults}} P(f) out(x,y,f) gap(x,y)$$

coverage = 1 -  $P(\text{secondary fails to handle error} \mid \text{error of primary}) =$

$$1 - \frac{\sum_{x \in \text{demands}} P(x) \sum_{y \in \text{responses}} iswrong(x,y) \sum_{f \in \text{faults}} P(f) out(x,y,f) gap(x,y)}{\sum_{x \in \text{demands}} P(x) \sum_{y \in \text{responses}} iswrong(x,y) \sum_{f \in \text{faults}} P(f) out(x,y,f)}$$

... both depend on distributions of demands *and* faults

- some errors in distribution may cause serious errors in predictions
  - even reversing ranking between two systems

12

## In practice ...

We need to sample the space of two independent variables:




- if possible we get the sample from real life
  - or a close simulation of it
  - in any case, faster to obtain sample of demands than of faults

13

## Sampling in practice - 2

We need to sample the space of faults and demands:



- if we need to sample artificially, a simplification is: stratify using "sensible" classifications of demands, faults
  - then try to verify that faults during measurement (e.g. injected faults) "qualitatively" (i.e. by categories) match real faults
    - + then assume *some* distribution within each stratum
  - hoping that wrong choices of distribution here are less of a problem than using arbitrary general distribution
  - i.e. that error likelihood and coverage are "reasonably constant" within each stratum

14

# How good is "practical" sampling of faults?

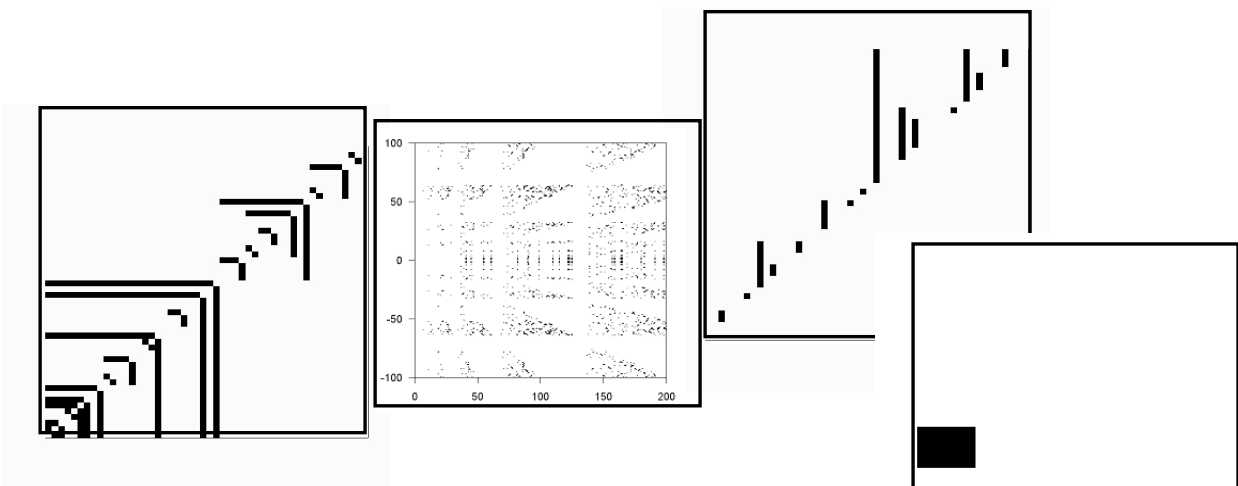
- in some cases, it "must" be good
  - e.g. hardware transients due to uniform radiation
- in some, it "must" be less good
  - e.g., software faults

(note: here we deal with "epistemic uncertainty"  
but this is not *the* difficulty)

15

## Software faults

- we have practical ways of classifying them
- but are these a useful stratification for sampling ...
- ... given that faults' mapping into "failure regions" is so often surprising



*(maps of failure regions in simple programs, courtesy of  
Meine van der Meulen - cf forthcoming article in TSE)*

## So, how full is the glass?

- great strides forward in measurement and injection techniques
- ... *but* inaccurate sampling of fault (or demand) distribution may make predictions of dependability and resilience wrong and even misleading
- this problem may be very serious in e.g. safety certification
  - but could also subvert guidance for design
  - or benchmarking - consensus does not help
- yet, across all uses of assessment and benchmarking
  - we can see that some predictions are sounder than others
  - I believe we could argue more clearly *why*, in each case
    - + e.g. do we trust that *this* fault profile has "acceptable realism"? (assessed on what metric?)
    - + or that coverage is "reasonably flat" on some set of faults? Why?

17

## Interesting ways forward through research?

- we could reduce the gap
  - between the "feeling" of security given by having *measurements* and the cold negativism of the basic formulas
- through internal validation: clarifying
  - assumptions that would justify modest claims like
    - + if system A's measure of reliability (coverage, ...) is better than system B's in test conditions, A will be better than B in real life
    - + or at least A "likely to be better" than B in real life?
    - + e.g., stochastic ordering given the future unknown regimes of operation (or systems)?
  - does one's intuition of a certain "plausible" fault profile satisfy any of these conditions?
    - + e.g. link empirical research on types of faults/errors to statistical statements of above assumptions
- or external validation
  - comparing predictions with outcomes
  - as engineers we are happy to use "wrong" models, *IF* their predictions are good enough

18

## Evolving needs, trends, challenges

there is a need to extend measurement and prediction towards

- a broader view of systems and threats
- a broader view of "resilience"

19

## A range of difficulty in sampling problems

For a complex system, failure scenarios range ...

- from "benign" ones: e.g., hardware in a shower of cosmic rays
  - + nature is "subtle, but not malicious"
- to some less so, e.g. design faults
- and some even more challenging
  - security:
    - + your enemy is subtle *and* malicious
    - + ... so, protective features actively cause changes in the fault profile
  - human resilience:
    - + a person's ability to react to an error varies with the likelihood of the error mode
    - + ... so, e.g., conservative predictions based on "improving a component will only improve the system" won't help

*but still well-definable problems in modelling, sampling, validation of models and distributions*

20

## More about "resilience"

- from the Latin verb for "to jump back": rebounding or springing back
  - ability to recover from adverse events; ability of a material to go back to original shape after being bent, compressed etc
- so, it means roughly "fault tolerance"
- but now often used to mean "the same but in a broader sense", e.g.
  - refusing some narrowing of meaning that those words may have undergone in a technical community
  - to refer to "actual" resilience rather than consensus view of "assessment of resilience"
  - to include more "unforeseen" events
  - to include effects of evolution (of system and environments of use)
  - to apply to largely open, self-evolving systems of systems

21

## Questions from broader concepts of "resilience"

- to tolerate broader ranges of events, recommendations include
  - more open-ended, flexible defences
    - + sounds reasonable... but empirical validation will be needed
  - extra resources/redundancy: increase coverage for *all* possible errors
- need to map trade-offs between horizon of prediction, precision of prediction, and thus inevitable risks (from waste, from failures)
- problems with "unforeseen events"?
  - identified and considered unlikely/unimportant?
    - + then, we mean "tolerance for faults under more variable fault profile"
    - + not a radical change: we need artificial sampling, ..... stratification ... more "what if" modelling ...
    - + and we can use that extra clarity about what makes a prediction process "good enough"
  - outside the boundaries of one's imagination?
    - + no remedy except stretching imagination

22

## Conclusions?

- prediction remains difficult
- I have discussed the problem of representativeness of measurement conditions
  
- there are promising directions for taming this problem
  - through more explicit description of factors
  - to extend trustworthiness of current techniques
  - and allow application to new needs with more complete system views, larger and less predictable systems