How do we know that our system is correct?

Hana Chockler Department of Informatics King's College, London





How to verify computerized systems?













Testing



Execute the program on a test suite and inspect manually or (semi-)automatically

"To know that we know what we know, and that we do not know what we do not know, that is true knowledge." Confucius

Simulation-based Verification and Testing

Observations:

- The simulation/testing process is sampling
- The sampling should be as exhaustive as possible
- The process can only find bugs there is no way to guarantee that the system under test is correct

How exhaustive is the sampling? Did we check all locations of most complex functionality?

Various coverage metrics are widely used as heuristic measures of exhaustiveness of verification







Pass in Model Checking: Is it really correct?

Model checking ≠ sampling The whole reachable state space is visited

Do we know that the system is correct if model checking passes?

verified

Did I check

everything I

wanted to

check?

Correctness of the "pass" result depends on correctness and exhaustiveness of the specification

Did I check what I wanted to check?

9

Suspecting a positive answer [IBM, Intel]



The story of vacuity







The story of vacuity





Suspecting a positive answer [IBM, Intel]



The story of coverage



17



The most general definition:

In model checking: an element is covered ↔ The system with the element modified (*mutated*) no longer satisfies the specification

> elements are small (atomic) mutations are small changes
> if a mutant system still satisfies the specification → we did not check this particular corner of the system

Coverage

φ = always (req -> eventually grant)



A non-covered mutation represents a "corner" of the system that was not verified (and hence may contain bugs)

The story of coverage



Different types of mutations

- Flipping the value of one output signal in one state (Intel)
- Same with control signals
- Freeing a signal (turning it to an input)
- VHDL code mutations (erasing or changing a line)
- Removing behaviors (by removing states or otherwise)
- Changing one net in the net-list



mutations are small changes either in the signals in the representation inside the tool, or in the code written by the designer

The story of coverage



timeline



Problem: how to compute coverage efficiently?

Model-checking each mutant system separately is infeasible

Useful observation - all mutant systems are similar!



This leads to feasible and efficient algorithms:

- Symbolic algorithms: represent all possible mutants together by adding a small number of symbolic variables
- Improving average complexity: consider possible mutations as unknown values and attempt to compute as much as possible without assigning them

Was implemented and actually works at IBM

- **Reusing** the results of verification of the original system:
 - o Interpolation-based coverage computation: save the proof and reuse it
 - o **ic3-based** coverage computation: save the high-level proof or symbolic partial counterexample and re-use it

Reusing the results of verification





A byproduct - efficient regression verification methodology





Proofs and counterexamples are general enough to work for many mutant systems – in particular, for a new version If they don't work, they can be "patched" to work – cheaper than re-verification



The story of coverage



Results of the feasibility check - done at IBM Mutations of a hardware design

There exist many non-covered mutations

 There are hundreds of thousands of automatically generated mutations

0

- Not all of them are interesting
- Those that are interesting should be examined by the verification engineer and/or the designer
- Not all non-covered mutations point to a bug - but some do



unhappy designer

How to minimize the manual effort in checking coverage results?



Automatic analysis of results

- Identify non-covered areas of the design
- Construct non-covered traces
- Suggest properties that have better coverage

The holy grail of sanity checks:



The story of coverage



Do we need a metric that refines coverage?

Problem: some specifications almost <u>always</u> have low coverage!

<u>Example:</u> φ = "f is computed at least twice" (for fault tolerance)



Replacing coverage with causality and responsibility (from AI):

And element is covered if its responsibility is 1.

<u>Example:</u> φ = "f is computed at least twice" (for fault tolerance)



for formal verification and other things

<u>Side note:</u> causality and responsibility are useful in general



 ϕ = G((-START \land -STATUS VALID \land END) \rightarrow X[-START U (STATUS VALID \land READY)]).







What if we can just generate a correct system automatically from the specification?







Suspecting a positive answer [IBM, Intel]





Synthesizing non-vacuous systems [BCES17]

Vacuity is a non-interesting pass of a specification in the system. But we don't have a system yet! What do we do?

- Given a specification φ , strengthen it so that it does not allow vacuous satisfaction.
 - o "There must exist an interesting behaviour".
- This results in a formula with existential quantifiers.
- Synthesize the resulting formula.
- Improve the system iteratively until there are many interesting behaviours.

Synthesizing non-vacuous systems [BCES17]

 φ = always (req-> eventually grant)

Must allow any behaviour of inputs (req) (that is, the printer is connected)



What about a printer that prints regardless of requests? ψ = always (eventually(grant))

Must be at least one execution where eventually it stops printing $\neg \psi = \exists$ (eventually (always(!grant)))



So can we switch to synthesis and stop designing systems?







- We only know how to do non-vacuous bounded synthesis
- Nobody knows how to synthesize software
- We don't know how to generate the best possible specifications automatically

But hopefully in the future





