Using Fault History to Improve Mutation Reduction

Laura Inozemtseva
University of Waterloo
lminozem@uwaterloo.ca

Hadi Hemmati
University of Manitoba
hemmati@cs.umanitoba.ca

Reid Holmes
University of Waterloo
rtholmes@uwaterloo.ca

Scan for preprint

Problem

Mutation testing can be too slow for practical use because of the large number of mutants.

New Idea

A reduction method that tries to select the most "interesting" mutants; that is, mutants that will survive mutation testing.

Results

74 percent of the time, our method chose a greater number of interesting mutants than random reduction did.

The old solution: reduce the mutants at random.

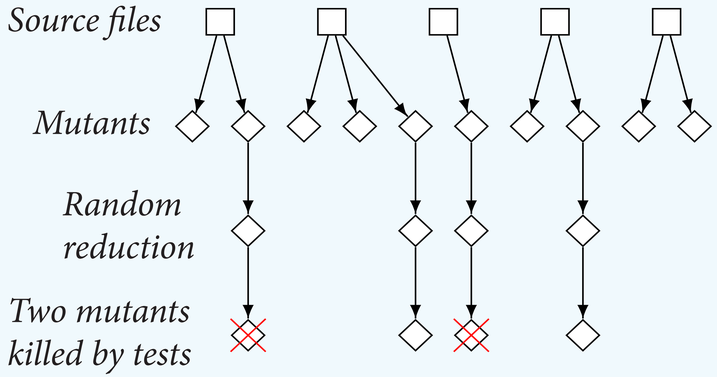

The surviving mutants indicate possible weaknesses in the test suite. They are therefore more interesting to developers than the killed mutants. Our goal was to reduce the number of mutants in a way that preserves these interesting ones.

We reduced the set of mutants to the *faulty-file mutants*, or mutants made by altering a file that has contained many faults in the past. We hypothesized that faulty file mutants are more likely to survive; that is, they are the more interesting mutants.

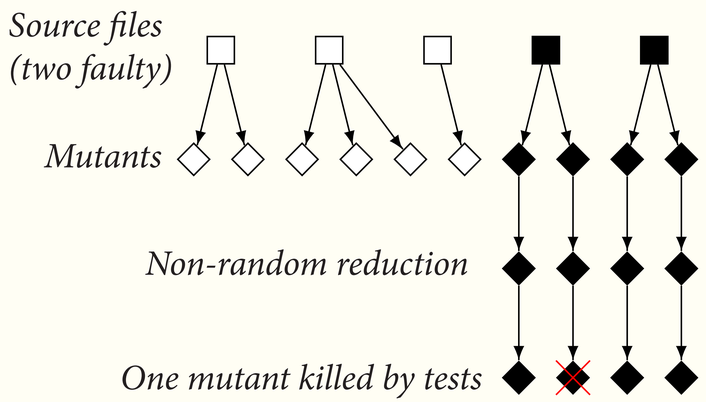

We ran 6000 test suites on Apache POI's faulty file mutants and on an equal number of randomly selected mutants. Our hypothesis was correct, but the results varied with suite size, ranging from 59% for size 100 suites to 97% for size 1000 suites.

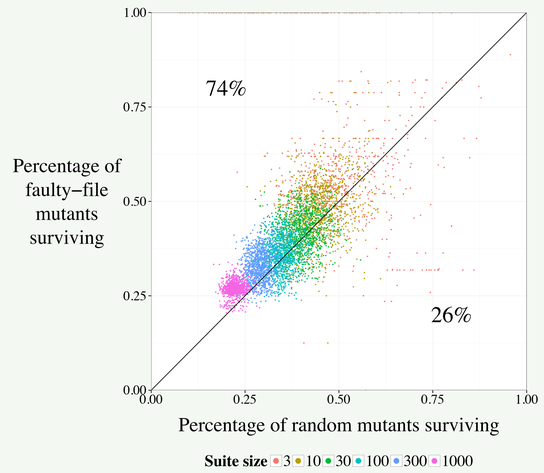

Comments