## Automated Program Repair

Opportunities, Challenges, Advances

**Chris Timperley** 



## About me...



ctimperley@cmu.edu

### **Research Interests:**

- Automated Program Repair
- Fault Localisation
- GI for Robotics
- Automated Test Generation

### **Recent Projects:**

- BugZoo: reproducible studies of historical bugs
- **Rooibos:** language-independent, syntax-aware search and transformation.
- Houston: automated testing for robotics

### Postdoc

Carnegie Mellon University, USA w/ Claire Le Goues



### MEng (2013), PhD (2017)

University of York, UK w/ Susan Stepney



**Purpose of Talk:** Challenge existing views, and identify opportunities.



# 2008

Genetic Programming to modify existing programs, rather than building them from scratch.

Demonstrates concept of automated program repair.



### **Evolutionary Repair of Faulty Software**

### Andrea Arcuri

The School of Computer Science, The University of Birmingham, Edgbaston, Birmingham B15 2TT, UK. Email: a.arcuri@cs.bham.ac.uk

### Abstract

Testing and fault localization are very expensive onflware engineering tasks that have been tried to be automated. Although many successful techniques have been designed, the actual change of the code for fixing the discovered faults is still a human only task. Even in the ideal case in which automated tools could tell us eatedly where the location of a fault is, it is not adways trivial how to fix the code. In this paper we analyse the possibility of automaing the complex task of fixing faults. We propose to model this task as a search problem, and hence to us for example evolutionary algorithms to solve it. We then discuss the potential of this approach and how its current limits can be addressed in the future. This task is extremely challenging and mainly unexplored in literature. Hence, this paper only covers an initial investigation and gives directions for future work. A research prototype called JAFF and a case study are presented to give first vidiation of this approach.

Keyword: Repair, Fault Localization, Automated Debugging, Genetic Programming, Search Based Software Engineering, Coevolution.

### 1 Introduction

Software testing is used to reveal the presence of faults in computer programs [50]. Even if no fault is found, testing cannot guarantee that the software is fault-free. However, testing can be used to increase our confidence in the software reliability. Unfortunately, testing is expensive, time consuming and tedious. It is estimated that testing requires around 50% of the total cost of software development [14]. This is the reason why there has been a lot of effort spent to automate this expensive offware engineering task.

Even if an optimal automated system for doing software testing existed, we still need to know where the faults are located, that in order to be able to fix them. Automated techniques can help the tester in this task [26, 65, 78].

Although in some cases it is possible to automatically locate the faults, there is still the need to modify the code to remove the faults. Is it possible to automate the task of fixing faults? This would be the natural next step if we seek a full automation of software engineering. And it would be particularly helpful in the cases of complex software in which, although the faulty part of code can be identified, it is difficult to provide a patch for the fault. This would also be a step forward to achieve corporate visions like for example IBM's Autonomic Computing [40].

There has been work on fixing code automatically (e.g., [63, 61, 68, 25]). Unfortunately, in that work there are heavy constraints on the type of modifications that can be automatically done on the source code. Hence, only limited classes of faults can be addressed. The reason for putting these constraints is that there are infinite ways to do modifications on a program, and checking all of them is impossible.

**Evolutionary repair of faulty software.** Andrea Arcuri. 2011. Applied Soft Computing 11, 4 (June 2011), 3494-3514.



PEOPLE EXPERIMENTS PUBLICATIONS VIDEOS REPOSITO

### GenProg

volutionary Program Repair

### **The Problem**

Software engineering is expensive, summing to over one half of the US GDP annually. Software maintenance accounts for o of that life cycle cost, and a key aspect of maintenance is fixin existing programs. Unfortunately, the number of reported bu available development resources. It is common for a popular hundreds of new bug reports filed every day.

Software maintenance is expensive. GenProg reduces softw: maintenance costs by automatically producing patches (repai defects.

### **Our Approach**

Many bugs can be fixed with just a few changes to a program's Human repairs often involve inserting new code and deleting existing code. GenProg uses those same building blocks to se automatically.

GenProg uses genetic programming to search for repairs. Ou

Automatically Finding Patches Using	Genetic Program	mming
-------------------------------------	-----------------	-------

 Westley Weimer
 ThanhVu Nguyen
 Claire Le Goues

 University of Virginia
 University of New Mexico
 University of Virginia

 weimer 8 virginia, edu
 tnguyen 8 cs.unn.edu
 legoues8 virginia.edu

Claire Le Goues Stephanie Forrest University of Virginia University of New Mexico egoues8 virginia.edu for rest 8 cs. unn.edu

To alleviate this burden, we propose an automatic tech-

nique for repairing program defects. Our approach does

not require difficult formal specifications, program annotations or special coding practices. Instead, it works on

off-the-shelf legacy applications and readily-available test-

cases. We use genetic programming to evolve program vari-

ants until one is found that both retains required function-

ality and also avoids the defect in question. Our technique

takes as input a program, a set of successful positive test-

negative testcase that demonstrates a defect.

cases that encode required program behavior, and a failing

Genetic programming (GP) is a computational method

inspired by biological evolution, which discovers computer

programs tailored to a particular task [19]. GP maintains a

population of individual programs. Computational analogs

of biological mutation and crossover produce program vari-

ants. Each variant's suitability is evaluated using a user-

defined fitness function, and successful variants are selected

A significant impediment for an evolutionary algorithm like GP is the potentially infinite-size search space it must

sample to find a correct program. To address this problem, we introduce two key innovations. First, we restrict the al-

gorithm to only produce changes that are based on structures in other parts of the program. In essence, we hypoth-

esize that a program that is missing important functionality (e.g., a null check) will be able to copy and adapt it from

another location in the program. Second, we constrain the

genetic operations of mutation and crossover to operate only

on the region of the program that is relevant to the error (that

is, the portions of the program that were on the execution

path that produced the error). Combining these insights, we

demonstrate automatically generated repairs for ten C pro-

We use GP to maintain a population of variants of that program. Each variant is represented as an abstract syn-

tax tree (AST) paired with a weighted program path. We modify variants using two genetic algorithm operations, crossover and mutation, specifically targeted to this repre-

grams totaling 63,000 lines of code

for continued evolution. GP has solved an impressive range of problems (e.g., see [1]), but to our knowledge it has not

been used to evolve off-the-shelf legacy software.

### Abstract

Automatic program repair has been a longstanding goal in software engineering, yet debugging remains a largely manual process. We introduce a fully automated method for locating and repairing bugs in software. The approach works on off-the-shelf legacy applications and does not require formal specifications, program annotations or special coding practices. Once a program fault is discovered, an extended form of genetic programming is used to evolve program variants until one is found that both retains required unctionality and also avoids the defect in question. Standard test cases are used to exercise the fault and to encode program requirements. After a successful repair has been discovered, it is minimized using structural differencing algorithms and delta debugging. We describe the proposed method and report experimental results demonstrating that it can successfully repair ten different C programs totaling 63 000 lines in under 200 seconds on avenue

### 1 Introduction

Fixing bugs is a difficult, time-consuming, and manual process. Some reports place software maintenance, traditionally defined as any modification made on a system after its delivery, at 90% of the total cost of a trypical software project [22]. Modifying existing cosh, repairing deficies, society [24]. The mathematical solution of the solution target of the solution of the solution of the solution and unknown bugs [21] because they lack the development resource to ded with every defect. For example, in 2032, bugs appear [...] for too much for only the Mozilla programment to handle [55, p. 8-33].

\*This research was supported in part by National Science Foundation Grants CNS 0627523 and CNS 0716478, Air Force Office of Scientific Research grant FA9550-07-1-0532, as well as gifts from Microsoft Research. No official endorement should be inferred.

ICSE'09, May 16-24, 2009, Vancouver, Canada 978-1-4244-3452-7/09/825.00 © 2009 IEEE

d Romoed use limited to: University of Virginia Libraries. Downloaded on November 5, 2009 at 11:05 from IEEE Xplore. Restrictions apply

364

## 2009

### GenProg demonstrates program repair on real-world C programs

### Automatically Finding Patches Using Genetic Programming.

Westley Weimer, ThanVu Nguyen, Claire Le Goues, Stephanie Forrest. Proceedings of International Conference on Software Engineering. ICSE '09. 2009.







## Where are the program repair bots?

Or, why aren't we all filthy rich yet?



### ICSE SEIP '18

### How to Design a Program Repair Bot? Insights from the Repairnator Project

Simon Urli University of Lille & Inria Lille, France simon.urli@inria.fr

Lionel Seinturier University of Lille & Inria Lille, France lionel.seinturier@inria.fr

ABSTRACT

Program repair research has made tremendous progress over the last few years, and software development bots are now being invented to help developers gain productivity. In this paper, we invesigate the concept of a "program repair bot" and present Repairnator. The Repairnator bot is an autonomous agent that constantly monitors test failures, reproduces bugs, and runs program repair tools against each reproduced bug. If a patch is found, Repairnator bot reports it to the developers. At the time of writing, Repairnator uses three different program repair systems and has been operating since February 2017. In total, it has studied 11317 test failures over 1609 open-source software projects hosted on GitHub, and has generated patches for 17 different bugs. Over months, we hit a number of hard technical challenges and had to make various design and engineering decisions. This gives us a unique experience in this area. In this paper, we reflect upon Repairnator in order to share this knowledge with the automatic program repair community.

#### ACM Reference Format:

Simon Urli, Zhongxing Yu, Lionel Seinturier, and Martin Monperrus. 2018. How to Design a Program Repair Bot'l Indjish from the Repairmote Project. In Proceedings of 4th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP '16). ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/rannum.nnmnn

### 1 INTRODUCTION

Program repair research has made temendous programs over the last few years (b, 5, 12, 11), as varied to dimensional end of the last few years (b, 5, 12, 11), as varied to dimensional (13, 12, 20), thus been shown that current program repair tendings in a called to synthesize patches for end logs in called gas regularities generdity only evaluate the capability of the repair distributions in the structure of program repair tendings in practice, second other perturgs are also accounted before an off the true of the core repair algorithm itself. To demonstrate the real potential of program repair infinitely, it is distributed to study the design and implementation of an end-to-end repair toolchin fin it is amenable to the mainstream development precise.

For bridging this gap between research and industrial use, we investigate the concept of a "program repair bot" in this paper. To us, a program repair bot is an autonomous agent that constantly

ESE-SEB '18, May 27-June 3 2018, Gothenburg, Sweden 2018. ACM ISBN 976-5-5558-542(Y)/ML...\$15.00 https://doi.org/10.1145/nnnnnn.nnnnnnn Zhongxing Yu University of Lille & Inria Lille, France zhongxing.yu@inria.fr

Martin Monperrus KTH Royal Institute of Technology, Sweden martin.monperrus@csc.kth.se

monitors test failures reproduces bugs, and nume program repair tools against each reproduced bug if a patch is found, the program repair but reports it to the developers. We envision that in ten years from now there will be hundreds of program repair bots that will work in concert with developers to maintain large code bases. But today, to the best of our knowledge, modely and ever today, to the best of our knowledge, modely and ever to the design and operation of aich a repair bot. The Repairment project is a project to design, implement and

operate a reput bof for Jone program. This reput bot itself is called after the project: "Equivators," and the name will only refer to the bot in the remaining of the paper. Repairmants constants monitors itself abates happening in continuous integration, also monitors in the abates happening in continuous antigration of the state of the state of the state of the papers on CL Repairment fruct itrus to reproduce the CL failure, then runs publicly available program repair tools to make a "engintic state".

At the time of writing, Regularizator uses there different program repair systems and has been operating usine reference y2017. In total, it has studied 11317 test failures over 1649 epen sources software projects based on a clifful, and has generated patches for 17 diffeent bugs. None of those patches were proposed to the developers because they all source from the overful time problem [72, 52, 60] they indeed fits the failing build but cannot be considered as a general, appropriate solution to the bug.

The design and operation of Repairmator has been challenging. Over months, we hit a number of hard technical challenges and had to make various design and engineering desions. This gives us a unique experience in this area. In this paper, we reflect upon Repairmator in order to share this knowledge with the automatic program regair community.

The pipeline of Repaintator is constituted by three stages: CI Boald Analysis, Bog Reproduction, and Platch Synthesia. For each of the three stages: (1) we present how it has been designed, aiming at inspiring the automode of upcoming pracis bots: (2) we report on results about the operation of Repaintator itself over 'n months' of operiments' and (1) we present and discuss activable in recomergorience in architecting and operating Repaintators experiment in architecting and operating Repaintators.

 a blueprint design of a program repair bot for continuous integration (CI) test failures;

# they exist!

### (note: we found out about them a few weeks ago)

### How to Design a Program Repair Bot? Insights from the Repairnator Project.

Simon Urli, Zhongxing Yu, Lionel Seinturier, and Martin Monperrus.

Proceedings of the International Conference on Software Engineering . ICSE '18.

## How do we deploy?

What do developers need?



**Easy Integration** 

Timeliness

**Bug Information** 

Patches

## The big assumption: The existence of tests



requires test cases

# Software engineering has changed since APR was introduced.







requires test cases



## What are the challenges?

## Challenges



Patch Quality

Scalability

Expressiveness

## Challenges: Patch Quality

	BUGGY				CORRECT				PLAUSIBLE				
max if( if(	<pre>max = 0; if(a &gt; b) max = a; if(b &gt; a) max = b;</pre>				<pre>max = a; if(a &gt; b) max = a; if(b &gt; a) max = b;</pre>				<pre>max = 7; if(a &gt; b) largest = a; if(b &gt; a) largest = b;</pre>				
а	b	max			a	b	max		a	b	max		
3	2	3			3	2	3		3	2	3		
4	5	5			4	5	5		4	5	5		
7	7	0	×		7	7	7		7	7	7		

## Not all good patches are correct.

CONTRONTRESIA

### ISSTA '15



### An Analysis of Patch Plausibility and Correctness for Generate-and-Validate Patch Generation Systems

Zichao Qi, Fan Long, Sara Achour, and Martin Rinard MIT EECS & CSAIL, USA {zichaoqi, fanl, sarachour, rinard}@csail.mit.edu

### ABSTRACT

We analyze reported patches for three existing generate-andvalidate patch generation systems (GenProg, RSRepair, and AE). The basic principle behind generate-and-validate systems is to accept only plausible patches that produce correct outputs for all inputs in the validation test suite. Because of errors in the patch evaluation infrastructure the majority of the reported patches are not plausible they do not produce correct outputs even for the inputs in the validation test suite. The overwhelming majority of the reported patches are not correct and are equivalent to a single modification that simply deletes functionality. Observed negative effects include the introduction of security vulnerabilities and the elimination of desirable functionality. We also present Kali, a generate-and-validate patch generation system that only deletes functionality. Working with a simpler and more effectively focused search space, Kali generates at least as many correct patches as prior GenProg. RSRepair, and AE systems. Kali also generates at least as many patches that produce correct outputs for the inputs in the validation test suite as the three prior systems We also discuss the natches nucleord by ClearView a generate-and-validate binary hot patching system that leverages learned invariants to produce patches that enable systems to survive otherwise fatal defects and security attacks.

Our analysis indicates that ClearView successfully patches 9

of the 10 accurity unbershilities used to evaluate the system.

At least 4 of these natches are correct

Categories and Subject Descriptors

1. INTRODUCTION

### Motivated to better understand the capabilities and po

Automatic patch generation holds out the promise of cor

recting defects in production software systems without the

time and expense required for human developers to under

stand, triage, and correct these defects. The prominent

inputs, at least one of which exposes a defect in the software. The patch generation system applies program mod-

ifications to generate a space of candidate patches, then

searches the generated patch space to find plausible natches

in the test suite. In this paper we start by considering the

GenProg [25], RSRepair [46], and AE [57] systems.

- i.e., patches that produce correct outputs for all inputs

generate-and-validate approach starts with a test suite o

Motivated to better understand the capabilities and potential of these systems, we performed an analysis of the patches that these systems produce. Enabled by the availability of the generated patches and the relevant patch generation and validation infrarenteure [6, 4, 3, 7, 1], our analysis was driven by the following research questions:

D.2.5 [Software Engineering]: Testing and Debugging General Terms RQ1: Do the reported GenProg. RSRepair, and AE patches produce correct outputs for the inpatch?

rastrutur

repair [25] auts for the

es that we

n the test

AE result pes correct

"In many cases the Kali patch cleanly identifies the exact functionality and location that the developer patch modifies"

"The Kali and developer patches typically modify common functionality and variables."

### GPCE '06

### Patches as Better Bug Reports

Westley Weimer University of Virginia weimer@virginia.edu

#### Abstract

Tools and analyses that find bags in software are becoming increasingly prevalent. However, even after the potential false a larmsraised by such tools are dealw with, many real reported errors may go unfixed. In such cases the programmers, have judged the benefit of fixing the bag to be less than the time cost of understanding and fixing it.

The true utility of a bug-finding tool lies not in the number of bugs it finds but in the number of bugs it causes to be fixed. Analyses that find safety-policy violations typically give error

reports as annotated backtraces or counterexamples. We propose that bug reports administrationally contain a specially-constructed patch describing an example way in which the program could be modified to avoid the reported policy violation. Programmers viewing the analysis output can use such patches as guides, starting points, or as an additional way of understanding what went wrong. We present an algorithm for automatically constructions such

We present an algorithm for automatically constructing such patches given model-rehecking and policy information typically already produced by most such analyses. We are not aware of any previous automatic techningers for generating patches in response not prevent in the original program, and can have help to explain bugs related to moising program elements. In addition, our patches do not introduce any new violations of the given safety policy. To evaluate our methods we generation also sholl by objects

revenues, applying our algorithm to over 70 bug reports produced by two off-the-shelf bug-finding tools running on large lava programs. Bug reports also accompanied by patches were three times as likely to be addressed as standard bug reports.

This work represents an early step toward developing new ways to report bugs and to make it easier for programmers to fix them. Even a minor increase in our ability to fix bugs would be a great increase for the quality of software.

Categories and Subject Descriptors D.2.5 [Software Engineering]: Testing and Debugging—Debugging aids: 1.2.2 [Automatic Programming]: Program Modification

General Terms Algorithms, Experimentation, Human Factors, Languages, Verification

Keywords bag, bug report, error, patch, counterexample, explanation, localization

### 1. Introduction

Tools and analyses that find bugs in software automatically are becoming increasingly prevalent. Such analyses usually proper false positives, and wading through spurious error report is one cost of using and sciola. Kives if false provides are controlled (e.g., the original programmer's invariable, the available programmer does not understand he bugs report on the specification, or the bugststand grant ding table tags associated to only eight be benefit understanding and ding table tags associated to only eight be benefit and transformer of the specification of the benefit to the specification of the specification of the benefit to the specification of the specification of the benefit to the specification of the specification of the benefit to the specification of the specification of the benefit to the specification of the benefit to the specification of the specification of the benefit to the specification of the specification of the benefit to the specification of the benefit to the specification of the specification of the benefit to the specification of the specification of the benefit to the specification of the specif

Analyses that find heps typically give error reports a annotated countercamples or hosticares (i.e., reading unth heoryth the perpara that demonstrate the bug). Utdivinuality, use h backraces are offend barg or afficient to instrate (if, a.g., Sky), by perporte that the describing an example way in which the program could be modified to avoid the reported policy visidatis without introducing any new large with respect to that policy. Our publics may clifted heighting the information is used places that policy outdoor without the program can use to patches as guides, starting points, or a san additional way of undorstanding what were invego where.

We present a novel algorithm for automatically constructing stuck patches given model-checking and policy information typically atready produced or required by the analysis. The algorithm builds on nearest accepting strings [86] and path predicates [31]. We then apply our algorithm to bag reports produced by two offhe-sheft blug finding tools [30, 34]. Our experiments indicate that bug reports also accompanied by our patches are more likely to be addressed than stundard bag reports.

In some commercial environments every big report from an official tool may be addressed on way or another, for example, some groups at Microsoft require that all PREfast to ESP [16], wrings be dealt with before a code change can be committed. We believe that in such a somatio our algorithm would reduce the time programmers spend addressing all of the bugs proposit (rather than increasing the number of bug reports deemed worth spending time on).

There are two main contributions of this work:

 A new algorithm for constructing candidate patches from the counterexamples produced by bugfinding tools. The patches can suggest inserting code not present in the original program. This is the first algorithm we are aware of that produces patches

"Bug reports also accompanied by [machine-generated] patches were three times as likely to be addressed as standard bug reports."

22

## Correctness is a major challenge, but overfitted patches can still be useful.

## Challenges



## Challenges: Expressiveness vs. Scalability



Expressiveness

Allows the program to be changed in a greater number of ways, increasing the odds of finding a modification that produces a repair.

- larger corpus of fix ingredients
- wider set of program transformations
- granular modifications to the program

## Challenges: Expressiveness vs. Scalability



Scalability

Time taken to discover a patch is a function of:

- patch size
- program size
- expressiveness
- ...





# How can we make APR both scalable and expressive?

## **Observation:** APR inherited most of its technologies.





**Abstract Syntax Tree** 



## Everyone is still using spectrum-based fault localisation!

- 2017: ssFix, Repairnator, NOPOL, ...
- 2016: History-Driven Program Repair, ...
- **2015:** Angelix, SearchRepair, ...
- 2014: Astor, RSRepair, ...
- 2013: SemFix, ...
- **2011:** AE, ...
- ...
- 2009: GenProg

### **Recap: Spectra-Based Fault Localisation**









Passing Test #1

Passing Test #2

Passing Test #3

$$\frac{e_f}{e_f + n_f + e_p} \qquad \frac{e_f}{\sqrt{(e_f + n_f) \cdot (e_f + e_p)}} \qquad \left| \begin{array}{c} \frac{e_f}{e_f + n_f} - \frac{e_p}{e_p + n_p} \right|$$

$$\text{Jaccard} \qquad \text{Ochiai} \qquad \text{AMPLE}$$

$$\begin{cases} 0, & \text{if } e_f + e_p = 0 \\ 1.0, & \text{if } e_f > 0 \land e_p = 0 \\ 0.1, & \text{otherwise} \qquad \frac{\frac{e_f}{e_f + n_f}}{\frac{e_p}{e_p + n_p} + \frac{e_f}{e_f + n_f}} \\ \text{Tarantula} \end{cases}$$

Shin Yoo. 2012. Evolving human competitive spectra-based fault localisation techniques. In Proceedings of the 4th international conference on Search Based Software Engineering (SSBSE'12), Gordon Fraser and Jerffeson Teixeira de Souza (Eds.). Springer-Verlag, Berlin, Heidelberg, 244-258.

```
if (namelist == NULL) {
   return -1;
if (!(dirp = opendir(dirname))) {
   return -1:
while (!php readdir r(dirp, (struct dirent *)entry, &dp) && dp) {
   size_t dsize = 0;
   struct dirent *newdp = NULL;
   if (selector && (*selector)(dp) == 0) {
   if (nfiles == vector_size) {
       struct dirent **newv;
       if (vector size == 0) {
           vector_size = 10;
           vector_size *= 2;
       newv = (struct dirent **) realloc (vector, vector_size * sizeof (struct dirent *));
       if (!newv) {
           return -1:
       vector = newv;
   dsize = sizeof (struct dirent) + ((strlen(dp->d name) + 1) * sizeof(char));
   newdp = (struct dirent *) malloc(dsize);
   if (newdp == NULL) {
       goto fail;
   vector[nfiles++] = (struct dirent *) memcpy(newdp, dp, dsize);
```

 $\otimes$ 2007 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for scale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

### On the Accuracy of Spectrum-based Fault Localization\*

Peter Zoeteweij

Rui Abreu

Arjan J.C. van Gemund

Software Technology Department Faculty of Electrical Engineering, Mathematics, and Computer Science Delft University of Technology P.O. Box 5031, NL-2600 GA Delft, The Netherlands {r.f.abreu, p.zoeteweij, a.j.c.vangemmed}}titudeff.nl

### Abstract

Spectrum-based fault localization shortens the testdiagnose-repair cycle by reducing the debugging effort. As a light-weight automated diagnosis technique it can easily be integrated with existing testing schemes. However, as no model of the system is taken into account, its diagnostic accuracy is inherently limited. Using the Siemens Set benchmark, we investigate this diagnostic accuracy as a function of several parameters (such as quality and quantity of the program spectra collected during the execution of the system), some of which directly relate to test design. Our results indicate that the superior performance of a particular similarity coefficient, used to analyze the program spectra, is largely independent of test design. Furthermore, nearoptimal diagnostic accuracy (exonerating about 80% of the blocks of code on average) is already obtained for low-quality error observations and limited numbers of test cases. The influence of the number of test cases is of primary importance for continuous (embedded) processing applications, where only limited observation

those faults that affect the user most can be solved before the release deadline, the efficiency with which faults can be diagnosed and repaired directly influences software reliability. Automated diagnosis can help to improve this efficiency.

Diagnosis techniques are complementary to testing in two ways. First, for test designed to verify correct behavior, they generate information on the root cause of test failures, forening the subsequent tests that are required to expose specific potential root cause, the extra information generated by diagnosis techniques can help to further reduce the set of remaining possible explanations. Given its incremental nature (i.e., taking into account the results of an entire sequence of tests), automated diagnosis alleviates much of the work of selecting tests in the latter category, and can hence have a profound impact on the test-diagnose-repair cycle.

An important part of diagnosis and repair consist in localizing faults, and several tools for automated debugging and systems diagnosis implement an approach

**Assumption:** several failing test cases (6 is optimal).

**Reality:** usually one failing test.

Takes > 12 hours to run GenProg and SearchRepair.

## Can we tailor fault localisation to CI-based program repair?









