

Bug Catching: Automated Program Verification

15414/15614 Fall 2018

Lecture 1: Introduction

Matt Fredrikson, Ruben Martins
{mfredrik, rubenm}@cs

August 28, 2018

Course Staff



Matt Fredrikson
Instructor



Ruben Martins
Instructor



Krishna Bagadia
TA



Rameel Rizvi
TA

- ▶ **April, 2014** OpenSSL announced critical vulnerability in their implementation of the Heartbeat Extension.



- ▶ **April, 2014** OpenSSL announced critical vulnerability in their implementation of the Heartbeat Extension.
- ▶ “The Heartbleed bug allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software.”



- ▶ **April, 2014** OpenSSL announced critical vulnerability in their implementation of the Heartbeat Extension.
- ▶ “The Heartbleed bug allows anyone on the Internet to read the memory of the systems protected by the vulnerable versions of the OpenSSL software.”
- ▶ “...this allows attackers to eavesdrop on communications, steal data directly from the services and users and to impersonate services and users.”



Heartbleed, explained




Image source: Randall Munroe, xkcd.com

Heartbleed, explained

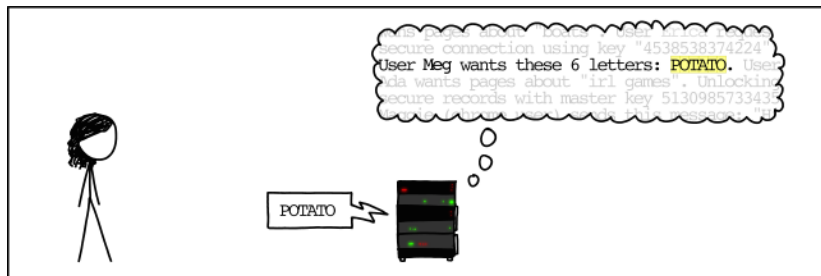


Image source: Randall Munroe, xkcd.com

Heartbleed, explained



Image source: Randall Munroe, xkcd.com

Heartbleed, explained




Image source: Randall Munroe, xkcd.com

Heartbleed, explained



Image source: Randall Munroe, xkcd.com

Heartbleed, explained

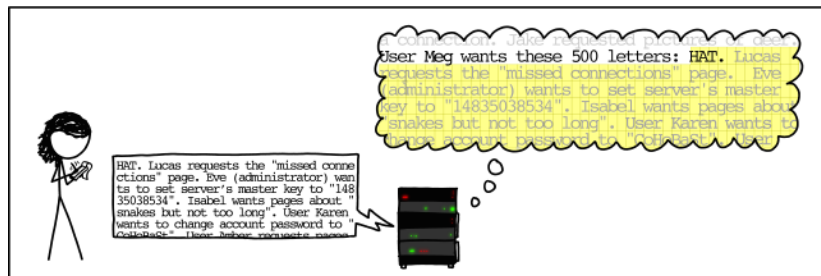


Image source: Randall Munroe, xkcd.com

Algorithms vs. code

```
1 int binarySearch(int key, int[] a, int n) {
2     int low = 0;
3     int high = n;
4
5     while (low < high) {
6         int mid = (low + high) / 2;
7
8         if(a[mid] == key) return mid; // key found
9         else if(a[mid] < key) {
10            low = mid + 1;
11        } else {
12            high = mid;
13        }
14    }
15    return -1; // key not found.
16 }
```

This is a correct binary search algorithm.

This is a correct binary search algorithm.

But what if `low + high > 231 - 1`?

This is a correct binary search algorithm.

But what if $\text{low} + \text{high} > 2^{31} - 1$?

Then $\text{mid} = (\text{low} + \text{high}) / 2$ becomes negative

This is a correct binary search algorithm.

But what if $\text{low} + \text{high} > 2^{31} - 1$?

Then $\text{mid} = (\text{low} + \text{high}) / 2$ becomes negative

- ▶ Best case: `ArrayIndexOutOfBoundsException`

This is a correct binary search algorithm.

But what if $\text{low} + \text{high} > 2^{31} - 1$?

Then $\text{mid} = (\text{low} + \text{high}) / 2$ becomes negative

- ▶ Best case: `ArrayIndexOutOfBoundsException`
- ▶ Worst case: undefined behavior

This is a correct binary search algorithm.

But what if $\text{low} + \text{high} > 2^{31} - 1$?

Then $\text{mid} = (\text{low} + \text{high}) / 2$ becomes negative

- ▶ Best case: `ArrayIndexOutOfBoundsException`
- ▶ Worst case: undefined behavior

Algorithm may be correct. But we want to run the code!

How do we fix it?

The culprit: `mid = (low + high) / 2`

How do we fix it?

The culprit: `mid = (low + high) / 2`

Need to make sure we don't overflow at any point

How do we fix it?

The culprit: `mid = (low + high) / 2`

Need to make sure we don't overflow at any point

Solution: `mid = low + (high - low)/2`

The fix

```
1 int binarySearch(int key, int[] a, int n) {
2     int low = 0;
3     int high = n;
4
5     while (low < high) {
6         int mid = low + (high - low) / 2;
7
8         if(a[mid] == key) return mid; // key found
9         else if(a[mid] < key) {
10            low = mid + 1;
11        } else {
12            high = mid;
13        }
14    }
15    return -1; // key not found.
16 }
```

The fix

```
1 int binarySearch(int key, int[] a, int n)
2 //@requires 0 <= n && n <= \length(A);
3 {
4     int low = 0;
5     int high = n;
6
7     while (low < high) {
8         int mid = low + (high - low) / 2;
9
10        if(a[mid] == key) return mid; // key found
11        else if(a[mid] < key) {
12            low = mid + 1;
13        } else {
14            high = mid;
15        }
16    }
17    return -1; // key not found.
18 }
```

The fix

```
1 int binarySearch(int key, int[] a, int n)
2 //@requires 0 <= n && n <= \length(a);
3 //@requires is_sorted(a, 0, n);
4 /*@ensures (\result == -1 && !is_in(key, A, 0, n))
5   @      // (0 <= \result, \result < n
6   @      && A[\result] == key); @*/
7 {
8     int low = 0;
9     int high = n;
10
11     while (low < high) {
12         int mid = low + (high - low) / 2;
13
14         if(a[mid] == key) return mid; // key found
15         else if(a[mid] < key) {
16             low = mid + 1;
17         } else {
18             high = mid;
19         }
20     }
21     return -1; // key not found.
22 }
```


How do we know if it's correct?

How do we know if it's correct?

One solution: test the code

How do we know if it's correct?

One solution: test the code

- ▶ Possibly incomplete \rightarrow uncertain answer
- ▶ Exhaustive testing not feasible

How do we know if it's correct?

One solution: test the code

- ▶ Possibly incomplete \rightarrow uncertain answer
- ▶ Exhaustive testing not feasible

Another: code review

How do we know if it's correct?

One solution: test the code

- ▶ Possibly incomplete → uncertain answer
- ▶ Exhaustive testing not feasible

Another: code review

- ▶ Correctness definitely important, but not the only thing
- ▶ Humans are fallable, bugs are subtle
- ▶ What's the specification?

How do we know if it's correct?

One solution: test the code

- ▶ Possibly incomplete \rightarrow uncertain answer
- ▶ Exhaustive testing not feasible

Another: code review

- ▶ Correctness definitely important, but not the only thing
- ▶ Humans are fallable, bugs are subtle
- ▶ What's the specification?

Better: prove correctness

Specification \iff *Implementation*

How do we know if it's correct?

One solution: test the code

- ▶ Possibly incomplete \rightarrow uncertain answer
- ▶ Exhaustive testing not feasible

Another: code review

- ▶ Correctness definitely important, but not the only thing
- ▶ Humans are fallable, bugs are subtle
- ▶ What's the specification?

Better: prove correctness

Specification \iff *Implementation*

- ▶ Specification must be precise
- ▶ Meaning of code must be well-defined
- ▶ Reasoning must be sound

Algorithmic Approaches

Formal proofs are tedious

We want algorithms to:

- ▶ Check our work
- ▶ Fill in low-level details
- ▶ Give diagnostic info
- ▶ Verify the system (if possible)

This is called
algorithmic verification




Image source: Daniel Kroening & Ofer Strichman,
Decision Procedures

Course objectives

- ▶ Identify and formalize program correctness
- ▶ Understand language semantics
- ▶ Apply mathematical reasoning to program correctness
- ▶ Learn how to write correct software, from beginning to end
- ▶ Use automated tools that assist verifying your code
- ▶ Understand how verification tools work

Functional Correctness

- ▶ Specification
- ▶ Proof

Specify behavior with logic

- ▶ Declarative
- ▶ Precise

Systematic proof techniques

- ▶ Derived from semantics
- ▶ Exhaustive proof rules
- ▶ Automatable*

```
1 int[] array_copy(int[] A, int n)
2 //@requires 0 <= n && n <= \length(A);
3 //@ensures \length(\result) == n;
4 {
5     int[] B = alloc_array(int, n);
6
7     for (int i = 0; i < n; i++)
8         //@loop_invariant 0 <= i;
9         {
10            B[i] = A[i];
11        }
12
13     return B;
14 }
```

Deductive verification platform

- ▶ Programming language
- ▶ Verification toolchain

Rich specification language

- ▶ Pre- and post-conditions, assertions
- ▶ Pure mathematical functions
- ▶ Termination metrics

Programmer writes specification, partial annotations

Compiler proves correctness automatically!

Systems that prove that programs match their specifications

Basic idea:

1. Translate programs into *proof obligations*
2. Encode proof obligations as satisfiability
3. Solve using a decision procedure

Systems that prove that programs match their specifications

Problem is undecidable!

1. Require annotations
2. Relieve manual burden by inferring some annotations

Basic idea:

1. Translate programs into *proof obligations*
2. Encode proof obligations as satisfiability
3. Solve using a decision procedure

Systems that prove that programs match their specifications

Problem is undecidable!

1. Require annotations
2. Relieve manual burden by inferring some annotations


Verifiers are complex systems

Basic idea:

1. Translate programs into *proof obligations*
2. Encode proof obligations as satisfiability
3. Solve using a decision procedure


Fully-automatic techniques for finding bugs (or proving their absence)

- ▶ Specifications written in *propositional temporal logic*
- ▶ Verification by exhaustive state space search
- ▶ Diagnostic counterexamples




Fully-automatic techniques for finding bugs (or proving their absence)

- ▶ Specifications written in *propositional temporal logic*
- ▶ Verification by exhaustive state space search
- ▶ Diagnostic counterexamples
- ▶ No proofs!



Fully-automatic techniques for finding bugs (or proving their absence)

- ▶ Specifications written in *propositional temporal logic*
- ▶ Verification by exhaustive state space search
- ▶ Diagnostic counterexamples
- ▶ No proofs!
- ▶ **Downside:** “State explosion”



Model Checking



Fully-automatic techniques for finding bugs (or proving their absence)

- ▶ Specifications written in *propositional temporal logic*
- ▶ Verification by exhaustive state space search
- ▶ Diagnostic counterexamples
- ▶ No proofs!
- ▶ **Downside:** “State explosion”

10^{70} atoms



10^{500000} states



Clever ways of dealing with state explosion:

Clever ways of dealing with state explosion:

- ▶ Partial order reduction
- ▶ Bounded model checking
- ▶ Symbolic representations
- ▶ Abstraction & refinement

Clever ways of dealing with state explosion:

- ▶ Partial order reduction
- ▶ Bounded model checking
- ▶ Symbolic representations
- ▶ Abstraction & refinement

Now widely used for verification & bug-finding:

- ▶ Hardware, software, protocols, ...
- ▶ Microsoft, Intel, Amazon, Google, NASA, ...

Clever ways of dealing with state explosion:

- ▶ Partial order reduction
- ▶ Bounded model checking
- ▶ Symbolic representations
- ▶ Abstraction & refinement

Now widely used for verification & bug-finding:

- ▶ Hardware, software, protocols, ...
- ▶ Microsoft, Intel, Amazon, Google, NASA, ...



Ed Clarke
Turing Award,
2007

Breakdown:

- ▶ 40% labs
- ▶ 25% written homework
- ▶ 30% exams (15% each, midterm and final)
- ▶ 5% participation

5 labs

Weekly written homework

In-class exams, closed-book

Participation:

- ▶ Come to lecture
- ▶ Ask questions, give answers
- ▶ Contribute to discussion

For the labs, you will:

- ▶ Implement some functionality
- ▶ Specify correctness for that functionality
- ▶ Use Why3 to prove it correct

For the labs, you will:

- ▶ Implement some functionality
- ▶ Specify correctness for that functionality
- ▶ Use Why3 to prove it correct

Most important criterion is **correctness**.

For the labs, you will:

- ▶ Implement some functionality
- ▶ Specify correctness for that functionality
- ▶ Use Why3 to prove it correct

Most important criterion is **correctness**.

Full points when you provide the following

- ▶ Correct implementation
- ▶ Correct specification
- ▶ Correct annotations
- ▶ Sufficient annotations for verification

For the labs, you will:

- ▶ Implement some functionality
- ▶ Specify correctness for that functionality
- ▶ Use Why3 to prove it correct

Most important criterion is **correctness**.

Full points when you provide the following

- ▶ Correct implementation
- ▶ Correct specification
- ▶ Correct annotations
- ▶ Sufficient annotations for verification

Partial credit depending on how many of these you achieve

Clarity & conciseness is necessary for partial credit!

Labs are intended to build proficiency in:

- ▶ Writing good specifications
- ▶ Applying formal proof to practice
- ▶ Making effective use of automated tools
- ▶ **Writing useful & correct code**

Labs are intended to build proficiency in:

- ▶ Writing good specifications
- ▶ Applying formal proof to practice
- ▶ Making effective use of automated tools
- ▶ **Writing useful & correct code**

Gradual progression to sophistication:

1. Familiarize yourself with Why3
2. Implement and prove something
3. Work with more complex data structures
4. Implement and prove something really interesting
5. Optimize your implementation, still verified

New this year: 1st Annual Verified SAT Competition

New this year: **1st Annual Verified SAT Competition**

Think your final lab is a cut above the rest?

1. Submit it to the competition
2. Come to the party (food & entertainment provided)
3. Run against a set of benchmarks from real applications
4. To win: correct answer on all benchmarks + fastest execution

New this year: **1st Annual Verified SAT Competition**

Think your final lab is a cut above the rest?

1. Submit it to the competition
2. Come to the party (food & entertainment provided)
3. Run against a set of benchmarks from real applications
4. To win: correct answer on all benchmarks + fastest execution

Winners receive **fame and glory**

New this year: **1st Annual Verified SAT Competition**

Think your final lab is a cut above the rest?

1. Submit it to the competition
2. Come to the party (food & entertainment provided)
3. Run against a set of benchmarks from real applications
4. To win: correct answer on all benchmarks + fastest execution

Winners receive **fame and glory**

- ▶ and also prizes

Written homeworks focus on theory and fundamental skills

Written homeworks focus on theory and fundamental skills

Grades are based on:

- ▶ Correctness of your answer
- ▶ How you present your reasoning

Written homeworks focus on theory and fundamental skills

Grades are based on:

- ▶ Correctness of your answer
- ▶ How you present your reasoning

Strive for **clarity & conciseness**

- ▶ Show each step of your reasoning
- ▶ State your assumptions
- ▶ Answers without these → no points

Late Policy

No late days on written homework

- ▶ Not intended to be time-intensive
- ▶ 25% deduction for each day past deadline

No late days on written homework

- ▶ Not intended to be time-intensive
- ▶ 25% deduction for each day past deadline

Can earn back missed points for proofs on labs

- ▶ Must submit original lab by the deadline
- ▶ Resubmit **once** within three days of deadline
- ▶ If proof is complete & correct, earn back points *only on the proof*

Website: <http://www.cs.cmu.edu/~15414>

Course staff contact: Piazza or
15414-staff@lists.andrew.cmu.edu

Lecture: Tuesdays & Thursdays, 10:30-11:50 GHC 4211

Matt Fredrikson, Ruben Martins

- ▶ Location: CIC 2126, GHC 9103
- ▶ Office Hours: TBD
- ▶ Email: mfredrik@cs, rubenm@cs

Krishna Bagadia, Rameel Rizvi

- ▶ Office Hours: TBD