

Program Analysis and Verification

0368-4479

Noam Rinetzky

Lecture 1: Introduction & Overview

Slides credit: Tom Ball, Dawson Engler, Roman Manevich, Erik Poll,
Mooly Sagiv, Jean Souyris, Eran Tromer, Avishai Wool, Eran Yahav

Admin



- Lecturer: Noam Rinetzky
 - *maon@cs.tau.ac.il*
 - <http://www.cs.tau.ac.il/~maon>
- 14 Lessons
 - Monday, 13:00-16:00, Shenkar-Physics 222
- 4 Assignments (30%)
 - 1 involves programming
- 1 Lesson summary (10%)
- Final exam (60%)
 - Must pass

Today

- Motivation
- Introduction

- Not technical

Software is Everywhere



Unreliable Software is Everywhere

Windows

A fatal exception 0E has occurred at 0020:C0011E36 in UXD UMM(01) + 00010E36. The current application will be terminated.

- * Press any key to terminate the current application.
- * Press CTRL+ALT+DEL again to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue

30GB Zunes all over the world fail en masse



December 31, 2008

Zune bug

```
1 while (days > 365) {
2   if (IsLeapYear(year)) {
3     if (days > 366) {
4       days -= 366;
5       year += 1;
6     }
7   } else {
8     days -= 365;
9     year += 1;
10  }
11 }
```



December 31, 2008

Zune bug

```
1 while (366 > 365) {  
2   if (IsLeapYear(2008)) {  
3     if (366 > 366) {  
4       days -= 366;  
5       year += 1;  
6     }  
7   } else {  
8     days -= 365;  
9     year += 1;  
10  }  
11 }
```



December 31, 2008

Suggested solution: wait for tomorrow

Patriot missile failure

On the night of the 25th of February, 1991, a Patriot missile system operating in Dhahran, Saudi Arabia, failed to track and intercept an incoming Scud. The Iraqi missile impacted into an army barracks, killing 28 U.S. soldiers and injuring another 98.



February 25, 1991

Patriot bug – rounding error

- Time measured in 1/10 seconds
- Binary expansion of 1/10:
0.0001100110011001100110011001100....
- 24-bit register
0.00011001100110011001100
- error of
 - 0.00000000000000000000000000000011001100... binary, or
~0.000000095 decimal
- After 100 hours of operation error is
 $0.000000095 \times 100 \times 3600 \times 10 = 0.34$
- A Scud travels at about 1,676 meters per second, and so travels more than half a kilometer in this time

Suggested solution: reboot every 10 hours

Toyota recalls 160,000 Prius hybrid vehicles



Programming error can activate all warning lights, causing the car to think its engine has failed

October 2005

Therac-25 leads to 3 deaths and 3 injuries



Software error exposes patients to radiation overdose (100X of intended dose)

1985 to 1987

Northeast Blackout



14 August, 2003

Unreliable Software is **Exploitable**

The Sony PlayStation Network breach: An identity-theft bonanza

Stuxnet Worm Still Out of Control at Iran's Nuclear Sites, Experts Say

The Stuxnet worm, named after initials found in its code, is **the most sophisticated cyberweapon** ever.

Security Advisory for Adobe Flash Player, Adobe Reader and Acrobat

This vulnerability could cause a crash and potentially **take control of the affected system**. There are reports that this vulnerability is being exploited in the (.swf) file embedded in a Microsoft Office document.

RSA hacked, information leaks

RSA's corporate network suffered what RSA describes as a successful advanced persistent threat attack, and "certain information" was stolen. The information affected an email account of SecurID authentication (May 2011)

RSA tokens may be behind major network security problems at Lockheed Martin

Lockheed Martin remote access network, protected by SecurID tokens, has been shut down (May 2011)

*Billy Gates why do you make this possible ? Stop making money
and fix your software!!*

(W32.Blaster.Worm)


August 13, 2003

Windows exploit(s)

Buffer Overflow

```
void foo (char *x) {  
    char buf[2];  
    strcpy(buf, x);  
}  
int main (int argc, char *argv[]) {  
    foo(argv[1]);  
}
```

```
./a.out abracadabra  
Segmentation fault
```



Buffer overrun exploits

```
int check_authentication(char *password) {
    int auth_flag = 0;
    char password_buffer[16];


    strcpy(password_buffer, password);
    if(strcmp(password_buffer, "brillig") == 0) auth_flag = 1;
    if(strcmp(password_buffer, "outgrabe") == 0) auth_flag = 1;
    return auth_flag;
}
int main(int argc, char *argv[]) {
    if(check_authentication(argv[1])) {
        printf("\n-----\n");
        printf("    Access Granted.\n");
        printf("-----\n"); }
    else
        printf("\nAccess Denied.\n");
}
```

Input Validation



1234567890123456

evil input
→



Access Granted.

Boeing's 787 Vulnerable to Hacker Attack



security vulnerability in onboard computer networks could allow passengers to access the plane's control systems

January 2008

What can we do about it?

What can we do about it?

I just want to say LOVE YOU SAN!!soo much

Billy Gates why do you make this possible ? Stop making money and fix your software!!


(W32.Blaster.Worm / Lovesan worm)

August 13, 2003

What can we do about it?

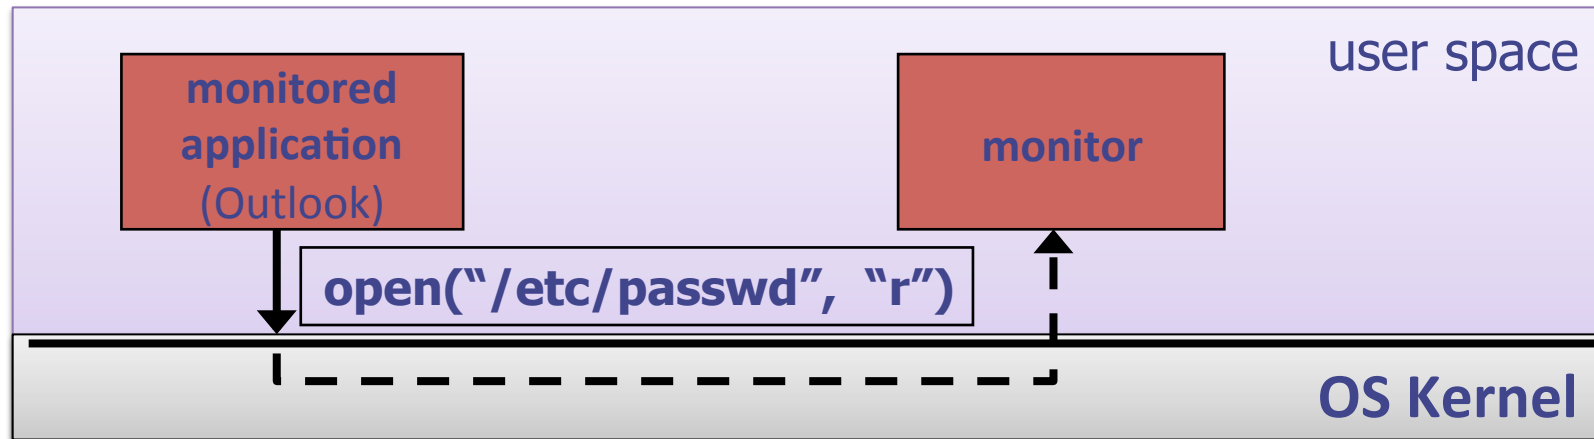
- Monitoring
- Testing
- Static Analysis
- Formal Verification
- Specification

Run time



Design Time


Monitoring (e.g., for security)



- StackGuard
- ProPolice
- PointGuard
- Security monitors (ptrace)

Testing

- build it; try it on a some inputs



- printf ("x == 0 => should not get that!")

Testing

- **Valgrind** memory errors, race conditions, taint analysis
 - Simulated CPU
 - Shadow memory

Invalid read of size 4

at 0x40F6BBCC: (within /usr/lib/libpng.so.2.1.0.9)

by 0x40F6B804: (within /usr/lib/libpng.so.2.1.0.9)

by 0x40B07FF4: read_png_image(QImageIO *) (kernel/qpngio.cpp:326)

by 0x40AC751B: QImageIO::read() (kernel/qimage.cpp:3621)

Address 0xBFFFF0E0 is not stack'd, malloc'd or free'd

Testing

- **Valgrind** memory errors, race conditions
- **Parasoft Jtest/Insure++** memory errors + visualizer, race conditions, exceptions ...
- **IBM Rational Purify** memory errors
- **IBM PureCoverage** detect untested paths
- **Daikon** dynamic invariant detection

Testing

- Useful and challenging
 - Random inputs
 - Guided testing (coverage)
 - Bug reproducing
- But ...
 - Observe **some** program behaviors
 - What can you say about **other** behaviors?


Testing is not enough

- Observe **some** program behaviors
- What can you say about **other** behaviors?
- Concurrency makes things worse
- Smart testing is useful
 - requires the techniques that we will see in the course

What can we do about it?

- Monitoring
- Testing
- Static Analysis
- Formal Verification
- Specification

Run time



Design Time

Program Analysis & Verification

```
x = ?  
if (x > 0) {  
    y = 42;  
} else {  
    y = 73;  
    foo();  
}  
assert (y == 42);
```



- Is assertion true?

Program Analysis & Verification

```
y = ?; x = y * 2
if (x % 2 == 0) {
    y = 42;
} else {
    y = 73;
    foo();
}
assert (y == 42);
```



- Is assertion true? Can we prove this? Automatically?
- Bad news: problem is generally undecidable

Formal verification

- Mathematical model of software
 - $\rho: \text{Var} \rightarrow \mathbb{Z}$
 - $\rho = [x \mapsto 0, y \mapsto 1]$
- Logical specification
 - $\{0 < x\} = \{\rho \in \text{State} \mid 0 < \rho(x)\}$
- Machine checked formal proofs

$$\{0 < x \wedge y = x\} \rightarrow \{0 < y\}$$

$$\{0 < x\} \quad y := x \quad \{0 < x \wedge y = x\} \quad \{0 < y\} \quad y := y + 1 \quad \{1 < y\}$$

$$\{ \quad ? \quad \} y := x ; y := y + 1 \quad \{1 < y\}$$

Formal verification

- Mathematical model of software
 - State = Var \rightarrow Integer
 - $S = [x \rightarrow 0, y \rightarrow 1]$
- Logical specification
 - $\{ 0 < x \} = \{ S \in \text{State} \mid 0 < S(x) \}$
- Machine checked formal proofs

$$\{ Q' \} \rightarrow \{ P' \}$$

$$\{ P \} \text{ stmt1 } \{ Q' \} \quad \{ P' \} \text{ stmt2 } \{ Q \}$$


$$\{ P \} \text{ stmt1; stmt2 } \{ Q \}$$

Program Verification

```
{true}
y = ?; x = 2 * y;
{x = 2 * y}
if (x % 2 == 0) {
  {x = 2 * y}
  y = 42;
  { ∃z. x = 2 * z ∧ y = 42 }
} else {
  {false }
  y = 73;
  foo();
  {false }
}
{ ∃z. x = 2 * z ∧ y = 42 }
assert (y == 42);
```

- Is assertion true? Can we prove this? Automatically?
- Can we prove this manually?

Central idea: use approximation



Program Verification

```
{true}
y = ?; x = 2 * y;
{x = 2 * y}
if (x % 2 == 0) {
  {x = 2 * y}
  y = 42;
  { ∃z. x = 2 * z ∧ y = 42 }
} else {
  {false }
  y = 73;
  foo();
  {false }
}
{ ∃z. x = 2 * z ∧ y = 42 } {x = ? ∧ y = 42 } {x = 4 ∧ y = 42 }
assert (y == 42);
```

- Is assertion true? Can we prove this? Automatically?
- Can we prove this manually?

L4.verified [Klein⁺, '09]

- Microkernel
 - IPC, Threads, Scheduling, Memory management
- Functional correctness (using Isabelle/HOL)
 - + No null pointer de-references.
 - + No memory leaks.
 - + No buffer overflows.
 - + No unchecked user arguments
 - + ...
- Kernel/proof co-design
 - Implementation - 2.5 py (8,700 LOC)
 - Proof – 20 py (200,000 LOP)

Static Analysis

- Lightweight formal verification
- Formalize software behavior in a mathematical model (semantics)
- Prove (selected) properties of the mathematical model
 - Automatically, typically with approximation of the formal semantics

Why static analysis?


- Some errors are hard to find by testing
 - arise in unusual circumstances/uncommon execution paths
 - buffer overruns, unvalidated input, exceptions, ...
 - involve non-determinism
 - race conditions
- Full-blown formal verification too expensive

Is it at all doable?

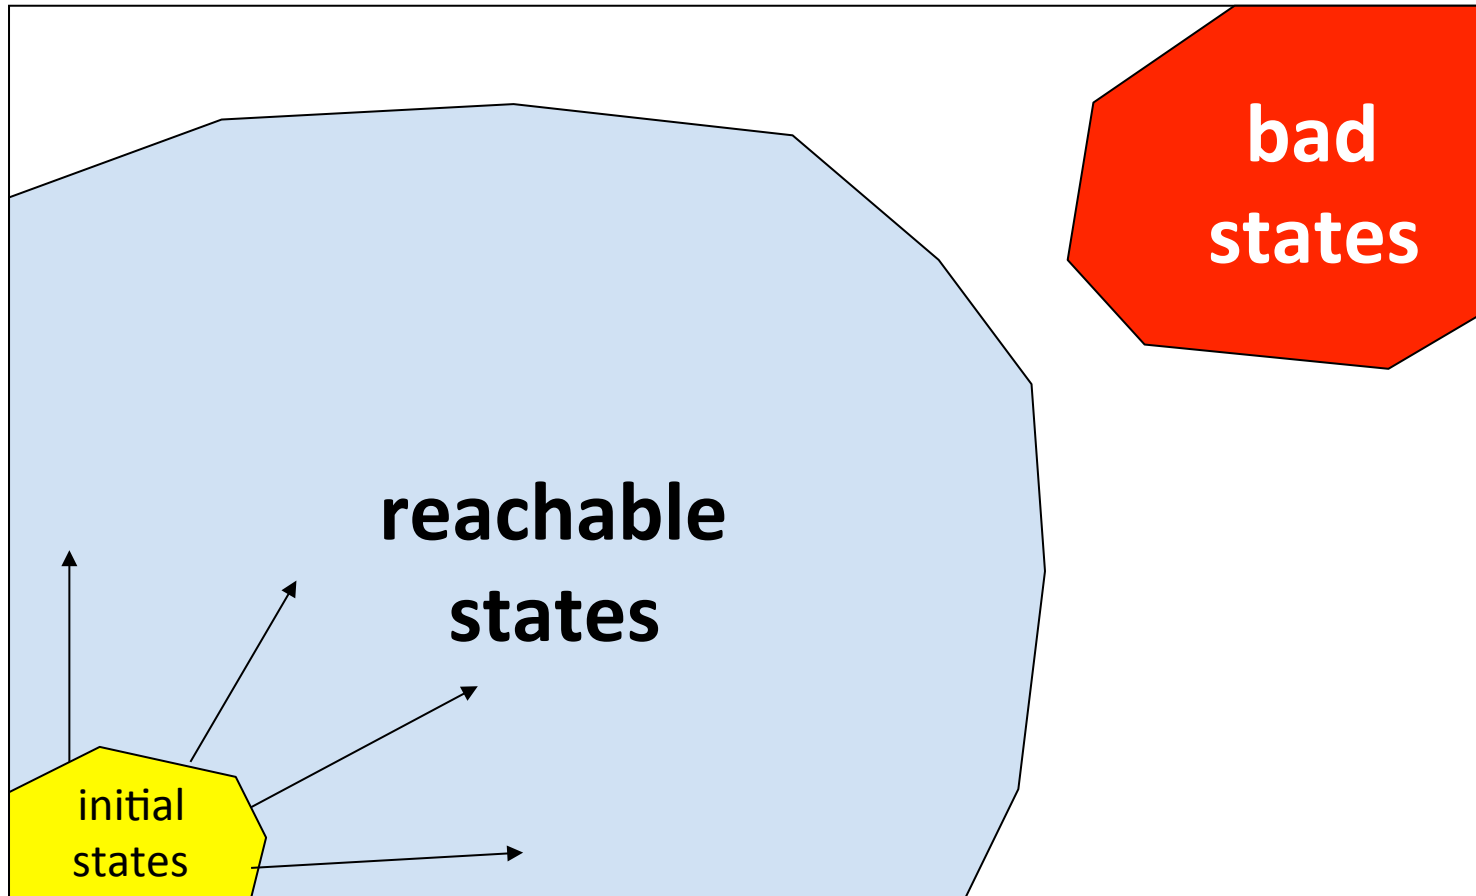
```
x = ?  
if (x > 0) {  
    y = 42;  
} else {  
    y = 73;  
    foo();  
}  
assert (y == 42);
```

Bad news: problem is generally undecidable


Central idea: use approximation




Goal: exploring program states




Technique: explore abstract states




Technique: explore abstract states




Technique: explore abstract states




Technique: explore abstract states



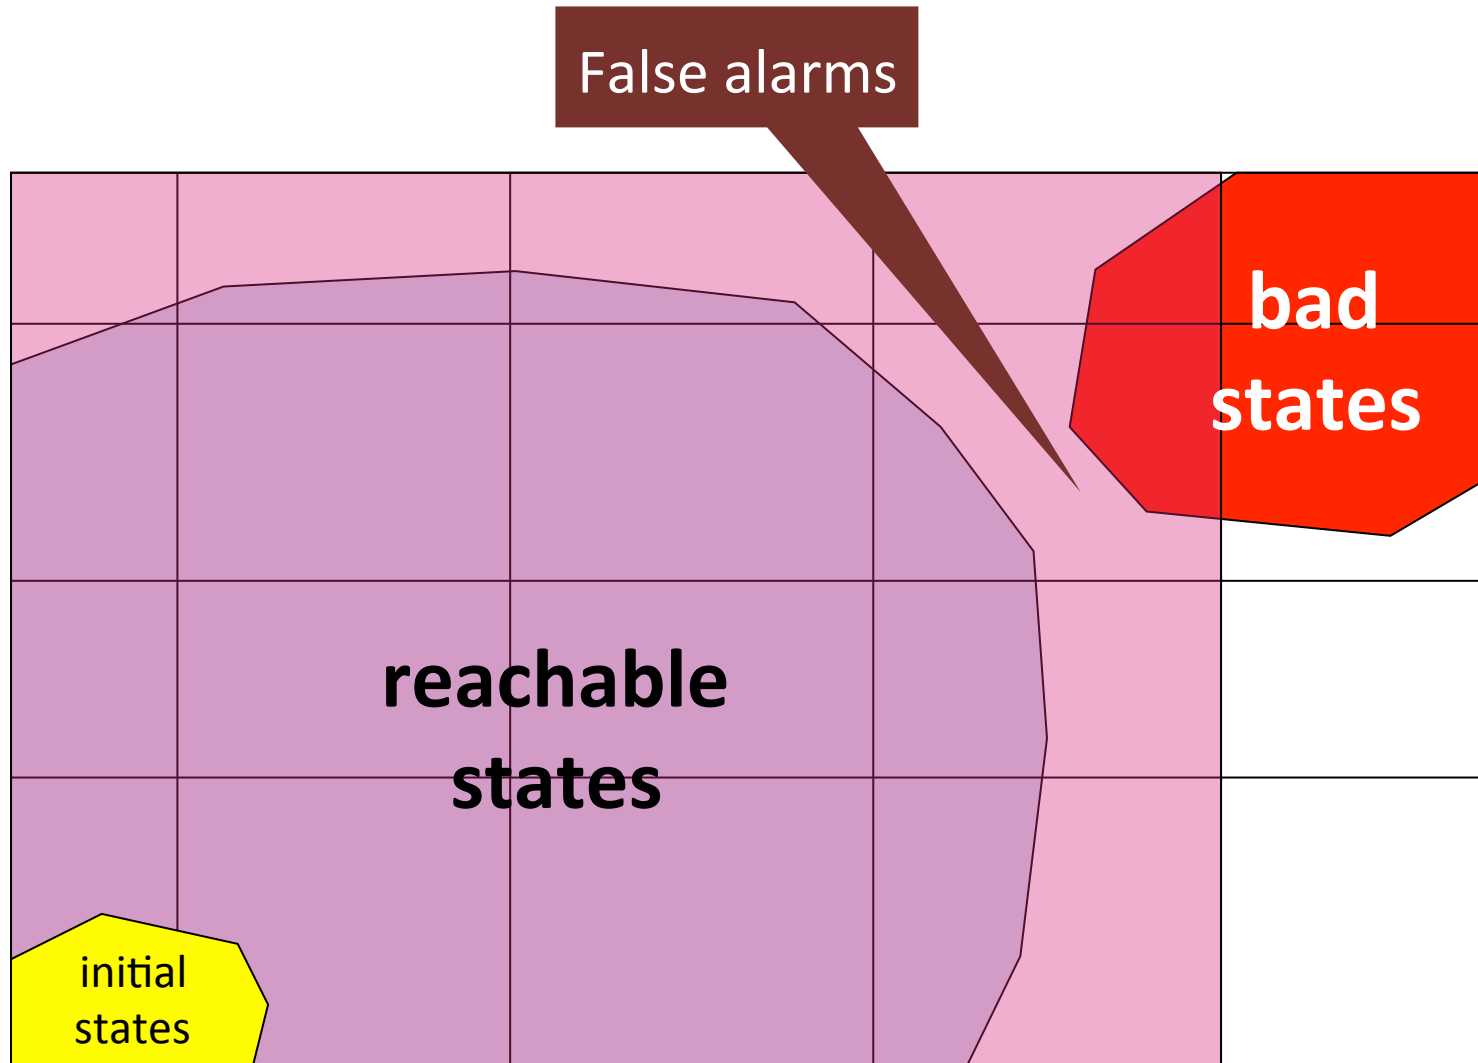
Sound: cover all reachable states



Unsound: miss some reachable states



Imprecise abstraction



A sound message

```
x = ?  
if (x > 0) {  
    y = 42;  
} else {  
    y = 73;  
    foo();  
}  
assert (y == 42); Assertion may be violated
```

Precision

- Avoid useless result

```
UselessAnalysis(Program p) {  
    printf("assertion may be violated\n");  
}
```

- Low false alarm rate
- Understand where precision is lost

A sound message


```
y = ?; x = y * 2
if (x % 2 == 0) {
    y = 42;
} else {
    y = 73;
    foo();
}
assert (y == 42);
```

Assertion is true


How to find “the right” abstraction?

- Pick an **abstract domain** suited for your property
 - Numerical domains
 - Domains for reasoning about the heap
 - ...
- Combination of abstract domains

Intervals Abstraction




Interval Lattice



Example


```
int x = 0;  
if (?) x++;  
if (?) x++;
```



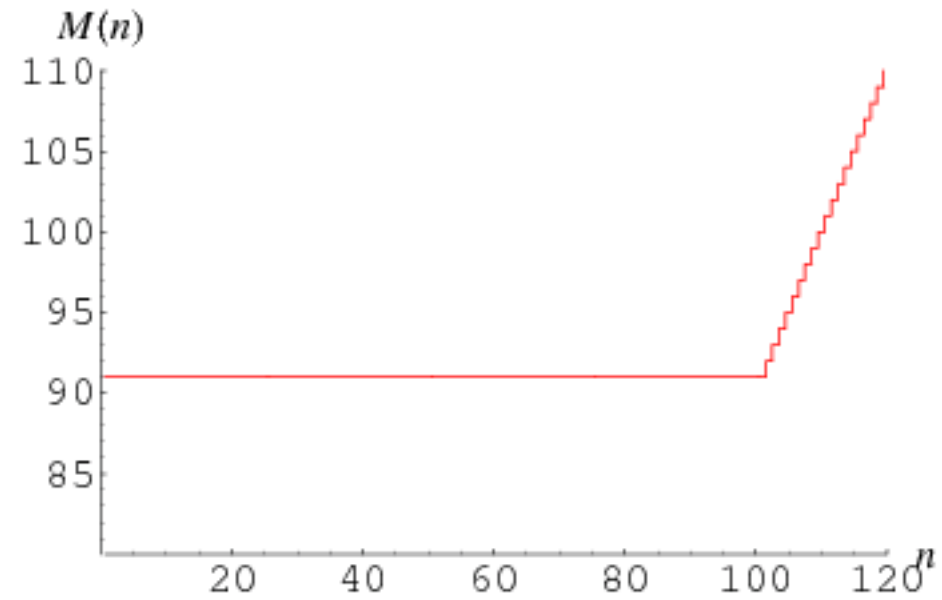
$$[a1,a2] \sqcup [b1,b2] = [\min(a1,b1), \max(a2,b2)]$$

Polyhedral Abstraction

- abstract state is an intersection of linear inequalities of the form $a_1x_1 + a_2x_2 + \dots + a_nx_n \leq c$
- represent a set of points by their convex hull



McCarthy 91 function



$$M(n) = \begin{cases} M(M(n+11)) & \text{for } n \leq 100 \\ n - 10 & \text{for } n > 100. \end{cases}$$

McCarthy 91 function

```
proc MC (n : int) returns (r : int) var t1 : int, t2 : int;
begin

    if n > 100 then

        r = n - 10;
    else

        t1 = n + 11;
        t2 = MC(t1);

        r = MC(t2);

    endif;
end

var a : int, b : int;
begin /* top */
    b = MC(a);
end
```

if (n>=101) then n-10 else 91

McCarthy 91 function

```
proc MC : int;
begin
    if (n >= 101) then n-10 else 91
        /* top */
    if n > 100 then
        /* [|n-101>=0|] */
        r = n - 10; /* [| -n+r+10=0; n-101>=0|] */
    else
        /* [| -n+100>=0|] */
        t1 = n + 11; /* [| -n+t1-11=0; -n+100>=0|] */
        t2 = MC(t1); /* [| -n+t1-11=0; -n+100>=0;
            -n+t2-1>=0; t2-91>=0|] */
        r = MC(t2); /* [| -n+t1-11=0; -n+100>=0;
            -n+t2-1>=0; t2-91>=0; r-t2+10>=0;
            r-91>=0|] */
    endif; /* [| -n+r+10>=0; r-91>=0|] */
end

var a : int, b : int;
begin /* top */
    b = MC(a); /* [| -a+b+10>=0; b-91>=0|] */
end
```

if (n >= 101) then n-10 else 91

McCarthy 91 function

```
proc MC (n : int) returns (r : int) var t1 : int, t2 : int;
begin
    /* (L6 C5) top */
    if n > 100 then
        /* (L7 C17) [|n-101>=0|] */
        r = n - 10; /* (L8 C14) [| -n+r+10=0; n-101>=0|] */
    else
        /* (L9 C6) [| -n+100>=0|] */
        t1 = n + 11; /* (L10 C17) [| -n+t1-11=0; -n+100>=0|] */
        t2 = MC(t1); /* (L11 C17) [| -n+t1-11=0; -n+100>=0;
            -n+t2-1>=0; t2-91>=0|] */
        r = MC(t2); /* (L12 C16) [| -n+t1-11=0; -n+100>=0;
            -n+t2-1>=0; t2-91>=0; r-t2+10>=0;
            r-91>=0|] */
    endif; /* (L13 C8) [| -n+r+10>=0; r-91>=0|] */
end

var a : int, b : int;
begin
    /* (L18 C5) top */
    b = MC(a); /* (L19 C12) [| -a+b+10>=0; b-91>=0|] */
end
```

What is Static analysis

- Develop **theory** and **tools** for program correctness and robustness
- Reason **statically** (at compile time) about the possible runtime behaviors of a program

“The **algorithmic discovery** of **properties** of a program by inspection of its source text¹”

-- Manna, Pnueli

¹ Does not have to literally be the source text, just means w/o running it

Static analysis definition

Reason **statically** (at compile time) about the possible runtime behaviors of a program

“The **algorithmic discovery** of **properties** of a program by inspection of its source text¹”
-- Manna, Pnueli

¹ Does not have to literally be the source text, just means w/o running it

Some automatic tools

Challenges


```
class SocketHolder { Socket s; }
Socket makeSocket() { return new Socket(); // A }
open(Socket l) { l.connect(); }
talk(Socket s) { s.getOutputStream().write("hello"); }

main() {
    Set<SocketHolder> set = new HashSet<SocketHolder>();
    while(...) {
        SocketHolder h = new SocketHolder();
        h.s = makeSocket();
        set.add(h);
    }
    for (Iterator<SocketHolder> it = set.iterator(); ...) {
        Socket g = it.next().s;
        open(g);
        talk(g);
    }
}
```


(In)correct usage of APIs

- **Application trend: Increasing number of libraries and APIs**
 - Non-trivial restrictions on permitted sequences of operations
- **Typestate:** Temporal safety properties
 - What sequence of operations are permitted on an object?
 - Encoded as DFA

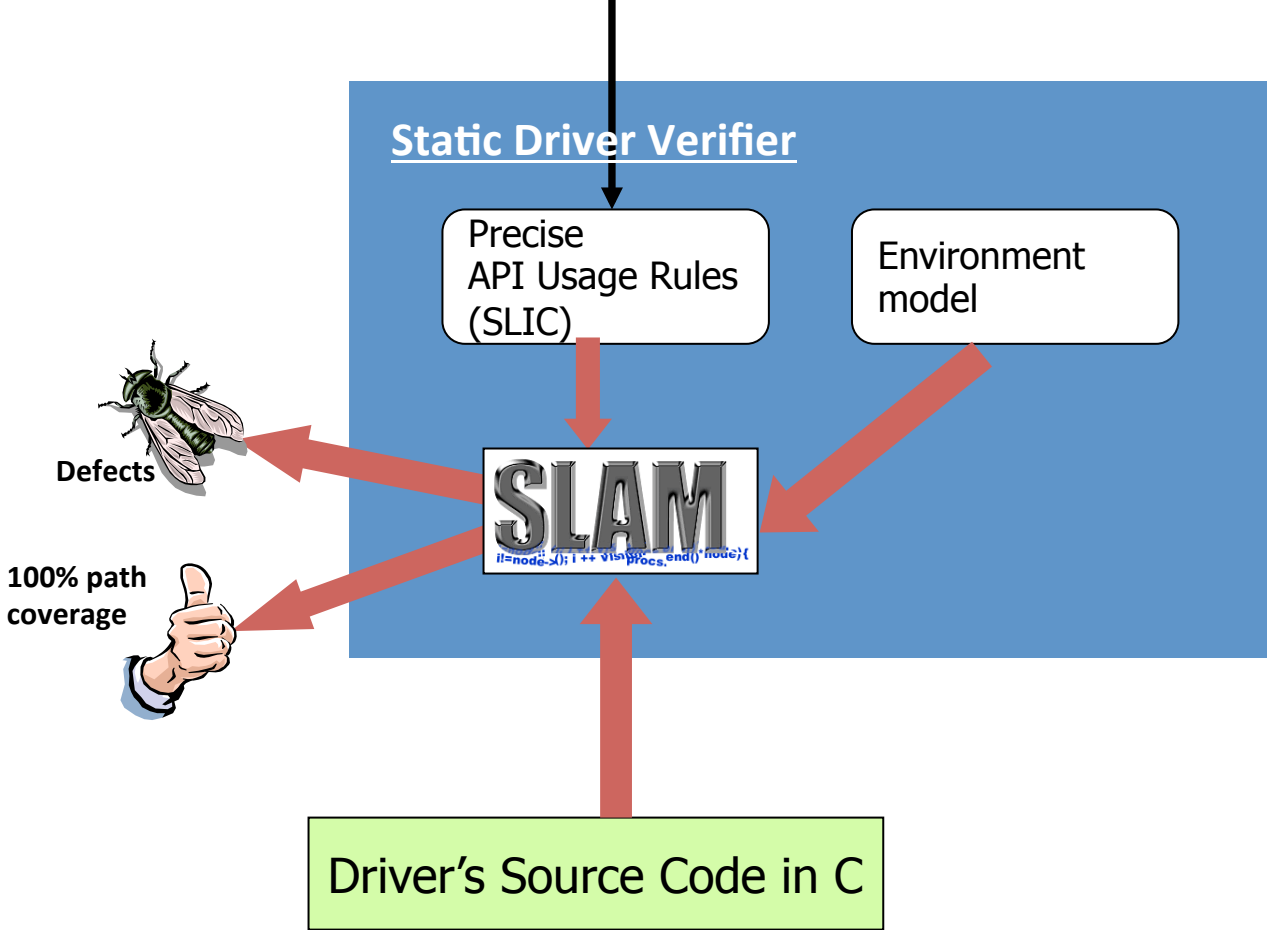
e.g. “Don’t use a Socket unless it is connected”



Static Driver Verifier




Rules



SLAM

State machine for locking



Locking rule in SLIC

```
state {  
    enum {Locked,Unlocked}  
    s = Unlocked;  
}  
  
KeAcquireSpinLock.entry {  
    if (s==Locked) abort;  
    else s = Locked;  
}  
  
KeReleaseSpinLock.entry {  
    if (s==Unlocked) abort;  
    else s = Unlocked;  
}
```

SLAM (now SDV) [Ball⁺, '11]

- 100 drivers and 80 SLIC rules.
 - The largest driver ~ 30,000 LOC
 - Total size ~450,000 LOC
- The total runtime for the 8,000 runs (driver x rule)
 - 30 hours on an 8-core machine
 - 20 mins. Timeout
- Useful results (bug / pass) on over 97% of the runs
- Caveats: pointers (imprecise) & concurrency (ignores)

The **Astrée** Static Analyzer

Patrick Cousot
Radhia Cousot
Jérôme Feret
Laurent Mauborgne
Antoine Miné
Xavier Rival


ENS France

Objectives of **Astrée**

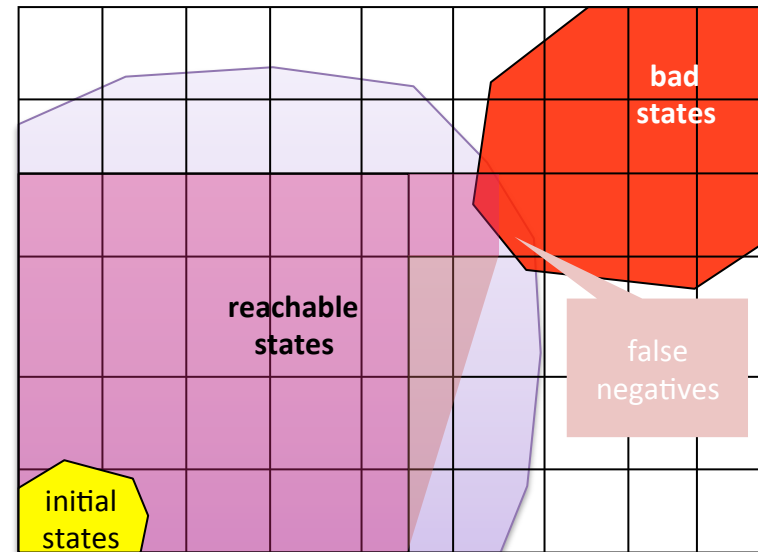
- Prove absence of errors in safety critical C code
- ASTRÉE was able to prove completely automatically the absence of any RTE in the primary flight control software of the Airbus A340 fly-by-wire system
 - a program of 132,000 lines of C analyzed



Scaling




Sound



Complete

Unsound static analysis



Unsound static analysis

- Static analysis
 - No code execution
- Trade soundness for scalability
 - Do not cover all execution paths
 - But cover “many”

FindBugs [Pugh+, '04]

- Analyze Java programs (bytecode)
- Looks for “bug patterns”
- Bug patterns
 - Method() vs method()
 - Override equal(...) but not hashCode()
 - Unchecked return values
 - Null pointer dereference

App17	KLOC	NP bugs	Other Bugs	Bad Practice	Dodgy
Sun JDK 1.7	597	68	180	594	654
Eclipse 3.3	1447	146	259	1079	653
Netbeans 6	1022	189	305	3010	1112
glassfish	2176	146	154	964	1222
jboss	178	30	57	263	214

PREfix [Pincus+, '00]

- Developed by Pinucs, purchased by Microsoft
- Automatic analysis of C/C++ code
 - Memory errors, divide by zero
 - Inter-procedural bottom-up analysis
 - Heuristic - choose “100” paths
 - Minimize effect of false positive

Program	KLOC	Time
Mozilla browser	540	11h
Apache	49	15m

2-5 warnings per KLOC

PREfast

- Analyze Microsoft kernel code + device drivers
 - Memory errors, races,
- Part of Microsoft visual studio
- Intra-procedural analysis
- User annotations

```
memcpy( __out_bcount( length ) dest, __in_bcount( length ) src, length );
```

PREfix + PREfast found 1/6 of bugs fixed in Windows Server'03

Coverity [Engler+, '04]

- Looks for bug patterns
 - Enable/disable interrupts, double locking, double locking, buffer overflow, ...
- Learns patterns from common
- Robust & scalable
 - 150 open source program -6,000 bugs
 - Unintended acceleration in Toyota

Sound SA vs. Testing

Sound SA

- Can find rare errors
Can raise false alarms
- Cost ~ program's complexity
- Can handle limited classes of programs and still be useful

Unsound SA

- Can miss errors
Can raise false alarms
- Cost ~ program's complexity
- No need to efficiently handle rare cases

Testing

- Can miss errors
Finds real errors
- Cost ~ program's execution
- No need to efficiently handle rare cases

Sound SA vs. Formal verification

Sound Static Analysis

- Fully automatic
- Applicable to a programming language
- Can be very imprecise
- May yield false alarms

Formal verification

- Requires specification and loop invariants
- Program specific
- Relatively complete
- Provides counter examples
- Provides useful documentation
- Can be mechanized using theorem provers

The End