

# Statically Detecting Likely Buffer Overflow Vulnerabilities

David Larochelle and David Evans  
USENIX'01

David Larochelle and David Evans  
IEEE Software Jan/Feb 2002

Presented by Adam Polyak  
30.03.2014

# Outline of talk

- Introduction
- Suggested Solution: Splint
- Evaluation
- Related work
- Conclusions

# Introduction

```
1  int B=0;  
2  char A[8]={};  
3  strcpy(A, "excessive");
```

- Before command:

Var name	A								B	
Value	[empty]								0	
Hex Value	00	00	00	00	00	00	00	00	00	00

- After command:

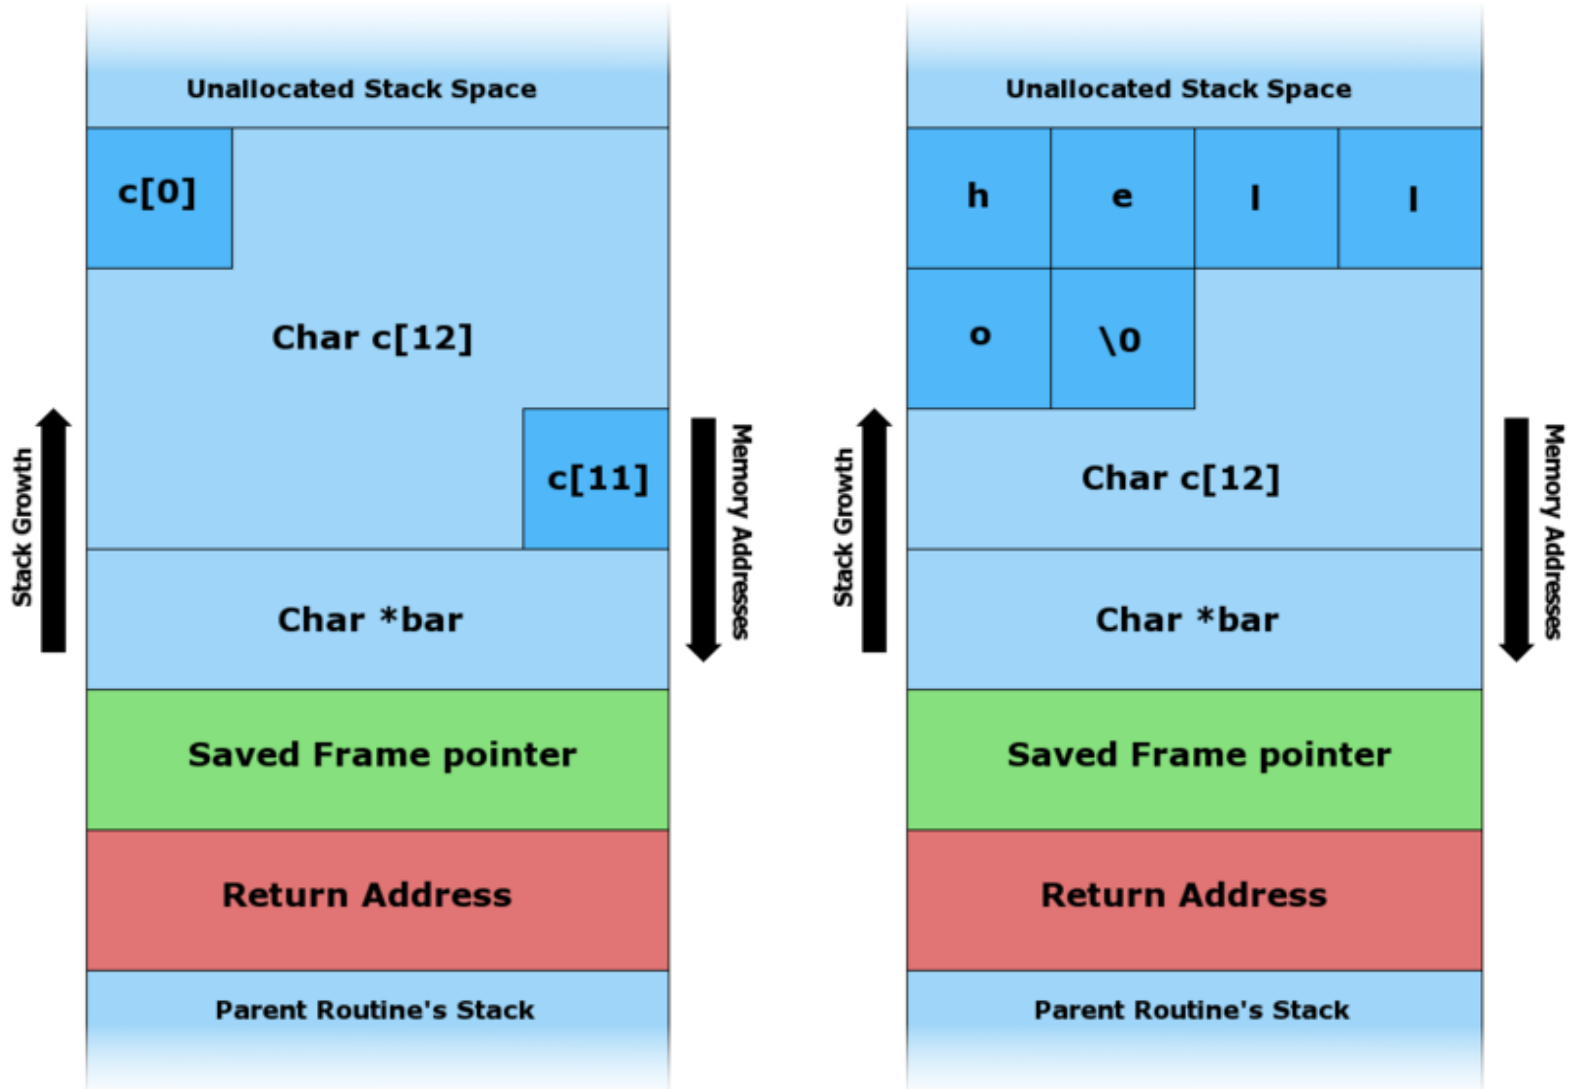
Var name	A								B	
Value	'e'	'x'	'c'	'e'	's'	's'	'i'	'v'	25856	
Hex Value	65	78	63	65	73	73	69	76	65	00

# It can get worse

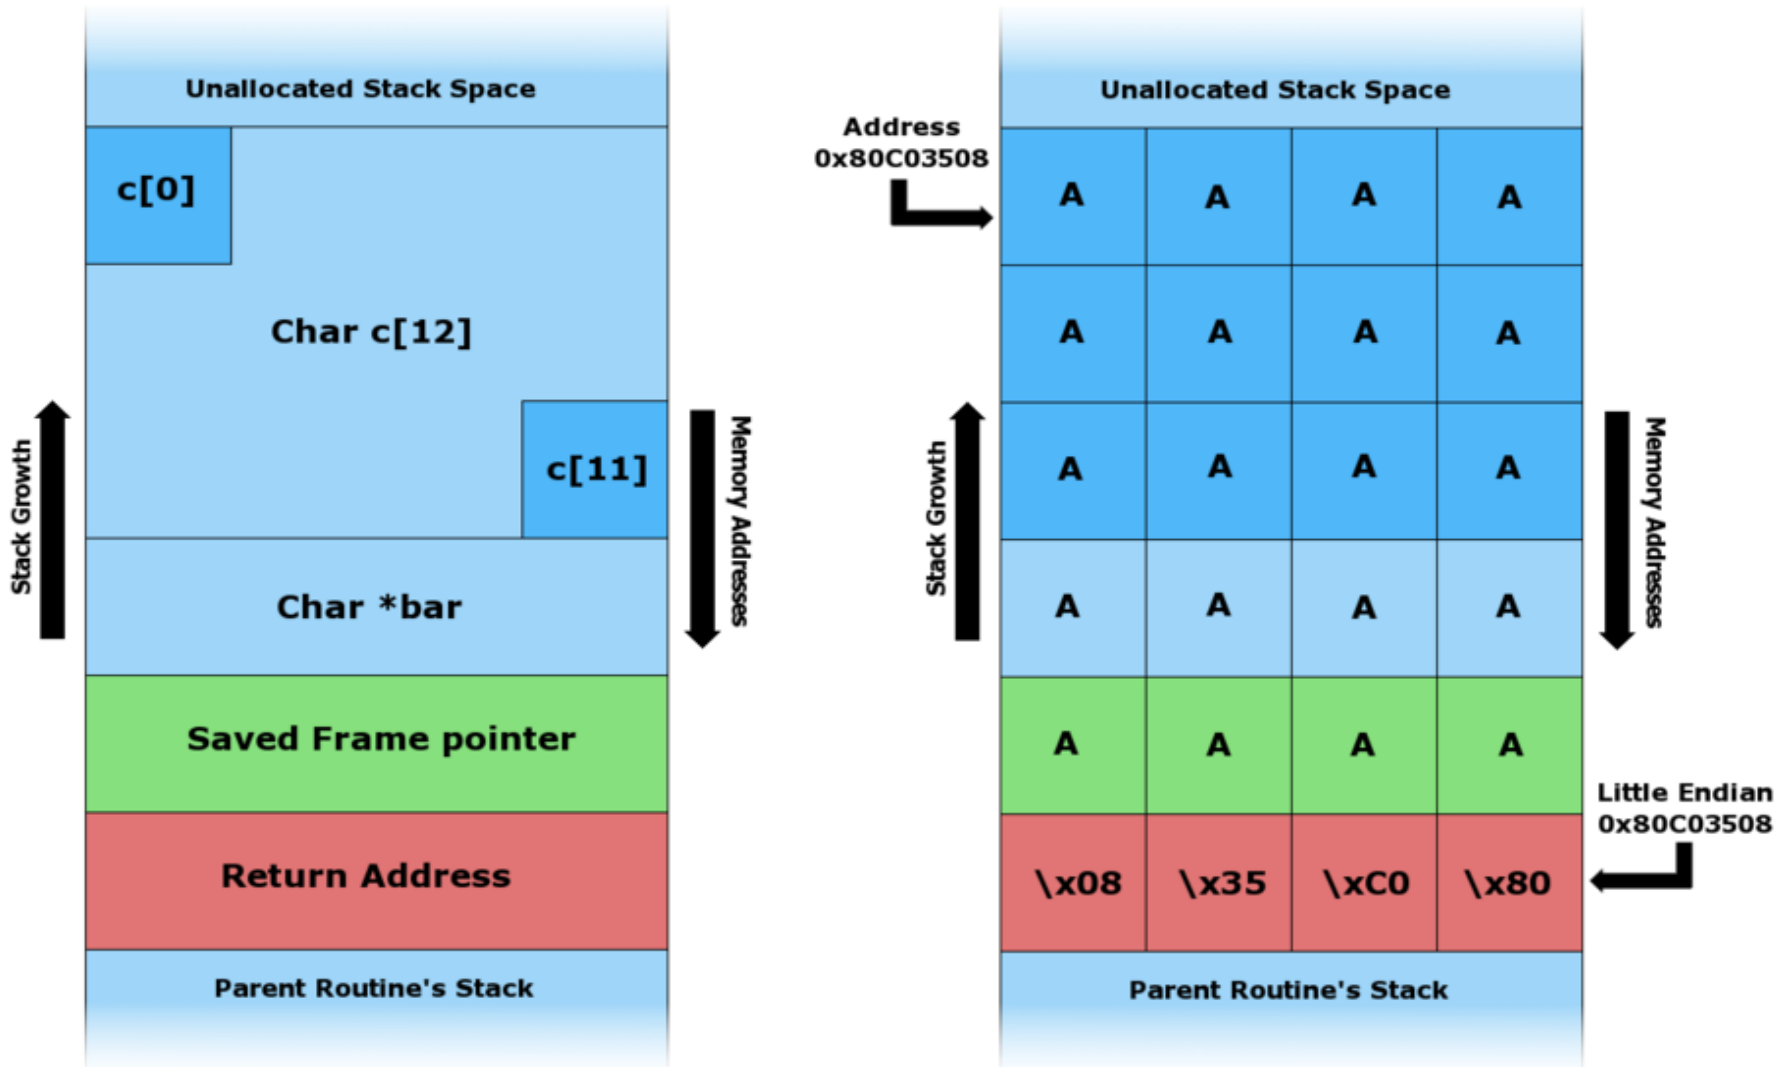
- Code execution

```
1     void foo (char *bar) {
2         char c[12];
3         strcpy(c, bar);
4     }
5     int main (int argc, char **argv) {
6         foo(argv[1]);
7     }
```

# It can get worse



# It can get worse




# It can get worse

- Why limit ourselves to stack? Heap buffer overflow
  - Heap is a linked list
  - What if we corrupt one of the links?

```
1  #define unlink( y, BK, FD ) {  
2      BK = P->bk;  
3      FD = P->fd;  
4      FD->bk = BK;  
5      BK->fd = FD;  
6  }
```

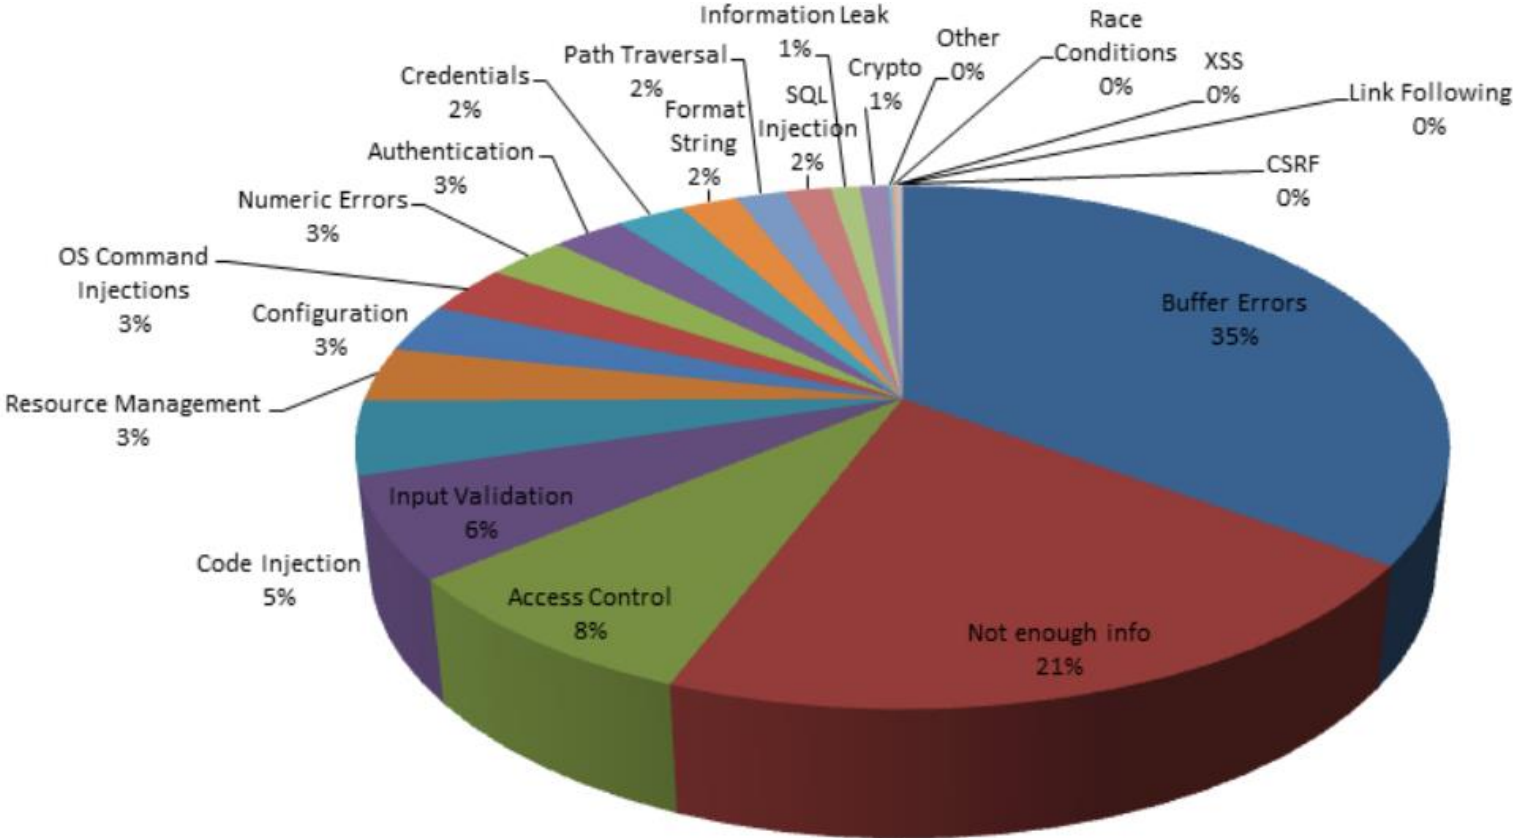
# Why Is It Important?



RSAConference2013 – 25 Years of Vulnerabilities



# Critical Vulnerabilities By Type



# Causes

- Programs written in C
  - Prefer space and performance over security
  - Unsafe language – direct access to memory
- Lack of awareness about security
  - Code is written to work
- Legacy code
  - Knowledge about code isn't preserved
- Inadequate development process

# Defense: Limiting Damage

- Mostly runtime methods:
  - Compiler modifications (stack cookies) - StackGuard
  - Safe libraries - Baratloo, Singh and Tsai
  - Modify program binaries - (assert) SFI
- Pros:
  - Minimal extra work is required from developers
- Cons:
  - Increase performance/memory overhead
  - Simply replace the flaw with a DOS vulnerability

# Defense: Eliminating Flaws

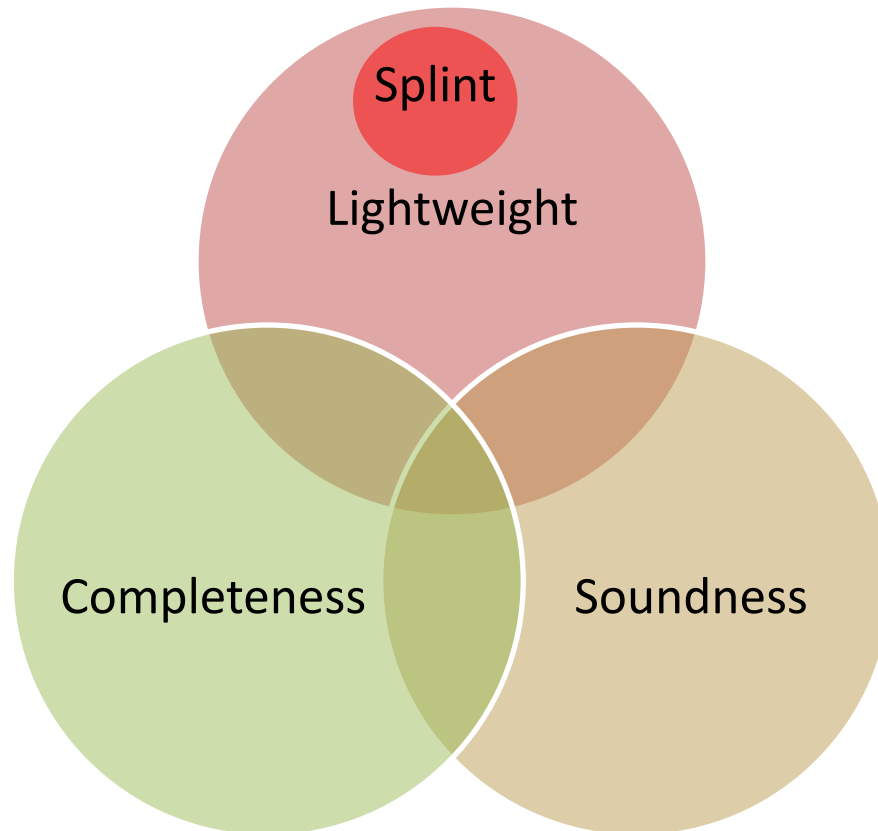
- Human code review
  - Better than automatic tools
  - Can overlook problems
- Testing
  - Mostly ineffective (security wise)
  - Checks expected program behavior
- Fuzzing – doesn't check all possibilities
- Static analysis
  - Allows to codify human knowledge

# Defense: Static Analysis

- Pros:
  - Analyze source code directly, lets us make claims about all possible program paths
  - It still possible to generate useful information
- Cons:
  - Detecting buffer overflows is an undecidable problem
  - Wide range of methods:
    - compilers (low-effort, simple analysis)
    - Full program verifiers (expensive yet effective)

# Solution: Splint – lightweight static analysis

- “Useful results for real programs with reasonable effort”



# Solution: Splint – main ideas

- Static Analysis:
  - Use semantic comments to enable checking of interprocedural properties
  - Use loop heuristics
- Lightweight:
  - good performance and scalability
  - sacrifices soundness and completeness

# Annotations

- Splint is based upon LCLint
  - Annotation assisted static checking tool
- Describe programmer assumptions and intents
- Added to source code and libraries
- Associate with:

Function parameters	Local variables
Function return values	Structure fields
Global variables	Type definitions



# Example: @null@

```
1  typedef /*@null@*/ char *mstring;
2  static mstring mstring_createNew (int x) ;
3  mstring mstring_space1 (void) {
4      mstring m = mstring_createNew (1);
5      /* error, since m could be NULL */
6      *m = ' '; *(m + 1) = '\0';
7      return m;
8  }
```

mstringnn.c: (in function mstring\_space1)

mstringnn.c:6,4: Dereference of possibly null  
pointer m: \*m

# Example: @nonnull@

```
1 static /*@nonnull@*/ mstring
2 mstring_createNewNN (int x) ;
3 mstring mstring_space2 (void) {
4     mstring m = mstring_createNewNN (1);
5     /* no error, because of nonnull annotation */
6     *m = ' ';
7     *(m + 1) = '\\0';
8     return m;
9 }
```

# Annotations: Buffer Overflow

- In LCLint – references are limited to a small number of possible states.
- Splint extends LCLint to support a more general annotation
- Functions `pre(requires)` and `post(ensures)` conditions
- Describe assumptions about buffers:

<code>minSet</code>	<code>maxSet</code>
<code>minRead</code>	<code>maxRead</code>

# Example: Buffer Annotations

- The declaration:

```
char buf[MAXSIZE];
```

Generates the **constraints**:

```
maxSet(buf) = MAXSIZE - 1
```

```
minSet(buf) = 0
```

- Functions **conditions**:

```
char *strcpy (char *s1, const char *s2)  
/*@requires maxSet(s1) >= maxRead(s2)@*/  
/*@ensures maxRead(s1) == maxRead(s2)  
\ result == s1@*/;
```

# Annotations: Buffer Overflow Cont.

- **Constraints** are used to validate **conditions**
- Conditions are found 3 ways:
  - By Splint:

<code>buf[i] = 'a'</code>	Precondition: <code>maxSet(buf) &gt;= i</code>
<code>char a = buf[i]</code>	Precondition: <code>maxRead(buf) &gt;= i</code>
<code>buf[i] = 'a'</code>	Postcondition: <code>maxRead(buf) &gt;= i</code>

- Library functions are annotated
- User generated

# Example: Buffer Overflow Conditions

```
1  void updateEnv(char * str)
2  {
3      char * tmp;
4      tmp = getenv(MYENV);
5      if (tmp != NULL)
6          strcpy(str, tmp);
7  }
```

Unable to resolve constraint:

requires  $\text{maxSet}(\text{str} @ \text{bounds.c:6}) \geq$   
 $\text{maxRead}(\text{getenv}(\text{"MYENV"}) @ \text{bounds.c:4})$

needed to satisfy precondition:

requires  $\text{maxSet}(\text{str} @ \text{bounds.c:6}) \geq \text{maxRead}(\text{tmp} @ \text{bounds.c:6})$

derived from strcpy precondition:

requires  $\text{maxSet}(\langle \text{parameter 1} \rangle) \geq \text{maxRead}(\langle \text{parameter 2} \rangle)$

# Analysis: Settings

- Static analysis is limited
  - Depends on several undecidable problems
- Unsound -> False positives
- Incomplete -> False negatives
- Scalability over precision
- Highly configurable for users

# Analysis: Settings cont.

- Analysis is mostly intraprocedural
  - i.e. at function level
  - Achieve interprocedural (dataflow between functions) using annotations
- Flow-sensitive (order of statements matters) with compromises
  - Handle loops using heuristics



# Analysis: Algorithm

- Programs are analyzed at the function level
- Generate constraints for each C statement
  - By conjoining the constraints of sub expressions
  - Simplify constraints:  $\text{maxSet}(\text{ptr} + i) = \text{maxSet}(\text{ptr}) - i$
- Constraints are resolved at call site
  - Done at statement level
  - Use function preconditions and postconditions of earlier statements

# Analysis: Algorithm Cont.

- All variables used in constraints have an associated location

```
1   t++;  
2   *t = x;  
3   t++;
```

- Leads to the constraints:  
requires  $\text{maxSet}(t @ 1:1) \geq 1$ ,  
ensures  $\text{maxRead}(t @ 3:4) \geq -1$  and  
ensures  $(t @ 3:4) = (t @ 1:1) + 2$

# Axiomatic semantics

- Mathematical logic for proving the correctness of computer programs

$$\{ P \} C \{ Q \}$$

precondition

statement  
a.k.a command

postcondition

- P & Q are state predicates, C a command
- ***If*** P holds and ***if*** C terminates, ***then*** Q will hold

# Analysis: Control Flow - If

- Condition and loops can make unsafe operations, safe
- Analyze an `if` based on the predicate, for example:

```
if (sizeof (s1) > strlen (s2))  
    strcpy (s1, s2);
```

- If the condition holds, the operation is safe

# Analysis: Control Flow - Loops

- Statically analyzing loops is hard. Simplify:
  - Analyze loops according to common idioms
  - First and last iteration matter
- First iteration:
  - Treated as `if` statement
- Last iteration:
  - Determine number of runs base on loop heuristics:
    - `for (index=0;expr;index++) body`
    - `for (init;*buf;buf++)`

# Evaluation

- Analyzed two popular programs
  - wu-ftpd
  - BIND
- Checked detection of known and unknown buffer overflows

# Evaluation: wu-ftpd

- Version wu-ftp-2.5.0 with known security vulnerabilities
- Execution:

Code size(lines)	Time(minutes)
17,000	1

- Source code wasn't modified – resulted 243 warnings(related to buffer overflow)
- Detected both known and unknown buffer overflows

# wu-ftpd: Unknown Bug

```
1 char ls_short[1024];
2 ...
3 extern struct aclmember *
4 getaclentry(char *keyword,
5             struct aclmember **next);
6 ...
7 int main(int argc, char **argv,
8         char **envp)
9 {
10 ...
11 entry = (struct aclmember *) NULL;
12 if (getaclentry("ls_short", &entry)
13     && entry->arg[0]
14     && (int)strlen(entry->arg[0]) > 0)
15 {
16 strcpy(ls_short,entry->arg[0]);
17 ...
```



# wu-ftpd: Unknown Bug Cont.

- Will generate the following warning:

Possible out-of-bounds store.

Unable to resolve constraint:

```
maxRead ((entry->arg[0] @ ftpd.c:1112:23)) <= (1023)
```

needed to satisfy precondition:

```
requires maxSet ((ls_short @ ftpd.c:1112:14))
```

```
>= maxRead ((entry->arg[0] @ ftpd.c:1112:23))
```

derived from strcpy precondition:

```
requires maxSet (<param 1>) >= maxRead (<param 2>)
```

# wu-ftpd: Known Bug

```
1 char mapped_path [200];
2 ...
3 void do_elem(char *dir) {
4 ...
5 if (!(mapped_path[0] == '/'
6     && mapped_path[1] == '\\0'))
7     strcat (mapped_path, "/");
8     strcat (mapped_path, dir);
9 }
```

- `dir` is entered by a remote user
- Reported buffer overflow:
  - <http://www.cert.org/historical/advisories/CA-1999-13.cfm>

# wu-ftpd: False positive

```
1  i = passive_port_max - passive_port_min + 1;
2  port_array = calloc (i, sizeof (int));
3
3  for (i = 3; ... && (i > 0); i--) {
4      for (j = passive_port_max
5           - passive_port_min + 1;
6           ... && (j > 0); j--) {
7          k = (int) ((1.0*j*rand()) /
8                  (RAND_MAX + 1.0));
9          pasv_port_array [j-1]
10         = port_array [k];
```

- Unable to determine that  $1 < k < j$
- Can be suppressed by the user

# wu-ftpd: Summary

- Unmodified code:
  - 243 warnings
- After adding 22 annotations
  - 143 warnings
  - Of these, 88 unresolved constraints involving `maxSet`
  - The rest, can be suppressed by the user
- 225 calls to unsafe functions(`strcat`, `strcpy`, ...)
  - Only 18 reported by Splint
  - 92% of the calls are safe by Splint

# Evaluation: BIND

- Berkley Internet Name Domain – reference for DNS implementations
- Version 8.2.2p7
- Execution:

Code size(lines)	Time(minutes)
47,000	3.5

- Check limited to a subset of code(~3,000 lines)
  - Because of time of human analysis

# BIND: Library Annotations

- Extensive use of internal libraries instead of C library functions
  - Requires annotating large code base
  - Iteratively run Splint and annotate
- To reduce human analysis required
  - Only interface library functions were annotated -> based on code comments and parameters names
- For example:
  - `int foo(char* p, int size)`
  - Resulted: `MaxSet(p) >= (psize - 1)`

# BIND: Library Annotations

- Extensive use of internal libraries instead of C library functions
  - Requires annotating large code base
  - Iteratively run Splint and annotate
- To reduce human analysis required
  - Only interface library functions were annotated -> based on code comments and parameters names
- For example:
  - `int foo(char* p, int size)`
  - Resulted: `MaxSet(p) >= (psize - 1)`

# BIND: req\_query

- Code called in response for querying the domain name server version
- Version string read from configuration file and appended to a buffer
  - OK if used with default
- However, sensitive to
  - Code modification
  - Configuration changes



# BIND Cont.

- BIND uses extensive run time bounds checking
  - This doesn't guarantee safety
  - Buffer overflow was detected, because buffer sizes were calculated incorrectly
- ns\_sign – the functions receives a buffer and its size
  - Splint detected that the size might be incorrect
  - Occurs, If the message contains a signature but the key isn't found
  - Bug was introduced after feature was added

# Related work – Lexical analysis

- Idea:
  - Put simply , use **grep** to find unsafe code
- Tool:
  - ITS4 - [VBKM00]
- Pros:
  - simple and fast
- Cons:
  - Very limited, deserts semantics and syntax

# Related work – Proof-carrying code

- Idea [NL 96, Necula97] :
  - Executable is distributed with a proof and verifier
  - Ensures the executable has certain properties
- Pros:
  - Sound solution
- Cons:
  - At time of paper, wasn't feasible automatically

# Related work – Integer range analysis

- Idea [Wagner et al]:
  - Treat strings as integer range
- Cons:
  - Non-character buffers are abandoned
  - Insensitive analysis – ignore loops and conditions

# Related work – Source transformation

- Idea [Dor, Rodeh and Sagiv]:
  - Instrument the code
  - Assert string operations
  - Then use integer analysis
- Pros:
  - Can handle complex properties – pointer overlapping
- Cons:
  - Doesn't scale

# Conclusions

- Splint isn't perfect but improves security with a reasonable amount of effort
- Splint is lightweight
  - Scalable
  - Efficient
  - Simple
- Hard to introduce to a large code base
  - However, done incrementally can be just right

# Reflections

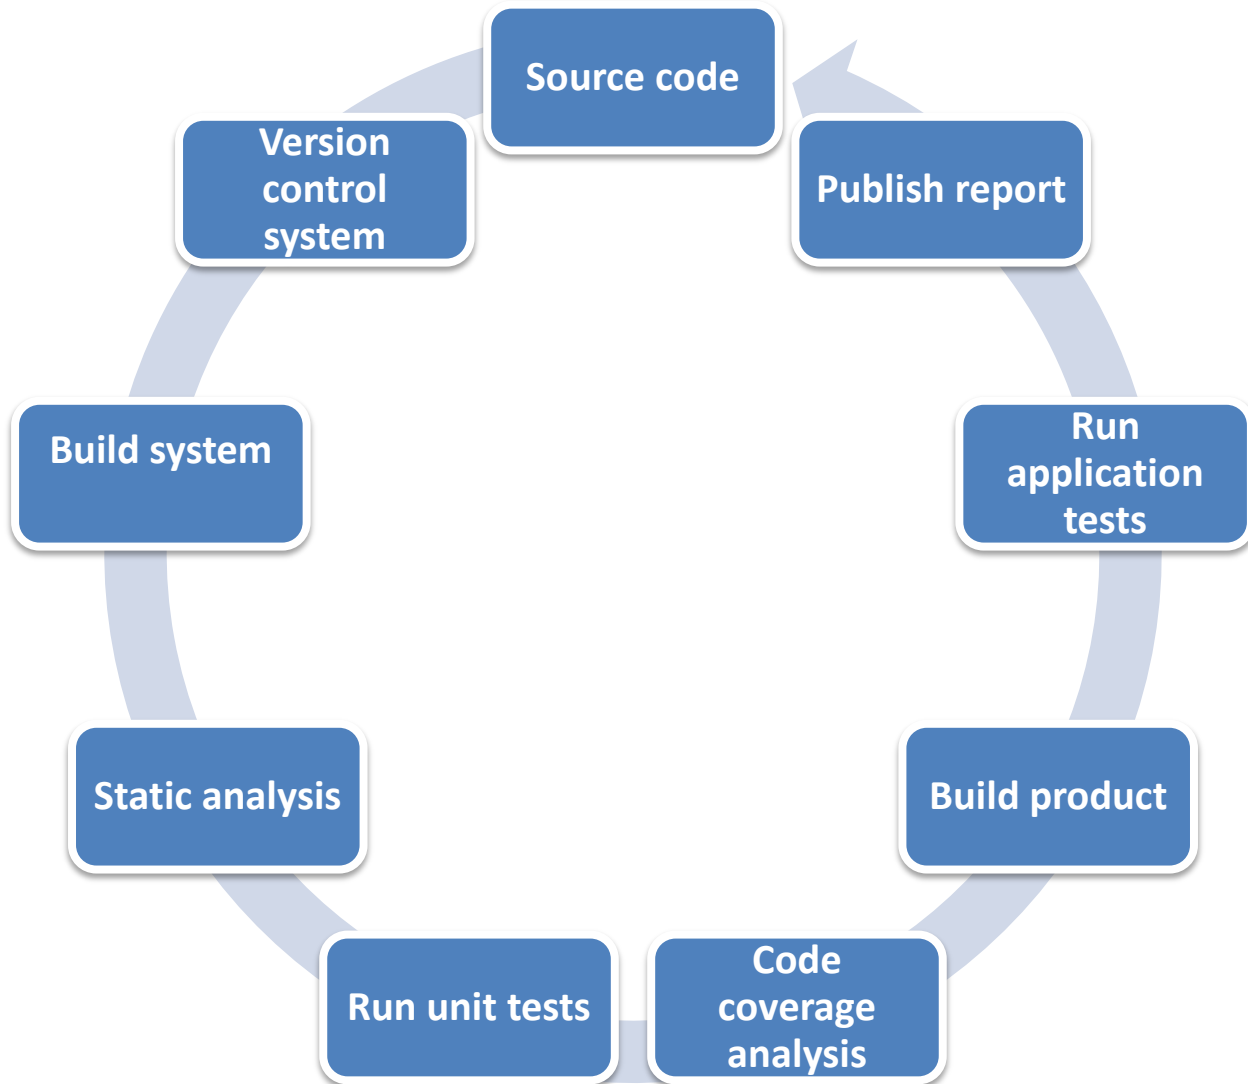
- Practical approach!
  - 80/20 principle in action
- Human effort vs security gained
- Improves code base readability and maintainability
- Can be leveraged as part of Continuous Integration process – Code Quality

# Continuous integration

- “The practice of frequently integrating one's new or changed code with the existing code repository”
- Some of its advantages:
  - Immediate unit testing of all changes
  - Early warning of broken/incompatible code
  - Frequent code check-in pushes developers to create modular, less complex code

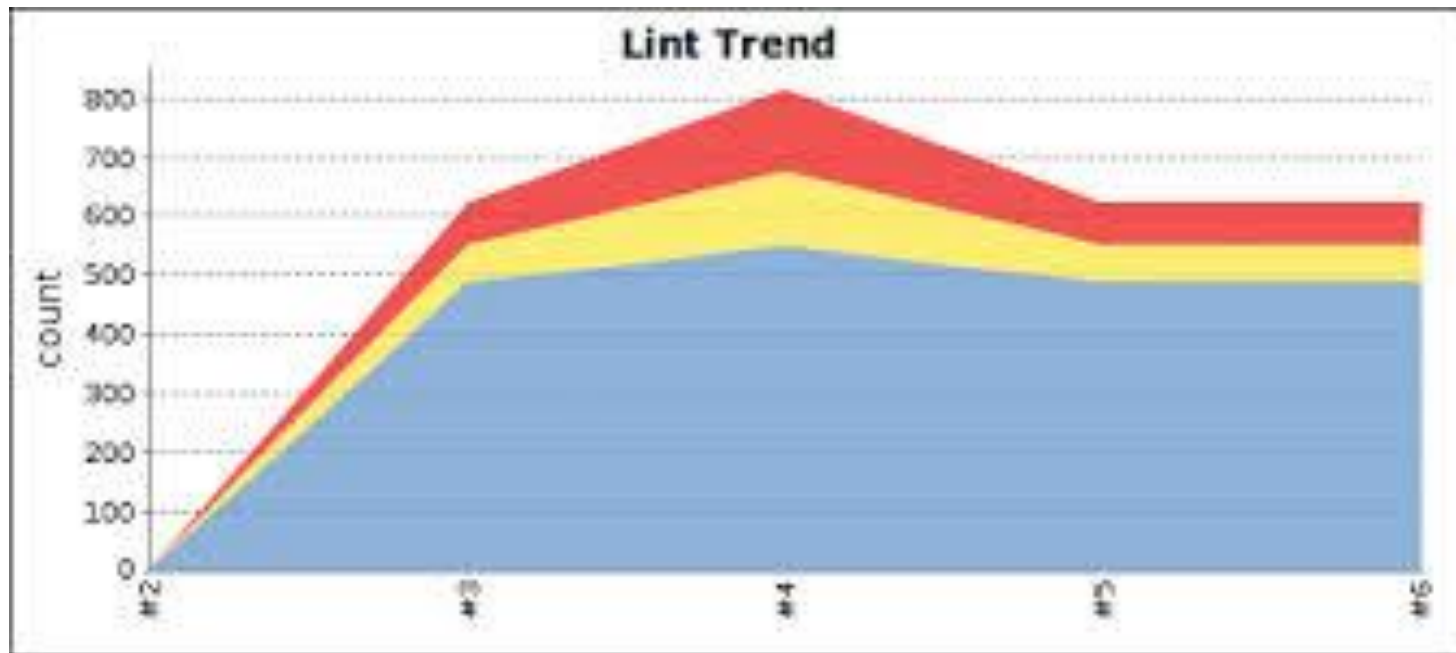


# Continuous integration



# Continuous integration

- Static analysis on source code:
  - Ensures coding standards
  - Assists in avoiding common bugs



Questions?



# Discussion

- Would you use Splint\LCLint in your projects?
  - Security wise?
  - Code quality wise?
- What about dynamic languages?
  - Not necessarily security