SYRCoSE-2016 Krasnovidovo, Moscow Region 30 May 2016



## Verification of Operating Systems

Alexey Khoroshilov khoroshilov@ispras.ru



Institute for System Programming of the Russian Academy of Sciences



## **Operating Systems**





## **Embedded Operating Systems**



Host System



#### Static Verification

#### **Runtime Verification**











Static Verification	<b>Runtime Verification</b>			
+ All paths at once	<ul> <li>One path only</li> </ul>			
<ul> <li>Hardware, test data and test environment is not required</li> </ul>	<ul> <li>Hardware, test data and test environment is required</li> </ul>			
<b>T</b> I <b>CI</b> '''				

There are false positives
 Almost no false positives



Static Verification	<b>Runtime Verification</b>
+ All paths at once	<ul> <li>One path only</li> </ul>
+ Hardware, test data and test environment <b>is not</b> required	<ul> <li>Hardware, test data and test environment is required</li> </ul>
<ul> <li>There are false positives</li> </ul>	+ Almost no false positives
<ul> <li>Checks for predefined set of bugs only</li> </ul>	The only way to show the code actually works



### **Verification Approaches**

all kinds of bugs One test 1 kind bugs in 1 execution in all executions



## **Operating Systems**





## **Functional Testing**

```
void test_memcpy(void)
```

```
char *dst, *src;
char *res;
```

```
dst = malloc(16);
assert(dst != NULL, "not enough memory");
```

free(dst);



## **Verification Approaches**





### T2C - Template2C





Institute for System Programming of the Russian Academy of Sciences



## **Requirements Catalogue**

The Open Group Base Specifications Issue 6 IEEE Std 1003.1, 2004 Edition Copyright © 2001-2004 The IEEE and The Open Group, All Rights reserved.

#### NAME

memcpy - copy bytes in memory

#### SYNOPSIS

#include <<u>string.h</u>>

void \*memcpy(void \*restrict s1, const void \*restrict s2, size t n);

#### DESCRIPTION

🖾 🖾 The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of IEEE Std 1003.1-2001 defers to the ISO C standard. 🖾

{memcpy.01} The memcpy() function shall copy n bytes from the object pointed to by s2 into the object pointed to by s1. {app.memcpy.02} If copying takes place between objects that overlap, the behavior is undefined.

#### RETURN VALUE

{memcpy.03} The memcpy() function shall return s1; no return value is reserved to indicate an error.



## Requirements Catalogue:Requality





## T2C test for memcpy

#### <CODE>

**char** \*res, \*buffer = NULL;

buffer = malloc( 200 );
if (buffer == NULL) ABORT\_TEST\_PURPOSE("Not enough memory");

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "", !are\_buffers\_overlapped(buffer+S1,buffer+S2,N));

```
res = <u>memcpy(</u> buffer, buffer + 100, 100 );
```

// The memcpy() function shall copy n bytes from the object pointed to by s2
// into the object pointed to by s1.
REQ("memcpy.01", "", buffer\_compare( buffer, buffer + 100, 100 ) == 0 );

// The memcpy() function shall return s1
REQ("memcpy.03", "", res == buffer );

```
if (buffer != NULL) free( buffer );
</CODE>
```



## T2C test for memcpy

#### <CODE>

char \*res, \*buffer = NULL;

buffer = malloc( 200 );
if (buffer == NULL) ABORT\_TEST\_PURPOSE("Not enough memory");

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "", !are\_buffers\_overlapped(buffer+S1,buffer+S2,N));

```
res = <u>memcpy(</u> buffer, buffer + 100, 100 );
```

// The memcpy() function shall copy n bytes from the object pointed to by s2
// into the object pointed to by s1.
REQ("memcpy.01", "", buffer\_compare( buffer, buffer + 100, 100 ) == 0 );

```
// The memcpy() function shall return s1
REQ("memcpy.03", "", res == buffer );
</CODE>
```

#### <FINALLY>

```
if (buffer != NULL) free( buffer );
```

#### </FINALLY>



## T2C test for memcpy

#### <CODE>

char \*res, \*buffer = NULL;

buffer = malloc( SIZE );
if (buffer == NULL) ABORT\_TEST\_PURPOSE("Not enough memory");

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "", !are\_buffers\_overlapped(buffer+S1,buffer+S2,N));

```
res = <u>memcpy(</u> buffer + S1, buffer + S2, N );
```

// The memcpy() function shall copy n bytes from the object pointed to by s2
// into the object pointed to by s1.
REQ("memcpy.01", "", buffer\_compare( buffer + S1, buffer + S2, N ) == 0 );

```
// The memcpy() function shall return s1
REQ("memcpy.03", "", res == buffer + S1 );
</CODE>
<FINALLY>
```

```
if (buffer != NULL) free( buffer );
</FINALLY>
```

```
      1000
      1000

      0
      50

      50
      0

      50
      50

      </PURPOSE>
```

<PURPOSE>

<PURPOSE>



#### **Test Generation by Template**





#### **T2C Main Elements**





#### **T2C Test Development Process**





## T2C Results – LSB Desktop

Target Library	Version	Interfaces (Tested of Total)	Requirements (Tested of Total)	Code Coverage (Lines of Code)	Bugs Found
libatk-1.0	1.19.6	222 of 222 (100%)	497 of 515 (96%)	-	11
libglib-2.0	2.14.0	832 of 847 (98%)	2290 of 2461 (93%)	12203 of 16263 (75.0%)	13
libgthread- 2.0	2.14.0	2 of 2 (100%)	2 of 2 (100%)	149 of 211 (70.6%)	0
libgobject- 2.0	2.16.0	313 of 314 (99%)	1014 of 1205 (84%)	5605 of 7000 (80.1%)	2
libgmodule- 2.0	2.14.0	8 of 8 (100%)	17 of 21 (80%)	211 of 270 (78.1%)	2
libfonconfig	2.4.2	160 of 160 (100%)	213 of 272 (78%)	-	11
Total		1537 of 1553 (99%)	4033 of 4476 (90%)	18168 of 23744 (76.5%)	39



#### OLVER - Model Based Testing





Institute for System Programming of the Russian Academy of Sciences



### **Open Linux VERification**





## **OLVER Process**





## **Requirements Catalogue**

The Open Group Base Specifications Issue 6 IEEE Std 1003.1, 2004 Edition Copyright © 2001-2004 The IEEE and The Open Group, All Rights reserved.

#### NAME

memcpy - copy bytes in memory

#### SYNOPSIS

#include <<u>string.h</u>>

void \*memcpy(void \*restrict s1, const void \*restrict s2, size t n);

#### DESCRIPTION

🖾 🖾 The functionality described on this reference page is aligned with the ISO C standard. Any conflict between the requirements described here and the ISO C standard is unintentional. This volume of IEEE Std 1003.1-2001 defers to the ISO C standard. 🖾

{memcpy.01} The memcpy() function shall copy n bytes from the object pointed to by s2 into the object pointed to by s1. {app.memcpy.02} If copying takes place between objects that overlap, the behavior is undefined.

#### RETURN VALUE

{memcpy.03} The memcpy() function shall return s1; no return value is reserved to indicate an error.

#### **ISP**RAS

## memcpy() specification template

#### pre

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "Objects are not overlapped", TODO\_REQ() );

```
return true;
```

#### post

```
/*The memcpy() function shall copy n bytes from the object
pointed to by s2 into the object pointed to by s1. */
REQ("memcpy.01", "s1 contain n bytes from s2", TODO_REQ() );
```

```
/* The memcpy() function shall return s1; */
REQ("memcpy.03", "memcpy() function shall return s1", TODO_REQ() );
```

```
return true;
```



## memcpy() precondition

#### specification

VoidTPtr memcpy\_spec( CallContext context, VoidTPtr s1, VoidTPtr s2, SizeT n )

#### pre

```
/* [Consistency of test suite] */
REQ("", "Memory pointed to by s1 is available in the context",
    isValidPointer(context,s1) );
REQ("", "Memory pointed to by s2 is available in the context",
    isValidPointer(context,s2) );
```

/\* [Implicit precondition] \*/
REQ("", "Memory pointed to by s1 is enough", sizeWMemoryAvailable(s1) >= n );
REQ("", "Memory pointed to by s2 is enough", sizeRMemoryAvailable(s2) >= n );

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "Objects are not overlapped", !areObjectsOverlapped(s1,n,s2,n) );

return true;



## **OLVER Distributed Architecture**





## memcpy() precondition

#### specification

VoidTPtr memcpy\_spec( CallContext context, VoidTPtr s1, VoidTPtr s2, SizeT n )

#### pre

```
/* [Consistency of test suite] */
REQ("", "Memory pointed to by s1 is available in the context",
    isValidPointer(context,s1) );
REQ("", "Memory pointed to by s2 is available in the context",
    isValidPointer(context,s2) );
```

/\* [Implicit precondition] \*/
REQ("", "Memory pointed to by s1 is enough", sizeWMemoryAvailable(s1) >= n );
REQ("", "Memory pointed to by s2 is enough", sizeRMemoryAvailable(s2) >= n );

// If copying takes place between objects that overlap, the behavior is undefined. REQ("app.memcpy.02", "Objects are not overlapped", !areObjectsOverlapped(s1,n,s2,n) );

return true;



## memcpy() postcondition

```
specification
VoidTPtr memcpy spec( CallContext context, VoidTPtr s1, VoidTPtr s2, SizeT n ) {
  post
    /*The memcpy() function shall copy n bytes from the object
      pointed to by s2 into the object pointed to by s1. */
    REQ("memcpy.01", "s1 contain n bytes from s2",
        equals( readCByteArray VoidTPtr(s1,n), @readCByteArray VoidTPtr(s2,n))
      );
    /* [The object pointed to by s2 shall not be changed] */
    REQ("", "s2 shall not be changed",
        equals( readCByteArray_VoidTPtr(s2,n), @readCByteArray_VoidTPtr(s2,n) ));
    /* The memcpy() function shall return s1; */
    REQ("memcpy.03", "memcpy() function shall return s1", equals VoidTPtr(memcpy spec, s1));
    /* [Other memory shall not be changed] */
    REQ("", "Other memory shall not be changed",
        equals( readCByteArray MemoryBlockExceptFor( getTopMemoryBlock(s1), s1, n),
              @readCByteArray MemoryBlockExceptFor( getTopMemoryBlock(s1), s1, n ) );
    return true;
```



#### **Test Oracle Generation**





## **Test Scenarios Generation**



- abstract state
- set of test actions
- Test Engine generates a sequence of actions
  - to ensure: each action is executed in each abstract state



## Bug Example - POSIX mq



- sending threads are blocked if queue is full
- receiving threads are blocked if queue is empty
- Test scenario abstract state:
  - number of messages in the queue
  - number of threads waiting to send
  - number of threads waiting to receive



(0,0,0)

# Bug Example - POSIX mq



## Bug Example - POSIX mq






	1	mmmmmmmm	1	(1,N,1)
--	---	----------	---	---------













5. 2<sup>nd</sup> sender never unblocked



## **OLVER Results**

- Requirements catalogue built for LSB and POSIX
  - 1532 interfaces
  - 22663 elementary requirements
- 97 deficiencies in specification reported
- Formal specifications and tests developed for
  - **1270 interface** (good quality)
  - + 260 interfaces (basic quality)
- 80+ bugs reported in modern distributions
- OLVER is a part of the official LSB Certification test suite http://ispras.linuxfoundation.org



- model based testing allows to achieve better quality using less resources
- maintenance of MBT is cheaper



- model based testing allows to achieve better quality using less resources if you have advanced test engineers
- maintenance of MBT is cheaper if you have advanced test engineers



- model based testing allows to achieve better quality using less resources if you have advanced test engineers
- maintenance of MBT is cheaper if you have advanced test engineers
- traditional tests are more useful for typical test engineers and developers



- model based testing allows to achieve better quality using less resources if you have advanced test engineers
- maintenance of MBT is cheaper if you have advanced test engineers
- traditional tests are more useful for typical test engineers and developers
- so, long term efficiency is questionable
- but...



40

Subject [lsb-bugs] [Bug 3285] New neurses test failures 11.05.2016 06:08 To lsb-bugs@lists.linuxfoundation.org Comment # 13 on bug 3285 from zhpeng@linx-info.com As the nature of the bug is that agent side pty state is uncertained, we have 2 solving methods. One is on agent side, the other one is on local model side. Method one is described in the patch and previous comments. Method tow is also Jeff's suggestion, I tried it but I found it is not sufficient. The bug in test rw case a, rw case b and rw canon can be fixed, but rw case d can not because its program structure is different from the other 3. In test rw case a, rw case b and rw canon, there is a common scenario to set the local termios, that is setattr a scen in rw case a, setattr b scen in rw case b, tcsetattr rw scen in rw canon. But rw case d has not such a setattr d scen between open d scen and read d scen. I added a setattr d scen following the other rw case\*, but the test still failed. I gave up this way. So I still advise to fix the bug in the agent side personally. The code for fixing the bug on local model side is as followed, they are suitable for test rw case a, rw case b and rw canon separately. Note that the csize and cread flag should be set explicitly. /home/rocky/source/original/lsb-test-olver-core-4.1.4/src/model/io/term/tests/rw case 2012-10-26 06:29:03.000000000 +0800 +++ rw case a scenario.sec 2016-05-11 10:13:01.000000000 +0800 @@ -241,6 +241,15 @@



### **OS** Verification



#### LSB Desktop

LSB Desktop 3.1 (18841 interfaces)





#### **Some Statistics**

	Release Date	System Calls	Libraries	Interfaces	Utilities
Debian 7.0	May 2013	~350	~1650	~ 720 000	~10 000
RTOS	Nov 2013	~200	-	~700	~80



#### **API Sanity Autotest**





Institute for System Programming of the Russian Academy of Sciences



#### Smoke Testing

Smoke testing – checks only main use cases for basic requirements, i.e. the system under test is not broken and its results looks like correct.



#### **API** Sanity Autotest





#### **API** Sanity Autotest





#### Additional semantic information

- How to initialize library
- Hot to get a valid data of a particular type
- What is a valid argument of a particular function
- Which check can be done for returned types if it is of a particular type



#### **Special Expressions**

```
• $(type) - create an object of particular type
void create_QProxyModel(QProxyModel* Obj) {
    Obj->setSourceModel($(QItemModel*));
}
```

• \$[interface] – call the interface with some valid arguments

```
xmlListPtr create_filled_list() {
    xmlListPtr l = $[xmlListCreate];
    int num = 100;
    xmlListPushBack(l,&num);
    return l;
}
```



#### Sources of Scalability

- Special expressions
- Extensive reuse
  - $\rightarrow$  Minimal duplication of code



#### Results

Libraries	Number of interfaces
libqt-mt	9 792
libQtCore	2 066
libQtGui	7 439
libQtNetwork	406
libQtOpenGL	84
libQtSql	362
libQtSvg	66
libQtXml	380
libxml2	1 284
Total	21 879



#### Bugs Found (1) Git index : freetype/freetype2.git

summary refs log tree **commit** diff about

author	Werner Lemberg <wl@gnu.org></wl@gnu.org>	2010-05-22 18:03:41 (GMT)
committer	Werner Lemberg <wl@gnu.org></wl@gnu.org>	2010-05-22 18:03:41 (GMT)
commit	e30de299f28370ed5aa65755c6be69da	a58eefc72 (patch)
tree	35355e4d7b42156baea7a21037d50f4c	:4ca5753c
parent	09344385ee652cb8df33d35f07627c93	3e827f10d (diff)
download	freetype2-e30de299f28370ed5aa657	755c6be69da58eefc72.tar.gz

#### Fix various memory problems found by linuxtesting.org.

```
* src/base/ftgxval.c (FT_TrueTypeGX_Free, FT_ClassicKern_Free),
src/base/ftotval.c (FT_OpenType_Free), src/base/ftpfr.c
(ft_pfr_check): Check `face'.
```

```
* src/base/ftobjs.c (FT_Get_Charmap_Index): Check `charmap' and
`charmap->face'.
(FT_Render_Glyph): Check `slot->face'.
(FT_Get_SubGlyph_Info): Check `glyph->subglyphs'.
```

Improve API documentation.

http://git.savannah.gnu.org/cgit/freetype/freetype2.git/commit/?id=e30de299f28370ed5aa65755c6be69da58eefc72



# Bugs Found (1)

#### LIBSSH2

Login OpenID Login Preferences

Timeline	Roadmap	Browse Source

#### Ticket #173 (closed defect: fixed)

Agent API does	n't call_libssh2_error co	onsistently	Opened <mark>6 weeks</mark> ago Last modified <mark>5 weeks</mark> ago
Reported by:	alamaison	Owned by:	bagder
Priority:	normal	Milestone:	1.2.7
Component:	API	Version:	1.2.6
Keywords:		Cc:	
Blocks:		Blocked By:	
Description			
All APIs should se most places.	et an error message with _	_libssh2_error but the lib	ossh2_agent_* API fails to do this in

Wiki

#### http://trac.libssh2.org/ticket/173



#### **OS** Verification





# Test Aspects (1)

	T2C	OLVER	Autotest
Monitoring Aspects			
Kinds of Observable Events			
interface events	+	+	+
internal events			
Events Collection			
internal	+	+	+
external			
embedded			
Requirements Specification			
in-place (local, tabular)	+		+
formal model (pre/post+invar	ants,)	+	
assertions/prohibited events	External	External	External
Events Analysis			
online	+	+	+
in-place	+		+
outside		+	
offline			

#### Test Aspects (2)

	T2C	OLVER	Autotest
Active Aspects			
Target Test Situations Set			
requirements coverage	+	+	
class equivalence coverage		+	
model coverage (SUT/reqs)		+	
source code coverage			
Test Situations Setup/Set Gen			
passive			
fixed scenario	+		+
manual	+		
pre-generated			
coverage driven			
random			+-
adapting scenario		+	
coverage driven		+	
source code coverage			
model/ coverage		+	
random			
TestActions			
application interface	+	+	+
HW interface			
internal actions			
inside			
outside			

#### **ISPRAS**



#### **Configuration Testing**





Institute for System Programming of the Russian Academy of Sciences



N₂	Platform	NP	Params	С	N1	N2	T1	T2
1	i386/i386	228	4 <sup>2</sup> 3 <sup>6</sup> 2 <sup>220</sup>	172	25	92	0.01	0.31
2	i386/x86	215	4 <sup>2</sup> 3 <sup>6</sup> 2 <sup>207</sup>	170	23	72	0.01	0.26
3	mips/bt205	373	17 <sup>2</sup> 15 <sup>1</sup> 8 <sup>8</sup> 7 <sup>2</sup> 6 <sup>1</sup> 5 <sup>3</sup> 4 <sup>19</sup> 3 <sup>21</sup> 2 <sup>316</sup>	279	322	670	1.57	5.07
4	mips/ bt23-202	331	17 <sup>2</sup> 15 <sup>1</sup> 13 <sup>4</sup> 9 <sup>4</sup> 7 <sup>2</sup> 5 <sup>3</sup> 4 <sup>7</sup> 3 <sup>10</sup> 2 <sup>298</sup>	267	319	653	1.07	4.58
5	mips64/ bt128	548	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>1</sup> 3 <sup>15</sup> 2 <sup>526</sup>	500	289	556	0.57	7.88
6	mips64/ bt211	334	17 <sup>2</sup> 15 <sup>1</sup> 13 <sup>4</sup> 9 <sup>4</sup> 7 <sup>2</sup> 5 <sup>3</sup> 4 <sup>5</sup> 3 <sup>10</sup> 2 <sup>302</sup>	277	319	726	0.78	4.08
7	mips64/ bt206	334	17 <sup>2</sup> 15 <sup>1</sup> 7 <sup>2</sup> 6 <sup>1</sup> 5 <sup>3</sup> 4 <sup>5</sup> 3 <sup>18</sup> 2 <sup>302</sup>	278	289	651	0.58	2.99
8	mips64 /mpon	301	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>1</sup> 3 <sup>15</sup> 2 <sup>279</sup>	251	289	574	0.33	2.15
9	mips64/ vmips	241	5 <sup>1</sup> 4 <sup>1</sup> 3 <sup>6</sup> 2 <sup>233</sup>	203	27	89	0.01	0.42
10	mips64/ cprio64	563	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>1</sup> 3 <sup>15</sup> 2 <sup>541</sup>	511	289	564	0.58	6.70
11	komdiv64/ bt128	548	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>1</sup> 3 <sup>15</sup> 2 <sup>526</sup>	500	289	556	0.56	7.80
12	R4000/bt128	419	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>2</sup> 3 <sup>14</sup> 2 <sup>397</sup>	377	289	541	0.44	7.00
13	R4000/bt206	220	17 <sup>2</sup> 15 <sup>1</sup> 14 <sup>14</sup> 7 <sup>2</sup> 6 <sup>1</sup> 5 <sup>3</sup> 4 <sup>6</sup> 3 <sup>17</sup> 2 <sup>174</sup>	155	823	826	1.28	1.65
14	R4000/mpon	183	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>2</sup> 3 <sup>14</sup> 2 <sup>161</sup>	138	289	543	0.23	0.56
15	R4000/ vmips	113	5 <sup>1</sup> 4 <sup>2</sup> 3 <sup>5</sup> 2 <sup>105</sup>	80	28	62	0.01	0.07
16	R4000/ cprio64	434	17 <sup>2</sup> 15 <sup>1</sup> 5 <sup>3</sup> 4 <sup>2</sup> 3 <sup>14</sup> 2 <sup>412</sup>	388	289	549	0.46	5.59

V.V. Kuliamin. Combinatorial generation of operation system software configurations. Proceedings of the Institute for System Programming Volume 23. 2012 y. pp. 359-370.



### **OS** Verification





# Test Aspects (1)

	T2C	OLVER	Autotest	Cfg
Monitoring Aspects				-
Kinds of Observable Events				
interface events	+	+	+	
internal events				
Events Collection				
internal	+	+	+	
external				
embedded				
Requirements Specification				
in-place (local, tabular)	+		+	lf
formal model (pre/post+invar	ants,)	+		lf
assertions/prohibited events	External	External	External	Co
Events Analysis				
online	+	+	+	
in-place	+		+	
outside		+		
offline				

#### Test Aspects (2)

	T2C	OLVER	Autotest	Cfg
Active Aspects				+-
Target Test Situations Set				cfgs
requirements coverage	+	+		
class equivalence coverage		+		
model coverage (SUT/reqs)		+		
source code coverage				
Test Situations Setup/Set Gen				
passive				
fixed scenario	+		+	
manual	+			
pre-generated				
coverage driven				+-
random			+-	
adapting scenario		+		
coverage driven		+		
source code coverage				
model/ coverage		+		
random				
TestActions				
application interface	+	+	+	
HW interface				
internal actions				
inside				
outside				

#### **ISP**RAS



#### **Robustness Testing**





Institute for System Programming of the Russian Academy of Sciences



# Fault Handling Code

- Is not so fun
- Is really hard to keep all details in mind
- Practically is not tested
- Is hard to test even if you want to
- Bugs seldom(never) occurs
   => low pressure to care



## Why do we care?

- It beats someone time to time
- Safety critical systems
- Certification authorities



# **Operating Systems**





# **Run-Time Testing of Fault Handling**

#### Manually targeted test cases

- + The highest quality
- Expensive to develop and to maintain
- Not scalable
- Random fault injection on top of existing tests
  - + Cheap
  - Oracle problem
  - No any guarantee
  - When to finish?


## Systematic Approach

#### Hypothesis:

- Existing test lead to more-or-less deterministic control flow in kernel code
- Idea:
  - Execute existing tests and collect all potential fault points in kernel code
  - Systematically enumerate the points and inject faults there



## Experiments – Outline

- Target code
- Fault injection implementation
- Methodology
- Results



## Experiments – Target

- Target code: file system drivers
- Reasons:
  - Failure handling is more important than in average
    - Potential data loss, etc.
  - Same tests for many drivers
  - It does not require specific hardware
  - Complex enough



#### Linux File System Layers





### File System Drivers - Size

File System Driver	Size, LoC
JFS	18 KLOC
Ext4	37 KLoC with jbd2
XFS	69 KLoC
BTRFS	82 KLoC
F2FS	12 KLoC



# File System Driver – VFS Interface

- file\_system\_type
- super\_operations
- export\_operations
- inode\_operations
- file\_operations
- vm\_operations
- address\_space\_operations
- dquot\_operations
- quotactl\_ops
- dentry\_operations

~100 interfaces in total



### FS Driver – Userspace Interface

File System Driver	ioctl	sysfs
JFS	6	-
Ext4	14	13
XFS	48	-
BTRFS	57	-



### FS Driver – Partition Options

File System Driver	mount options	mkfs options
JFS	12	6
Ext4	50	~30
XFS	37	~30
BTRFS	36	8



## FS Driver – On-Disk State

- File System Hierarchy
- \* File Size
- \* File Attributes
- \* File Fragmentation
- \* File Content (holes,...)



## FS Driver – In-Memory State

- Page Cache State
- Buffers State
- Delayed Allocation



#### Linux File System Layers





## FS Driver – Fault Handling

- Memory Allocation Failures
- Disk Space Allocation Failures
- Read/Write Operation Failures



## Fault Injection - Implementation

#### Based on KEDR framework\*

- intercept requests for memory allocation/bio requests
  - to collect information about potential fault points
  - to inject faults
- also used to detect memory/resources leaks



#### **KEDR Workflow**



http://linuxtesting.org/project/kedr



### Experiments – Oracle Problem

- Assertions in tests are disabled
- Kernel oops/bugs detection
- Kernel assertions, lockdep, memcheck, etc.
- Kernel sanitizers
- KEDR Leak Checker



# Methodology – The Problem

- Source code coverage is used to measure results on fault injection
- If kernel crashes code, coverage results are unreliable



# Methodology – The Problem

- Source code coverage is used to measure results on fault injection
- If kernel crashes code, coverage results are unreliable
- As a result
  - Ext4 was analyzed only
  - XFS, BTRF, JFS, F2FS, UbiFS, JFFS2 crashes and it is too labor and time consuming to collect reliable data



#### **Experiment Results**



#### Systematic vs. Random

	Increment new lines	Time min	Cost second/line
Xfstests without fault simulation	-	2	-
Xfstests+random(p=0.005,repeat=200)	411	182	27
Xfstests+random(p=0.01,repeat=200)	380	152	24
Xfstests+random(p=0.02,repeat=200)	373	116	19
Xfstests+random(p=0.05,repeat=200)	312	82	16
Xfstests+random(p=0.01,repeat=400)	451	350	47
Xfstests+stack filter	423	90	13
Xfstests+stackset filter	451	237	31



## Systematic vs. Random

- + 2 times more cost effective
- + Repeatable results
- Requires more complex engine

- + Cover double faults
- Unpredictable
- Nondeterministic

#### **ISPRAS**

No.	Туре	Brief	Added on	Accepted	Status	
F0011	Crash	ext4: When mounted with backup superblock online resize leads to BUG_ON or causes filesystem corruption	2014-12-27	http://www.spinics.net/lists /linux-ext4/msg46743.html commit	Fixed kernel 3.19-rc4	in
F0010	Crash	f2fs: Possible use-after-free when umount filesystem	2014-07-25	https://lkml.org/lkml/2014 /7/21/198 commit	Fixed kernel 3.17-rc1	in
F0009	Crash	ext4: Destruction of ext4_groupinfo_caches during one mount causes BUG_ON for other mounted ext4 filesystems	2014-05-12	https://lkml.org/lkml/2014 /5/12/147 commit	Fixed kernel 3.16-rc1	in
F0008	Crash	f2fs: BUG_ON() is triggered in recover_inode_page() when mount valid f2fs filesystem	2014-04-18	https://lkml.org/lkml/2014 /4/14/189 commit	Fixed kernel 3.17-rc1	in
F0007	Crash	f2fs: f2fs unmount hangs if f2fs_init_acl() fails during mkdir syscall	2014-02-17	https://lkml.org/lkml/2014 /2/6/18 commit	Fixed kernel 3.15-rc1	in
F0006	Deadlock	f2fs: a deadlock in mkdir if ACL is enabled	2013-10-28	https://lkml.org/lkml/2013 /10/26/163 commit	Fixed kernel 3.12-rc3	in
F0005	Crash	ext4: system hangs after failure in ext4_mb_new_preallocation()	2013-07-01	https://lkml.org/lkml/2013 /5/5/64 commit	Fixed kernel 3.10-rc3	in
F0004	Deadlock	ext4: deadlocks after allocation failure in ext4_init_io_end()	2013-06-04	https://lkml.org/lkml/2013 /5/13/426 commit	Fixed kernel 3.10-rc3	in
F0003	Crash	jfs: Several bugs in jfs_freeze() and jfs_unfreeze()	2013-05-24	https://lkml.org/lkml/2013 /5/24/76 commit	Fixed kernel 3.10-rc3	in
F0002	Crash	ext4: NULL dereference in ext4_calculate_overhead()	2012-11-28	https://lkml.org/lkml/2012 /11/28/354 commit	Fixed kernel 3.8-rc1	in
F0001	Crash	ext4: NULL pointer dereference in mount_fs() because of ext4_fill_super() wrongly reports	2012-11-08	https://bugzilla.kernel.org /show_bug.cgi?id=48431	Fixed kernel	in



### **OS** Verification





## Test Aspects (1)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC
Monitoring Aspects				-	-	
Kinds of Observable Events						
interface events	+	+	+			
internal events						+
Events Collection						
internal	+	+	+			+
external						
embedded						
Requirements Specification						Specific
in-place (local, tabular)	+		+	lf	Dis	
formal model (pre/post+invari	ants,)	+		lf	Co	
assertions/prohibited events	External	External	External	Co	Co	
Events Analysis						
online	+	+	+			
in-place	+		+			+
outside		+				
offline						

#### Test Aspects (2)

**ISP**RAS

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC
Active Aspects				+-	+	-
Target Test Situations Set				cfgs		
requirements coverage	+	+				
class equivalence coverage		+				
model coverage (SUT/reqs)		+				
source code coverage				al	most	
Test Situations Setup/Set Gen						
passive						
fixed scenario	+		+			
manual	+					
pre-generated						
coverage driven				+-		
random			+-			
adapting scenario		+				
coverage driven		+				
source code coverage				al	most	
model/ coverage		+				
random				as	opti	on
TestActions						
application interface	+	+	+			
HW interface						
internal actions					+	
inside					+	
outside						



#### Concolic Testing (S2E)





Institute for System Programming of the Russian Academy of Sciences



# **Concolic Testing**

#### Concolic = Symbolic + Concrete

- SUT runs in concrete and in symbolic modes
- Symbolic execution is used to collect conditions and branches of the current path
- Collected data is used to generate new input data to cover more execution paths



### **Concolic Tools**

Tool	Language	Platform	Constraint Solver
DART	С	NA	lp_solver
SMART	С	Linux	lp_solver
CUTE	С	Linux	lp_solver
CREST	С	Linux	Yices
EXE	С	Linux	STP
KLEE	C (LLVM bitcode)	Linux	STP
Rwset	С	Linux	STP
PathCrawler	С	NA	NA
SAGE	Machine code	Windows	Disolver



#### S2E

- based on KLEE
- uses patched Qemu
  - source code is not required
- supports plugins

(\*) https://s2e.epfl.ch/



### S2E – OS Level Checks

- lockdep
- asserts
- memleak
- sanitizers





#### S2E – VM Level Checks

many pathsrequires plugins





## S2E for Linux File Systems

- Instrumentation FS driver with SystemTap
- Mark metadata readed from disk as symbolic

```
static inline void s2e_make_symbolic(void *buf, int size, char *name){
    asm____volatile__(
    ".byte 0x0f, 0x3f\n"
    ".byte 0x00, 0x03, 0x00, 0x00\n"
    ".byte 0x00, 0x00, 0x00, 0x00\n"
    : : "a" (buf), "b" (size), "c" (name) );
}
probe vfs.read, vfs.write
{
    s2e_make_symbolic(
    &($file->f_path->dentry->d_inode->i_count),
    4,
    "i_count")
}
```



### **OS** Verification





## Test Aspects (1)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E
Monitoring Aspects				-	-	+	+-
Kinds of Observable Events							
interface events	+	+	+				
internal events						+	+
Events Collection							
internal	+	+	+			+	
external							+
embedded							
Requirements Specification						Specific	Plugin
in-place (local, tabular)	+		+	lf	Dis		Dis
formal model (pre/post+invar	ants,)	+		lf	Co		Со
assertions/prohibited events	External	External	External	Co	Co		Со
Events Analysis							
online	+	+	+				
in-place	+		+			+	
outside		+					
offline							

#### Test Aspects (2)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E
Active Aspects				+-	+	-	+
Target Test Situations Set				cfgs			
requirements coverage	+	+					
class equivalence coverage		+					
model coverage (SUT/reqs)		+					
source code coverage				al	most		+
Test Situations Setup/Set Gen							
passive							
fixed scenario	+		+				
manual	+						
pre-generated							
coverage driven				+-			
random			+-				
adapting scenario		+					
coverage driven		+					
source code coverage				al	most		+
model/ coverage		+					
random				as	opti	on	
TestActions							
application interface	+	+	+				
HW interface							
internal actions					+		+
inside					+		
outside							+

**ISP**RAS



#### Race Hound





Institute for System Programming of the Russian Academy of Sciences



#### Data race

- Data race is a situation if two threads access one location in memory
  - at least one of them is a write
  - they are not both synchronization accesses
  - can be executed on a multiprocessor in such a way that two conflicting memory accesses are performed simultaneously




### Race Hound

idea close to DataCollider (Microsoft Research)



https://forge.ispras.ru/projects/race-hound



### Race Hound

idea close to DataCollider (Microsoft Research)



https://forge.ispras.ru/projects/race-hound



## Race Hound Results

- works on x86/x86-64 multicore
- Linux kernel >= 3.2
- acknowledged 3 data race in Intel e1000 Linux device driver
- ~10 data race in total



# Test Aspects (1)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E	RH
Monitoring Aspects				-	-	+	+-	+
Kinds of Observable Events								
interface events	+	+	+					
internal events						+	+	+
Events Collection								
internal	+	+	+			+		
external							+	
embedded								
Requirements Specification						Specific	Plugin	Specific
in-place (local, tabular)	+		+	lf	Dis		Dis	
			1					
formal model (pre/post+invar	ants,)	+	•	lf	Co		Co	
formal model (pre/post+invar assertions/prohibited events	ants,) External	+ External	External	lf Co	Co Co		Co Co	
formal model (pre/post+invar assertions/prohibited events Events Analysis	ants,) External	+ External	External	If Co	Co Co		Co Co	
formal model (pre/post+invar assertions/prohibited events Events Analysis online	ants,) External +	+ External +	External	lf Co	Co Co		Co Co	
formal model (pre/post+invar assertions/prohibited events Events Analysis online in-place	ants,) External + +	+ External +	External + +	lf Co	Co	+	Co Co	+
formal model (pre/post+invar assertions/prohibited events Events Analysis online in-place outside	ants,) External + +	+ External +	External + +	If Co	Co	+	Co	+



# Test Aspects (1)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E	RH	KStrider
Monitoring Aspects				-	-	+	+-	+	+-
Kinds of Observable Events									
interface events	+	+	+						
internal events						+	+	+	+
Events Collection									
internal	+	+	+			+			+
external							+		
embedded									
Requirements Specification						Specific	Plugin	Specific	Specific
in-place (local, tabular)	+		+	lf	Dis		Dis		
formal model (pre/post+invar	ants,)	+		lf	Co		Со		
assertions/prohibited events	External	External	External	Co	Co		Со		
Events Analysis									
online	+	+	+						
in-place	+		+			+		+	
outside		+							
offline									+

#### Test Aspects (2)



	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E	RH
Active Aspects				+-	+	-	+	+
Target Test Situations Set				cfgs				Specific
requirements coverage	+	+						
class equivalence coverage		+						
model coverage (SUT/reqs)		+						
source code coverage				a	most		+	
Test Situations Setup/Set Gen								
passive								+-
fixed scenario	+		+					
manual	+							
pre-generated								
coverage driven				+-				
random			+-					
adapting scenario		+						
coverage driven		+						
source code coverage				a	most		+	
model/ coverage		+						
random				as	s opti	on		
TestActions								
application interface	+	+	+					
HW interface								
internal actions					+		+	+
inside					+			+
outside							+	



### Kernel Strider





Institute for System Programming of the Russian Academy of Sciences



## Kernel Strider

- KEDR collects data on all memory access from target module and calls to synchronization primitives
  - annotation of Linux synchronization primitives
  - annotation of nonstandard happens-before dependencies in module code
- ThreadSanitizer in Offline mode is used to find data races

https://forge.ispras.ru/projects/kernel-strider



# Test Aspects (1)

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E	RH	KStrider
Monitoring Aspects				-	-	+	+-	+	+-
Kinds of Observable Events									
interface events	+	+	+						
internal events						+	+	+	+
Events Collection									
internal	+	+	+			+			+
external							+		
embedded									
Requirements Specification						Specific	Plugin	Specific	Specific
in-place (local, tabular)	+		+	lf	Dis		Dis		
formal model (pre/post+invar	ants,)	+		lf	Co		Со		
assertions/prohibited events	External	External	External	Co	Co		Со		
Events Analysis									
online	+	+	+						
in-place	+		+			+		+	
outside		+							
offline									+

#### Test Aspects (2)

#### **ISPRAS**

	T2C	OLVER	Autotest	Cfg	FI	KEDR-LC	S2E	RH	KStrider
Active Aspects				+-	+	-	+	+	-
Target Test Situations Set				cfgs				Specifi	\$
requirements coverage	+	+							
class equivalence coverage		+							
model coverage (SUT/reqs)		+							
source code coverage				a	most		+		
Test Situations Setup/Set Gen									
passive								+-	
fixed scenario	+		+						
manual	+								
pre-generated									
coverage driven				+-					
random			+-						
adapting scenario		+							
coverage driven		+							
source code coverage				a	most		+		
model/ coverage		+							
random				as	s opti	on			
TestActions									
application interface	+	+	+						
HW interface									
internal actions					+		+	+	
inside					+			+	
outside							+		



# **Operating Systems**





### **SVACE - Static Analysis**





Institute for System Programming of the Russian Academy of Sciences



## SVACE

- Static analysis of C/C++/Java code, Linux/Windows
- 150+ kinds of defects
  - Buffer overflows, NULL-pointer dereferences
  - Memory management, tainted input
  - Concurrency issues
  - Lightweight analysis of semantic patterns
- Eclipse plugin or WebUI



## **OS** Verification





### Linux Driver Verification





Institute for System Programming of the Russian Academy of Sciences



## Commit Analysis<sup>(\*)</sup>

- All patches in stable trees (2.6.35 3.0) for 1 year:
  - 26 Oct 2010 26 Oct 2011
- 3101 patches overall

(\*) Khoroshilov A.V., Mutilin V.S., Novikov E.M. Analysis of typical faults in Linux operating system drivers. Proceedings of the Institute for System Programming of RAS, volume 22, 2012, pp. 349-374. (In Russian) http://ispras.ru/ru/proceedings/docs/2012/22/isp\_22\_2012\_349.pdf Raw data: http://linuxtesting.org/downloads/ldv-commits-analysis-2012.zip



### **Commit Analysis**

- All patches in stable trees (2.6.35 3.0) for 1 year:
  - 26 Oct 2010 26 Oct 2011
- 3101 patches overall

Unique commits to drivers (1503 ~ **50%**)

Support of a new functionality (321 ~ **20%**)

Bug fixes (1182 ~ **80%**)



### **Commit Analysis**

- All patches in stable trees (2.6.35 3.0) for 1 year:
  - 26 Oct 2010 26 Oct 2011
- 3101 patches overall



#### Taxonomy of Typical Bugs ISPRAS

Rule classes	Types	Number of bug fixes	Percents	Cumulative total percents
	Alloc/free resources	32	~18%	~18%
	Check parameters	25	~14%	~32%
	Work in atomic context	19	~11%	~43%
	Uninitialized resources	17	~10%	~53%
Correct usage of	Synchronization primitives in one thread	12	~7%	~60%
	Style	10	~6%	~65%
AFI (176~50%)	Network subsystem	10	~6%	~71%
(170 % 50 %)	USB subsystem	9	~5%	~76%
	Check return values	7	~4%	~80%
	DMA subsystem	4	~2%	~82%
DMA sub Core driv	Core driver model	4	~2%	~85%
	Miscellaneous	27	~15%	100%
	NULL pointer dereferences	31	~30%	~30%
	Alloc/free memory	24	~24%	~54%
Generic	Syntax	14	~14%	~68%
(102 ~ 30%)	Integer overflows	8	~8%	~76%
	Buffer overflows	8	~8%	~83%
	Uninitialized memory	6	~6%	~89%
	Miscellaneous	11	~11%	100%
Synchronization	Races	60	~85%	~85%
(71 ~ 20%)	Deadlocks	11	~15%	100%



### Software Model Checking

Reachability problem



error location



## Verification Tools World





# Device Driver World

```
int usbpn_open(struct net_device *dev) { ... };
int usbpn_close(struct net_device *dev) { ... };
struct net_device_ops usbpn_ops = {
   .ndo_open = usbpn open, .ndo_stop = usbpn close
};
int usbpn probe(struct usb interface *intf, const struct usb device id *id){
   dev->netdev ops = &usbpn ops;
   err = register netdev(dev);
void usbpn disconnect(struct usb interface *intf){...}
struct usb driver usbpn struct = {
   .probe = usbpn probe, .disconnect = usbpn disconnect,
};
int init usbpn init(void){ return usb register(&usbpn struct);}
void exit usbpn exit(void){usb deregister(&usbpn struct );}
                                                       No explicit calls to
module init(usbpn_init);
                                                       init/exit procedures
module exit(usbpn exit);
```



# Device Driver World

```
int usbpn_open(struct net_device *dev) { ... };
int usbpn_close(struct net_device *dev) { ... };
struct net device ops usbpn ops = {
   .ndo_open = usbpn open, .ndo_stop = usbpn close
};
int usbpn probe(struct usb interface *intf, const struct usb device id *id){
   dev->netdev ops = &usbpn ops;
   err = register netdev(dev);
                                                       Callback interface
                                                       procedures registration
void usbpn_disconnect(struct usb_interface *intf){...}
struct usb_driver usbpn_struct = {
   .probe = usbpn_probe, .disconnect = usbpn_disconnect,
int init usbpn init(void){ return usb_register(&usbpn_struct);}
void exit usbpn exit(void){usb deregister(&usbpn struct );}
                                                      No explicit calls to
module init(usbpn_init);
                                                      init/exit procedures
module_exit(usbpn_exit);
```



# Device Driver World





```
Driver Environment Model
int main(int argc, char* argv[])
{
 usbpn init()
  for(;;) {
    switch(*) {
     case 0: usbpn probe(*,*,*);break;
     case 1: usbpn open(*,*);break;
  usbpn exit();
```



# Driver Environment Model (2)

- Order limitation
  - open() after probe(), but before remove()
- Implicit limitations
  - read() only if open() succeed
- and it is specific for each class of drivers



# Model Checking and Linux Kernel

Reachability problem



error location



## **Error Location?**

```
int f(int y)
```

```
struct urb *x;
```

```
x = usb_alloc_urb(0,GFP_KERNEL);
...
usb_free_urb(x);
```

```
return y;
```



## **Error Location?**

```
int f(int y)
```

```
struct urb *x;
```

```
x = usb_alloc_urb(0,GFP_KERNEL); // allocate new URB
...
usb_free_urb(x); // deallocate URB: assert(x is NULL or previously allocated URB)
return y;
```

// after module exit: assert( all allocated URBs are deallocated)



### Instrumentation

int f(int y)

struct urb \*x;

```
x = usb_alloc_urb(0,GFP_KERNEL);
```

usb\_free\_urb(x);

return y;

```
set URBS = empty;
int f(int y)
{
  struct urb *x;
  x = usb_alloc_urb();
  add(URBS, urb);
...
  assert(contains(URBS, x));
  usb_free_urb(x);
```

```
remove(URBS, urb);
```

```
return y;
```

// after module exit
assert(is\_empty(URBS));



#### **Aspect-Oriented Notation**

// Model state: set of allocated URBs
set URBS = empty;

```
// Model functions
struct urn * ldv_usb_alloc_urb(void)
{
```

```
void *urb;
urb = ldv_alloc();
if (urb) {
   add(URBS, urb);
}
```

```
return urb;
```

```
void ldv_usb_free_urb(struct urb *urb)
```

```
if (urb) {
  remove(URBS, urb);
```

// Pointcut declarations
pointcut USB\_ALLOC\_URB:
 call(struct urb \*usb\_alloc\_urb(int, gfp\_t));
pointcut USB\_FREE\_URB:
 call(void usb\_free\_urb(struct urb \*));

```
// Update model state
around: USB_ALLOC_URB {
  return ldv_usb_alloc_urb();
```

```
around: USB_FREE_URB {
    Idv_usb_free_urb($arg1);
```

```
// Assertions
before: USB_FREE_URB {
   assert(contains(URBS, $arg1));
```

```
after: MODULE_EXIT {
    assert(is_empty(URBS));
```



## Instrumentation

int f(int y)

struct urb \*x;

```
x = usb_alloc_urb(0,GFP_KERNEL);
...
usb_free_urb(x);
```

usb\_free\_urb(x);

return y;

int f(int y)
{
 struct urb \*x;

```
x = ldv_usb_alloc_urb();
```

```
...
assert(contains(URBS, x));
ldv_usb_free_urb(x);
```

```
return y;
```

// after module exit
assert(is\_empty(URBS));



## Rule Instrumentor

- CIF C Instrumentation Framework
  - gcc-based aspect-oriented programming tool for C language
  - available under GPLv3: http://forge.ispras.ru/projects/cif



# Model Checking and Linux Kernel

Reachability problem





### **Linux Driver Verification**





### **SVCOMP'12 Results**

Competition candidate	BLAST 2.7	CPAchecker ABE 1.0.10	CPAchecker Memo 1.0.10	ESBMC 1.17	FShell 1.3	LLBMC 0.9	Predator 20111011	QARMC -HSF	SATabs 3.0	Wolverine 0.5c
Affiliation	Moscow, Russia	Passau, Germany	Paderborn, Germany	Southampton, UK	Vienna, Austria	Karlsruhe, Germany	Brno, Czechia	Munich, Germany	Oxford, UK	Princeton, USA
ControlFlowInteger 93 files, max score: 144	71 9900 s	<b>141</b> 1000 s	140 3200 s	102 4500 s	28 580 s	100 2400 s	17 1100 s	140 4800 s	75 5400 s	<b>39</b> 580 s
DeviceDrivers 59 files, max score: 103	72 30 s	51 97 s	51 93 s	<b>63</b> 160 s	20 3.5 s	80 1.6 s	80 1.9 s		71 140 s	68 65 s
DeviceDrivers64 41 files, max score: 66	55 1400 s	26 1900 s	49 500 s	10 870 s	0 0 s	1 110 s	0 0 s		<b>32</b> 3200 s	16 1300 s
HeapManipulation 14 files, max score: 24		4 16 s	4 16 s	1 220 s		17 210 s	20 1.0 s			
SystemC 62 files, max score: 87	33 4000 s	45 1100 s	<b>36</b> 450 s	<b>67</b> 760 s		8 2.4 s	<b>21</b> 630 s	<b>8</b> 820 s	57 5000 s	36 1900 s
Concurrency 8 files, max score: 11		0 0 s	0 0 s	<b>6</b> 270 s	0 0 s		0 0 s		<b>1</b> 1.4 s	
Overall 277 files, max score: 435	231 15000 s	267 4100 s	280 4300 s	249 6800 s	48 580 s	206 2700 s	138 1700 s	148 5600 s	236 14000 s	159 3800 s


## **SVCOMP'14 Results**

Competition candidate	BLAST 2.7.2	CBMC	CPAchecker	CPAlien	CSeq-Lazy	CSeq-MU	ESBMC 1.22	FrankenBit	LLBMC	Predator	Symbiotic 2	Threader	UFO	Ultimate Automizer	Ultimate Kojak
Representing Jury Member	Vadim Mutilin	Michael Tautschnig	Stefan Löwe	Petr Muller	Bernd Fischer	Gennaro Parlato	Lucas Cordeiro	Arie Gurfinkel	Stephan Falke	Tomas Vojnar	Jiri Slaby	Corneliu Popeea	Aws Albarghouthi	Matthias Heizmann	Alexander Nutz
Affiliation	Moscow, Russia	London, UK	Passau, Germany	Brno, Czechia	Stellenbosch, South Africa	Southampton, UK	Manaus, Brazil	Pittsburgh, USA	Karlsruhe, Germany	Brno, Czechia	Brno, Czechia	Munich, Germany	Pittsburgh, USA	Freiburg, Germany	Freiburg, Germany
BitVectors 49 tasks, max. score: 86		86 2 300 s	78 690 s				77 1 500 s		86 39 s	-92 28 s	39 220 s			-	-23 1 100 s
Concurrency 78 tasks, max. score: 136		128 29 000 s	0 0.0 s		<b>136</b> 1 000 s	136 1 200 s	32 30 000 s		0 0.0 s	0 0.0 s	-82 5.7 s	100 3 000 s			0 0.0 s
ControlFlow 843 tasks, max. score: 1261	508 32 000 s	397 42 000 s	<b>1009</b> 9 000 s	455 6 500 s			949 35 000 s	986 6 300 s	<b>961</b> 13 000 s	511 3 400 s	41 39 000 s		912 14 000 s	164 6 000 s	214 5 100 s
ControlFlowInteger 181 tasks, max. score: 255	64 7 800 s	-298 35 000 s	179 4 800 s	121 3 400 s			85 24 000 s	149 5 300 s	74 10 000 s	-28 2 200 s	-151 22 000 s		184 9 500 s	33 5 800 s	57 5 000 s
Loops 65 tasks, max. score: 99	25 320 s	99 1 100 s	68 600 s	-16 91 s			88 3 600 s	76 50 s	95 160 s	27 14 s	26 4.9 s		44 44 s	26 170 s	29 150 s
ProductLines 597 tasks, max. score: 929	639 24 000 s	918 6 600 s	928 3 500 s	715 3 100 s		-	928 7 500 s	905 950 s	925 2 600 s	929 1 200 s	347 17 000 s		927 4 800 s	0 0.0 s	0 0.0 s
DeviceDrivers64 1428 tasks, max. score: 2766	2682 13 000 s	2463 390 000 s	2613 28 000 s	-			2358 140 000 s	2639 3 000 s	0 0.0 s	50 9.9 s	980 2 200 s		2642 5 700 s		0 0.0 s
HeapManipulation 80 tasks, max. score: 135		132 12 000 s	107 210 s	71 70 s			97 970 s		<b>107</b> 130 s	111 9.5 s	105 15 s				18 35 s
MemorySafety 61 tasks, max. score: 98		4 11 000 s	95 460 s	9 690 s			- <b>136</b> 1 500 s		38 170 s	14 39 s	-130 7.5 s				0 0.0 s
Recursive 23 tasks, max. score: 39		<b>30</b> 11 000 s	0 0.0 s		-		-53 4 900 s		3 0.38 s	-18 0.12 s	6 0.93 s			12 850 s	9 54 s
SequentializedConcurrent 261 tasks, max. score: 364		237 47 000 s	97 9 200 s				244 38 000 s	-	<b>208</b> 11 000 s	-46 7 700 s	-32 770 s		83 4 800 s	49 3 000 s	9 1 200 s
Simple 45 tasks, max. score: 67	30 5 400 s	<b>66</b> 15 000 s	<b>67</b> 430 s				31 27 000 s	37 830 s	0 0.0 s	0 0.0 s	-22 13 s		<b>67</b> 480 s		0 0.0 s
Overall 2868 tasks, max. score: 4718		3 501 560 000 s	<b>2 987</b> 48 000 s				975 280 000 s		<b>1 843</b> 24 000 s	-184 11 000 s	-220 42 000 s			<b>399</b> 10 000 s	139 7 600 s



## **SVCOMP'15** Results

Competition candidate	AProVE	Beagle	BLAST 2.7.3	Cascade	CBMC	CRAchecker	CMrec	ESBMC 1.24.1	FOREST	Forester	Function	HIPTNT+	Lazy-CSeg	Map2Check	MU-CSeg	Perentie	Predator	SeaHorn	SMACK+Corral	Ultimate Automizer	Ultimate Kojak	Unbounded Lazy-CSeg
Representing Jury Member	Thomas Ströder	Dexi Wang	Vadim Mutilin	Wei Wang	Michael Tautschnig	Matthias Dangl	Ming-Hsien Tsai	Jeremy Morse	Pablo Sánchez	Ondrej Lengal	Caterina Urban	Ton-Chanh Le	Gennaro Parlato	Herbert Oliveira Rocha	Bernd Fischer	Franck Cassez	Tomas Vojnar	Arie Gurfinkel	Zvonimir Rakamaric	Matthias Heizmann	Alexander Nutz	Salvatore La Torre
Affiliation	Aachen, Germany	Beijing, China	Moscow, Russia	New York, USA	London, UK	Passau, Germany	Taipei, Taiwan	Bristol, UK	Cantabria, Spain	Bmo, Czechia	Paris, France	Singapore, Singapore	Southampton, UK	Manaus, Brazil	Stellenbosch, South Africa	Sydney, Australia	Bmo, Czechia	Pittsburgh, USA	Salt Lake City, USA	Freiburg, Germany	Freiburg, Germany	Southampton, UK
Anays 86 tasks, max. score: 145	-	-	-	-	-134 2 500 s	2 62 s	-	-206 5.5 s	-	-	-	-	-	-	-	-	-	0 0.61 s	48 400 s	2 64 s	2 5.9 s	-
BitVectors 47 tasks, max. score: 83	-	4 58 s	-	52 16 000 s	68 1 800 s	58 870 s	-	69 470 s	-	-	-	-	-	-	-	-	-	-80 550 s	-	5 170 s	-62 120 s	-
Concurrency 1 003 tasks, max. score: 1 222	-	-	-	-	1 039 70 000 s	0	-	1 014 13 000 s	-	-	-	-	1 222 5 600 s	-	1 222 16 000 s	-	-	- 8 973 42 s	-	-	-	984 36 000 s
ControlFlow 1 927 tasks, max. score: 3 122	-	-	983 33 000 s	537 43 000 s	158 570 000 s	2 317 47 000 s	-	1 968 59 000 s	-	-	-	-	-	-	-	-	-	2 169 30 000 s	1 691 78 000 s	1 887 54 000 s	872 10 000 s	-
ControlPhoentegor 48 Lades, max. score 78	-	-	51 3 300 s	38 11 000 5	62 1 300 6	77 1800 s	-	78 875	-			-		-	-	-	-	77 640 s	61 300 s	78 990 s	43 1 000 5	-
8CA 1140 laks, max some 1874	-		11 9 800 s	0	-2 334 340 000 6	987 39 000 s	-	523 58 000 s	-	-	-			-		-		518 25 000 s	113 45 000 s	632 46 000 s	1 14 s	-
Loops 142 Tasks, max some 215	-	-	44 790.6	46 32 000 5	59 6 900 s	118 2 800 s	-	66 620 5	168 43 000 s	-	-	-	•	-	-	125 1 600 5		130 1 100 s	86 370 6	115 2 900 6	109 4 300 s	-
ProductLines 597 Lado, max.score 929		-	837 19 000 s	0	199 1900 s	901 4 100 s	-	917 430 s	-	-	-						-	913 3 300 s	917 32 000 s	554 3 300 s	87 5 000 s	
DeviceDrivers64 1 650 tasks, max. score: 3 097	-	-	2 736 11 000 s	-	2 293 380 000 s	2 572 39 000 s	-	2 281 36 000 s	-	-	-	-	-	-	-	-	-	2 657 16 000 s	2 507 72 000 s	274 850 s	82 270 s	-
Roats 81 tasks, max. score: 140	-	-	-	-	129 15 000 s	78 5 100 s	-	-12 5 300 s	-	-	-	-	-	-	-	-	-	-164 5.9 s	-	-	-	-
HeapManipulation 80 tasks, max. score: 135	-	-	-	70 6 000 s	100 13 000 s	96 930 s	-	79 37 s	-	32 1.8 s	-	-	-	-	-	-	111 140 s	-37 14 s	109 820 s	84 460 s	84 420 s	-
MemorySafety 205 tasks, max. score: 361	-	-	-	200 82 000 s	-433 14 000 s	326 5 700 s	-	-	-	22 25 s	-	-	-	28 2 100 s	-	-	221 460 s	0	-	95 13 000 s	66 4 800 s	-
Recursive 24 tasks, max. score: 40	-	6 22 s	-	-	0 10 000 s	16 31 s	18 140 s	-	-	-	-	-	-	-	-	-	-	-88 2.3 s	27 2 300 s	25 310	10 220	-
Sequentialized 261 tasks, max. score: 364	-	-	-	-	-171 39 000 s	130 11 000 s	-	193 9 600 s	-	-	-	-	-	-	-	-	-	-59 5 800 s	-	15 8 600 s	-10 7 000 s	-
Simple 46 tasks, max. score: 68	-	-	32 4 200 s	-	51 16 000 s	54 4 000 s	-	29 990 s	-	-	-	-	-	-	-	-	-	65 1 400 s	51 5 100 s	0 1 800 s	3 140 s	-
Termination 393 tasks, max. score: 742	610 5 400 s	-	-	-	-	0 0 s	-	-	-	-	350 61 s	545 300 s	-	-	-	-	-	0	-	565 8 600 s	-	-
Overall 5 803 tasks, max. score: 9 562	-	-	-	-	1 731 1 100 000 s	4 889 110 000 s	-	- 2 161 130 000 s	-	-	-	-	-	-	-	-	-	- 6 228 53 000 s	-	2 301 87 000 s	231 23 000 s	-
			B	'AS	ア	C,	Ach	<sup>ec</sup> ke	37													

BLAST



## **Error Trace Visualizer**

Rule: Mutex lock/unlock

	Erro	or trace		Source code							
🗹 Fi	unction bodies	✓ Blocks	Others	car	l9170.h	<pre>main.c.common.c</pre>	wlan.h	rcupdate.h			
3182 3191 3195 3195 3195 3198 3200 3280	LDV_IN_INTERRU <u>+ldv_initializ</u> tmp8 = nondet assert(tmp8 tmp7 = nondet assert(tmp7 assert(tmp7	<pre>PT = 1; re_FOREACH(); _int() { /* The ! = 0); _int() { /* The ! = 0); ! = 0); ! = 1);</pre>	function body function body	1026 1027 1028 1029 1030 1031 1032	<pre>static int {     struct an     int err =</pre>	carl9170_op_set_key(struct struct ieee80 struct ieee80 struct ieee80 struct ieee80 r9170 *ar = hw->priv; = 0, i:	ieee80211_hw *hw, 211_vif *vif, 211_sta *sta, 211_key_conf *key)	enum set_key_cr≏			
3360	assert(tmp7	!= 2);		1033	u8 ktype;			=			
3440 3520 3600 3680 3760 3840 3920 4000 4080 4130 1031	<pre>assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 assert(tmp7 <u>carl9170_op_s</u> { ar = *(hw) err = 0;</pre>	<pre>!= 3); != 4); != 5); != 5); != 6); != 7); != 8); != 9); != 10); != 11); ret_key(var_groups); .priv;</pre>	= oupl /* hw */	1034 1035 1036 1037 1038 1039 1040 1041 1042 1043 1044 1045 1046	if (ar->c /* * We hav * the us * to mon * * This 1 * the hi */	disable_offload    !vif) return -EOPNOTSUPP; we to fall back to software ser choose to participates i re than one network. is very unfortunate, because igh througput speed in 802.1	encryption, whenev n an IBSS or is co some machines can ln networks.	ver onnected onot handle			
1035 1035 1047 1047 1159 1163	assert(*(ar assert(vif <u>+tmp_7</u> = assert(tmp_ assert(*(ar <u>+mutex_unlo</u> }	).disable_off != 0); is_main_vif(a _7 == 0); ).rx_software ck_mutex(&(ar	<pre>load == 0); r /* ar */, v _decryption )-&gt;mutex /*</pre>	1047 1048 1049 1050 1051 1052 1053	if (!is_n /* * While * group * decide	main_vif(ar, vif)) goto err_softw; the hardware supports *catc key en-/de-cryption. The wa es which keyId maps to which	h-all* key, for of y of how the hardw key, remains a my	floading ware vstery			
<			>	<	*/ III			>			

### **ISP**RAS

# Bugs Found

### http://linuxtesting.org/results/ldv >230 patches already applied

#### Problems in Linux Kernel

This section contains information about problems in Linux kernel found within Linux Driver Verification program.

No.	Туре	Brief	Added on	Accepted	Status	
L0212	Deadlock	nfit: acpi_nfit_notify(): Do not leave device locked	2015-12-11	https://lkml.org/lkml/2015 /12/11/781 commit	Fixed kernel 4.4-rc6	in
L0211	Crash	USB: whci-hcd: no check for dma mapping error	2015-12-01	http://linuxtesting.org /pipermail/ldv-project/2015- November/000558.html commit	Fixed kernel 4.4-rc5	in
L0210	Crash	vmxnet3: fix checks for dma mapping errors	2015-11-28	https://lkml.org/lkml/2015 /11/27/498 commit	Fixed kernel 4.4-rc4	in
L0209	Crash	sound: fix check for error condition of register_chrdev()	2015-11-07	https://lkml.org/lkml/2015 /11/6/914 commit	Fixed kernel 4.4-rc1	in
L0208	Crash	<pre>mcb: Do not return zero on error path in mcb_pci_probe()</pre>	2015-10-28	https://lkml.org/lkml/2015 /10/17/238 commit	Fixed kernel 4.4-rc1	in
L0207	Crash	staging: r8188eu: _enter_critical_mutex() error handling	2015-10-28	https://www.spinics.net/lists /kernel/msg2094451.html commit	Fixed kernel 4.4-rc1	in
L0206	Deadlock	usb: gadget: pch-udc: fix deadlock in pch-udc	2015-09-18	https://lkml.org/lkml/2015 /9/28/256 commit	Fixed kernel 4.4-rc1	in
L0205	Leak	<pre>mcb: leaks in mcb_pci_probe()</pre>	2015-09-16	https://lkml.org/lkml/2015 /7/8/1041	Fixed kernel	in



## **OS** Verification





## **OS** Verification





## **Deductive Verification**





Institute for System Programming of the Russian Academy of Sciences

### **Deductive verification**

Alan Turing — Lecture to the London Mathematical Society 1947

- 1970. R. Floyd / C.A.R. Hoare Methods
- Tools for deductive verification of programs in C, Java, C# 2000
  - SunRise, ESC/Java, Frama-C, LOOP, Boogie/VCC
  - Industrial applications
    - Nuclear power (UK, France)
    - Avionics (Airbus, NASA, UK Air Traffic Control)
    - Operating systems (seL4, PikeOS, Hyper-V)



```
/*@ predicate is divisor(int m, int n) { if (m) then (n % m == 0) else \false } */
/*@ predicate is_gcd(int z, int x1, int x2) {
            is divisor(z,x1)
  0
         && is divisor(z,x2)
  0
         && \forall int i; is divisor(i,x1) && is divisor(i,x2) => (i <= z) }
  0
  @*/
/*@ logic int gcd(int m, int n) */
/*@ axiom gcd_def: \forall int m; \forall int n; is_gcd(gcd(m,n),m,n) */
/* The function returns the greatest common divisor of x1 and x2 */
/*@ requires (x1 > 0) && (x2 > 0)
  @ ensures is gcd(\result,x1,x2)
  @*/
int nod(int x1, int x2)
int y1 = x1;
int y_2 = x_2;
int tmp = 0;
  if (y1 > y2) {
    y1 = x2;
    y^2 = x^1;
  }
  /*@
    @ invariant ((0 < y1 <= y2) && gcd(x1,x2) == gcd(y1,y2))</pre>
             || ((y1 == 0) && gcd(x1,x2) == y2)
    0
    @ variant y1
    @*/
  while (y1 != 0) {
    tmp = y1;
    y1 = y2\%y1;
    y^2 = tmp;
```

### **ISPRAS**

Why3 Interactive Pro	oof Session		En 👣 🖂 📢) 19:57 (共
Context	Theories/Goals	Status Time	21775 integer of int32
Unproved goals	<ul> <li>parsec.mlw</li> </ul>	0	21776 05 = 0 ->
O All cools	Jessie_model	0	21777 (forall us retres:
	Lemma not_less_max_int, lemma	0	21779 us retres =
Provers	Lemma not_more_zero_int, lemma	0	21780 05->
Alt-Ergo (0.95.2)	Lemma and int, lemma	0	21781 (Torall return:
01(22(2.4.4)	🙆 Lemma or_int, lemma	0	21783 return =
CVC3 (2.4.1)	Lemma not_less_max_unsigned_int, lemma	0	21784 us_retres ->
CVC4 (1.2)	Lemma not_more_zero_unsigned_int, lemma	a 🥝	21/85 Integer of Int32
Coc (9 4pl2)	Lemma and unsigned_int, lemma	0	21787 0 /\
Cod (0.4b(2)	Lemma or_unsigned_int, lemma	0	21788 valid_parsec_mac_write_up
PVS (6.0)	Lemma not_less_max_long, lemma	0	21/99 else forall 03:int32.
72 (4 2 1)	Lemma not_more_zero_long, lemma	0	21791 integer of int32 o3 = 0 ->
23 (4.3.1)	Lemma and long, lemma	0	21792 (forall us_retres:int32.
Transformations	Lemma or long, lemma	0	21/93 Us_retres = 03 -> 21794 (forall return: int32.
Split	Lemma not less max unsigned long, lemm	na 🧑	21795 return = us_retres ->
Inline	Lemma not more zero unsigned long, lemm	na 🙆	21796 integer of int32 return = 0 /\
Inune	Lemma and unsigned long, lemma	0	2//9/ Valid_parsec_mac_write_up 2/798 us anonstruct atomic t i parsec mac write up 1 alloc table))))
Tools	Lemma or unsigned long, lemma	0	21799 end
Edit	V V Jessie program	<u>0</u>	21800
Replay	VC for mac_file_permission_ensures_ALLOW	3	<pre>102 int mac_file_permission(const parsec_mac_t *s, const parsec_mac_label_t *o, 103</pre>
riepioy	VC for mac_file_permission_ensures_DEFAUL	и 🥝	104 { 104 {
Cleaning	Alt-Ergo (0.95.2)	2.10	
Remove	<ul> <li>VC for mac_file_permission_ensures_DENY</li> </ul>	0	106 ASSERI(s); 107 ASSERI(o):
Clean	split_goal_wp	0	103
	VC for mac_file_permission_ensures_default	. 🥥	109 if (mask & MAY_READ) {
Proof monitoring Waiting 0	<ul> <li>C for mac_file_permission_safety</li> </ul>	0	110 IT ( (!O-STYPE & MAC ATTR IGNORER LVL) &
Scheduled: 0			112 (!(o->type & MAC_ATTR_IGNORER_CAT) &&
Running: 0			<pre>113 (s-&gt;category &amp; o-&gt;mac.category) != o-&gt;mac.category) )</pre>
Interrupt			114 return - EPEKN; 115 }
			116
			117 if (mask & MAY WRITE) {
			118 int may write up = atomic read(sparsec mac write up); 119 if (mask & MAY WRITEUP) [1 unlikely(may write up)) {
			120 if ( (!(o->type & MAC_ATTR_IGNOREW_LVL) &&
			121 (s-slevel > o-smac.level))
			122 (10-stype of Mac_artin_lowner_car) on $(123)$ (s-scategory) (123) (s-scategory) (123)
			124 return - EPERM;
			125 } else {
			127 (s-2)cycl = o -mac_here[)
			128 (!(o->type & MAC_ATTR_IGNOREW_CAT) &&
			129 (s->category != o->mac.category))
			132 }
			133 134 if (mack & MAY EVEC) /
			135 if (!!o-type & MAC ATTR IGNOREX LVL) &&
			136 (s->level < o->mac.level)
			137 ( $(\circ, \sim)$ type & MAC_ATTR IGNOREX_CAT) & 138 ( $(\circ, \sim)$ category & $(\circ, \sim)$ mac_category) = $(\circ, \sim)$ mac_category)
			rite: /nome/astrauser/workspace/astraver-misc/parsec.c

### **ISPRAS**

# **Astraver Project**



- Deductive Verification of Linux Security Module
  - Joint project with NPO RusBITech
  - Formal security model MROSL-DP
- Assumptions
  - Linux kernel core conforms with its specifications
    - It is not target to prove
- Target code
  - Code is hardware independent
  - Verification unfriendly
  - Out of control



# MROSL DP

- Operating system access control model
  - Hierarchical Role-Based Access Control (RBAC)
  - Mandatory Access Control (MAC)
  - Mandatory Integrity Control (MIC)
- Implemented as Linux Security Module (LSM) for Astra Linux
- ~150 pages in mathematical notation

### AstraVer Project







## Problems with the tools

- Memory model limitations
  - Arithmetics with pointers to fields of structures (container\_of)
  - Prefix structure casts
  - Reinterpret casts
- Integer model problems
- Limited code support
  - Functional pointers
  - String literals
- Scalability problems
- Usability problems



Institute for System Programming of the Russian Academy of Sciences

### VERIFICATION CENTER LINUX



Q Search

▼ C<sup>I</sup>

#### About Us

€

- About Center
- Our Team
- News
- Partners
- Contacts

#### Projects

- Linux Kernel Space Verification
- LSB Infrastructure
- Testing Technologies
- Tests and Frameworks
- Portability Tools
- Contribution
- Publications
- Events

Results

- khoroshilov
- My account
- User list
- Create content
- Feed aggregator
- Administer
- Log out

#### 18-Feb-2015: The first public release of Astraver Toolset

View Edit Track Translate

#### Submitted by Mikhail Mandrykin on Wed, 18/02/2015 - 18:30

We are happy to announce the first public release of **Astraver Toolset 1.0** that is built on top of the '**Frama-C** + **Jessie** + **Why3 IDE**' deductive verification toolchain. The toolchain was adapted, so it can be used to specify and prove properties of Linux kernel code. The most of our modifications go to the Jessie plugin, while the Frama-C front-end and the Why3 platform have got just minor fixes or improvements. Some of our modifications were already applied upstream, while the rest is available in **our public repositories**.

The most important modifications are described below.

#### C Language Support

- Low-level reinterpret type casts between pointers to integral types. This feature required modification of the Jessie memory model as described in our paper "Extended High-Level C-Compatible Memory Model with Limited Low-Level Pointer Cast Support for Jessie Intermediate Language". The overall idea can be summarized as an ability to do certain ghost re-allocations of memory blocks in explicitly specified points in order to transform arrays of allocated objects (structures) from one type to another. WARNING. Discriminated unions support is not yet fully adapted to the modified memory model.
- Prefix type casts between outer structures and their corresponding first substructures (through field inlining and structure inheritance relation in Jessie).
- Kernel memory (de)allocating functions kmalloc()/kzalloc(), kfree().
- Builtin C99 \_\_Bool type.
- Standard library functions memcpy(), memmove(), memcmp() and memset(). The support for these functions is
  implemented through type-based specialization of several pre-defined pattern specifications. (\*)
- Function pointers (through exhaustive may-aliases checking). (\*)
- Variadic functions (through additional array argument). (\*)
- Inline assembly (through undefined function calls). (\*)

(\*)The main purpose of implementing support for these features was the ability to use the tools on our target code without the need for its significant preliminary modification. As a result the support is not complete enough to be usable for verification of code that significantly relies on these features. For instance:



# **OS** Verification





# **Operating Systems**





## Conclusions

- Test Execution System
- Benchmarking

# Thank you!

Alexey Khoroshilov khoroshilov@ispras.ru http://linuxtesting.org/



Institute for System Programming of the Russian Academy of Sciences



## Math



Institute for System Programming of the Russian Academy of Sciences



## Test Results: Details

