

Is It Possible to Automatically Port Kernel Modules?

Yanjie Zhen, Wei Zhang, **Zhenyang Dai**, Junjie Mao, Yu Chen

Tsinghua University

August 28, 2018

Outline

Background


Introduction

Experiment & Results

Discussion

Kernel modules (kmods) are essential

- ▶ A flexible way to extend the functionality of Linux kernel.
- ▶ Over 70% of Linux source codes are kernel modules.



Porting kernel modules is necessary but hard

Porting kmods is necessary:

- ▶ Kernel interfaces (APIs) are not stable.
- ▶ Forward porting: to enjoy enhanced security, boosted performance and new functionality from the fast evolving kernel.
- ▶ Back porting: for stability/compatibility requirements.

But, it is also hard:

- ▶ The Linux kernel is large.
- ▶ It is complex.
- ▶ And it evolves fast.

Existing approaches for porting kernel modules

Manually:

- ▶ A trial-error-fix approach.
- ▶ Using “git log” or Google.

Existing tools:

- ▶ Based on automated program repair techniques: rely heavily on compilation error/warning messages
- ▶ Based on instance-based inference techniques: infer change patterns from source code modifications

Outline

Background

Introduction

Experiment & Results

Discussion

Our work

An empirical study:

- ▶ Focus on the characteristics and representations of kernel API changes.
- ▶ Dataset: 200 patches related to 10 active kmods from Linux repository over the last 7 year.

Discussion:

- ▶ Feasibility and challenges of porting kmods automatically.
- ▶ Our insights in building effective automated tools.

Change patterns

Change patterns, or “patterns” for short, are a core concept of our work.

- ▶ Describe how the usage of an API should be updated.
- ▶ API changes $\xrightarrow{\text{cause}}$ “change instances” where the API is used $\xrightarrow{\text{infer}}$ change patterns.
- ▶ Provide reference for automatically porting.
- ▶ One patch may involve multiple patterns, and most patterns have multiple instances.

In our experiment, we infer change patterns manually. Semantic Patch Language (SmPL) is used as the representation of change patterns.

Change patterns: a concrete example

```
1 #define d_inode(d) ((d)->i)
2
3 -int f2(dentry *d) {
4 +int f2(dentry *d, inode *i){
5 -   if(d_inode(d)->flags & F2) { ... }
6 +   if(i->flags & F2) { ... }
7 -   return d->d_sb->mode;
8 +   return dsb(d)->mode;
9 }
10 struct ops ops1 = { .f1 = f1, .f2 = f2,
11
12 -int f3(dentry *x) {
13 +int f3(dentry *x, inode *i) {
14     ...
15 -   return __f3(x->parent->d_sb);
16 +   return __f3(dsb(x->parent));
17 }
18 struct ops ops2 = { .f2 = f3, };
19
20 -int f4(dentry *d) {
21 +int f4(dentry *d, inode *i) {
22     int ret;
23 -   if (d_inode(d)->flags & F3) { ... }
24 +   if (i->flags & F3) { ... }
25     return ret;
26 }
27 struct ops ops3 = { .f2 = f4, };
```

- ▶ f is a function having an argument d of type $dentry*$. f acts as the $f2$ callback field of some ops struct.
- ▶ For such f , add a new parameter i with type $inode*$, and replace $d_inode(d)->flags$ with $i->flags$.
- ▶ Replace $e->d_sb$ with $dsb(e)$, where e is any expression.

Change patterns: a concrete example

```
1 #define d_inode(d) ((d)->i)
2
3 -int f2(dentry *d) {
4 +int f2(dentry *d, inode *i){
5 -   if(d_inode(d)->flags & F2) { ... }
6 +   if(i->flags & F2) { ... }
7 -   return d->d_sb->mode;
8 +   return dsb(d)->mode;
9 }
10 struct ops ops1 = { .f1 = f1, .f2 = f2,
11
12 -int f3(dentry *x) {
13 +int f3(dentry *x, inode *i) {
14   ...
15 -   return __f3(x->parent->d_sb);
16 +   return __f3(dsb(x->parent));
17 }
18 struct ops ops2 = { .f2 = f3, };
19
20 -int f4(dentry *d) {
21 +int f4(dentry *d, inode *i) {
22   int ret;
23 -   if (d_inode(d)->flags & F3) { ... }
24 +   if (i->flags & F3) { ... }
25   return ret;
26 }
27 struct ops ops3 = { .f2 = f4, };
```

```
1 @rule@
2 identifier s, f;
3 @@
4 struct ops s = { .f2 = f, };
5 @@
6 typedef dentry, inode;
7 identifier rule.f;
8 identifier _arg1, _arg2;
9 @@
10 - f (dentry *_arg1) {
11 + f (dentry *_arg1, inode *_arg2) {
12 <...
13 -   d_inode(_arg1)->flags
14 +   _arg2->flags
15 ...>
16 }
17
18 @@
19 expression E
20 @@
21 - E->d_sb
22 + dsb(E)
23 }
```

Outline

Background

Introduction

Experiment & Results

Discussion

Dataset

- ▶ 10 kernel modules from `fs/` and `drivers/`
- ▶ 1753 patches related to API changes
- ▶ 200 patches chosen randomly due to limited time
- ▶ 407 patterns inferred manually

Target modules	Commits
<code>fs/btrfs</code>	282
<code>fs/ext2</code>	141
<code>fs/ext4</code>	342
<code>fs/jfs</code>	104
<code>fs/xfs</code>	219
<code>drivers/block/drbd</code>	137
<code>drivers/gpu/drm/radeon</code>	291
<code>drivers/net/.../ixgbe</code>	68
<code>drivers/usb/gadget</code>	367
<code>drivers/input/keyboard</code>	130
Total	2081


Questions to answer

Validity of the assumptions on inputs i.e. patches

- ▶ How many change patterns are there in a single commit?
- ▶ How many instances does each pattern have?
- ▶ Does the compiler always complain about out-of-date API usages?


Q1-1: How many change patterns are there in a single commit?

- ▶ 36.5% (73 out of 200) of the commits involve multiple patterns
- ▶ 111 commits (55.5%) contain changes irrelevant to API usage updates → *w irrelevant*



Q1-2: How many instances does each pattern have?


- ▶ 87.5% of patterns (356 out of 407) have multiple instances.



Q1-3: Does the compiler always complain about out-of-date API usages?

When some pattern is missed,

- ▶ 51 patterns (12.5%) change runtime behavior without triggering any compiler message → *runtime*
- ▶ 105 patterns (25.8%) neither cause compiler error/warning nor alter program runtime behavior → *future & unknown*



(c) Does compiler always complain

Questions to answer

Validity of the assumptions on inputs i.e. patches


- ▶ How many change patterns are there in a single commit?
- ▶ How many instances does each pattern have?
- ▶ Does the compiler always complain about out-of-date API usages?

Representation of changes (within 356 patterns with multiple instances)

- ▶ What kinds of code are changed by the patterns?
- ▶ How often do the patterns involve macro invocations?
- ▶ What kinds of contexts do the patterns require?


Q2-1: What kinds of code are changed by the patterns?

- ▶ 86.5% of the patterns (308 out of 356) touch function definitions. → *fun (intra)* & *fun (inter)*
- ▶ 6.4% of the patterns (23 out of 356) update definitions of global variables. → *global*



Q2-2: How often do the patterns involve macro invocations?

- ▶ 33.4% (119 out of 356) of the patterns involve macro invocations. → *both & deleted & inserted*




Change patterns: a concrete example

```
1 #define d_inode(d) ((d)->i)
2
3 -int f2(dentry *d) {
4 +int f2(dentry *d, inode *i){
5 -   if(d_inode(d)->flags & F2) { ... }
6 +   if(i->flags & F2) { ... }
7 -   return d->d_sb->mode;
8 +   return dsb(d)->mode;
9 }
10 struct ops ops1 = { .f1 = f1, .f2 = f2,
11
12 -int f3(dentry *x) {
13 +int f3(dentry *x, inode *i) {
14   ...
15 -   return __f3(x->parent->d_sb);
16 +   return __f3(dsb(x->parent));
17 }
18 struct ops ops2 = { .f2 = f3, };
19
20 -int f4(dentry *d) {
21 +int f4(dentry *d, inode *i) {
22   int ret;
23 -   if (d_inode(d)->flags & F3) { ... }
24 +   if (i->flags & F3) { ... }
25   return ret;
26 }
27 struct ops ops3 = { .f2 = f4, };
```

```
1 @rule@
2 identifier s, f;
3 @@
4 struct ops s = { .f2 = f, };
5 @@
6 typedef dentry, inode;
7 identifier rule.f;
8 identifier _arg1, _arg2;
9 @@
10 - f (dentry *_arg1) {
11 + f (dentry *_arg1, inode *_arg2) {
12 <...
13 -   d_inode(_arg1)->flags
14 +   _arg2->flags
15 ...>
16 }
17
18 @@
19 expression E
20 @@
21 - E->d_sb
22 + dsb(E)
23 }
```

Q2-3: What kinds of contexts do the patterns require?

- ▶ 36.5% (130 out of 356) of changes are context dependent. → *from the second to the last column*
- ▶ Callbacks (81 out of 356) and global variables (21 out of 356) account for the majority (78.5%). → *callback & global*



Outline

Background

Introduction

Experiment & Results

Discussion

Discussion

Feasibility:

- ▶ Based on compilation messages: severely limited
- ▶ Based on pattern inference: promising

Challenges:

- ▶ Noises: multiple patterns in a single commit; modification of irrelevant code
- ▶ Macros: infer pattern with macros
- ▶ Contexts: especially callbacks and global variables

Thank you!
More in the paper!
Any questions?