Increasing Automation in the Backporting of Linux Drivers Using Coccinelle

Luis R. Rodriguez, Julia Lawall

► To cite this version:

Luis R. Rodriguez, Julia Lawall. Increasing Automation in the Backporting of Linux Drivers Using Coccinelle. 11th European Dependable Computing Conference - Dependability in Practice, Nov 2015, Paris, France. hal-01213912

HAL Id: hal-01213912

<https://hal.inria.fr/hal-01213912>

Submitted on 12 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Increasing Automation in the Backporting of Linux Drivers Using Coccinelle

Luis R. Rodriguez

Rutgers University/SUSE Labs

mcgrof@winlab.rutgers.edu

mcgrof@suse.com, mcgrof@do-not-panic.com

Julia Lawall

Sorbonne Universités/Inria/UPMC/LIP6

Julia.Lawall@lip6.fr

Abstract—Software is continually evolving, to fix bugs and add new features. Industry users, however, often value stability, and thus may not be able to update their code base to the latest versions. This raises the need to selectively backport new features to older software versions. Traditionally, backporting has been done by cluttering the backported code with preprocessor directives, to replace behaviors that are unsupported in an earlier version by appropriate workarounds. This approach however involves writing a lot of error-prone backporting code, and results in implementations that are hard to read and maintain. We consider this issue in the context of the Linux kernel, for which older versions are in wide use. We present a new backporting strategy that relies on the use of a backporting compatibility library and on code that is automatically generated using the program transformation tool Coccinelle. This approach reduces the amount of code that must be manually written, and thus can help the Linux kernel backporting effort scale while maintaining the dependability of the backporting process.

Keywords—Linux, backports, program transformation

I. INTRODUCTION

Linux is an open-source operating system kernel that has been under development since 1991. Its reliability, customizability, and low cost have made it a popular choice for an operating system kernel across the computing spectrum, from smartphones and tablets based on the Android distribution, to desktops based on popular distributions such as Debian and Ubuntu, to supercomputers based on Enterprise Linux releases. The Linux kernel is evolving rapidly, with a major release roughly every 2.5 months. Between the recent major releases v3.14 and v3.19 (January 2014 - February 2015), on average, 8,400 lines of code were added, 5,300 lines of code were removed, and 2,100 lines of code were modified *every day*.¹ This rapid rate of change combined with frequent releases allows the Linux kernel to keep up to date with bug fixes, new functionalities, and new services, such as new CPU architectures, device drivers and filesystems.

A. The dilemma for silicon manufacturers

While the rapid evolution of the Linux kernel has many benefits, it can be problematic for certain classes of users. Some system integrators may have invested heavily in testing a specific release and may wish to avoid regressions due

to a kernel upgrade. Upgrading a kernel may also require experience to understand what features to enable, disable, or tune to meet existing deployment criteria. In the worst case, some systems may rely on components that have not yet been merged into the mainline Linux kernel, potentially making it impossible to upgrade the kernel without cooperation from the component vendor or a slew of partners that need to collaborate on developing a new productized image for a system. As an example, development for 802.11n AR9003 chipset support on the mainline ath9k device driver started on March 20, 2010 with an early version of the silicon, at which point the most recent major release of the Linux kernel was v2.6.32. One of the first products to ship with this driver was the Google Chrome CR48, using ChromeOS, which started selling in retail in May 2011. The latest kernel release at this point was v2.6.38, but ChromeOS was still based on the v2.6.32 kernel, the release for which it was originally developed.

The reluctance of users to keep up to date with the latest kernel poses a dilemma for silicon manufacturers, who need to make available device drivers so that their devices can be used on Linux systems. One approach is to develop drivers explicitly for the kernel releases that their clients are currently using. However, even clients who value stability may eventually need to modernize. Doing so then incurs the burden of a full rewrite or port of each driver to the newly adopted kernel release. And even once the driver is successfully ported to a more modern kernel, the result will only be usable by those who are currently using the same kernel release.

An alternative to targeting a device driver to a specific kernel version is *upstream-first* development, in which code is initially developed only for the latest major kernel release, and then is submitted for inclusion *upstream*, i.e., into the Linux git repository maintained by Linus Torvalds, allowing the device driver to be included in the coming major release. While achieving inclusion upstream can be a challenge for silicon manufacturers, due to the strict coding guidelines of the Linux kernel, once it is achieved, the developers at the silicon company that upstreamed the device driver can then benefit from help from the Linux community in reviewing changes to the device driver as the Linux kernel evolves [1]. The silicon manufacturer needs to contribute only one complete version of the code; since it is upstream it will then be part of all future kernel releases and all future Linux distributions.

¹<https://github.com/gregkh/kernel-development/raw/6cd82e40742008a0189fc2e57476928bd011d20f/kernel-development.pdf>

B. Why and how Linux is backported

The upstream-first model makes device drivers available for users of future kernels, but leaves out those who must remain with older releases. If a device driver is only supported upstream, a Linux distribution or system integrator has no option but to backport that device driver down to the kernel release of interest. Backporting strategies have typically consisted of augmenting each affected file with `#ifdefs` that handle what is required for each kernel release on each target file. As each file is augmented individually, there is no code sharing, even within a single Linux distribution’s codebase. Files also become harder to read, as the original code is interspersed with new code flows required to support each kernel release.

Since 2007, the Linux kernel backports project has promoted an alternative backport strategy, with the goal of maximizing code sharing and minimizing disruption to the individual driver source files, and with the goal of enabling upstream-first development of new drivers and features by making backports of the resulting code available to everyone, regardless of the Linux distribution used. The main innovation is to move the changes required to backport each driver out of the individual driver files and into a *backports* library, providing a set of helper functions. Indeed, typically, for a given class of devices, the drivers use a similar set of API functions and coding strategies, and these functions and coding strategies can all be backported in the same way. Rather than distributing the handling for each older release in every driver file, all of these variations are encapsulated into a backports library function. This approach was used in the ath9k support for ChromeOS noted above. The ath9k device driver was extended with AR9003 family chipset support upstream, and this support was incorporated as part of the v2.6.38 release at the time of the release of the Google Chrome CR48. Support for the AR9003 family of chipsets on ath9k on ChromeOS, however, was backported onto the ChromeOS v2.6.32 based kernel using the Linux backports library.

The use of a backports library can dramatically reduce the amount of code changes required to support a class of drivers. Nevertheless, some changes per driver are still needed, amounting to *glue code*, to invoke the library functions and to *e.g.*, modify type definitions, which cannot be encapsulated into a function definition. These changes must initially be made manually in each supported file to create *patches*,² which are made available to users, and these patches need to be maintained as the kernel evolves. Making these changes and maintaining the resulting patches is tedious and error prone, and limits the number of drivers that the backports project can support. A solution was thus needed to ease and improve the dependability of the process of introducing this glue code.

C. Our contributions

We report on a new methodology for backports adopted by the Linux backports project that combines the use of a backports library with the use of Coccinelle [3], [4]. Coccinelle is a program matching and transformation tool for C code that

has been specifically designed to meet the needs and requirements of Linux developers. Matches and transformations are described in terms of an extension of the patch notation with generic features, resulting in *semantic patches*. Unlike standard patches, which are restricted to specific positions in specific files, a semantic patch expresses a change in a generic way, allowing it to be applied across an entire code base and to adapt to minor changes in this code base, as the code base evolves over time. In the context of backports, we automate the integration of the glue code into each supported driver using Coccinelle. Now the glue code needs to be specified only once, and then this specification can be applied automatically to all supported drivers over multiple releases, thus further reducing the amount of maintenance work to support a backport, while increasing the dependability of the backport process.

The main contributions of this paper are as follows:

- We show that in practice the changes required in drivers to support the backport backports library are often systematic. Indeed, currently, 70 (50%) of the 140 patches distributed by the backports project affect multiple files, but do so in a similar way, and thus are candidates for automation.
- We show that the Coccinelle transformation language is expressive enough to describe these changes in a driver-independent way. Specifically, we are able to replace the 631 additions and 112 deletions lines involved in the common changes by 22 Coccinelle *semantic patch* specifications, consisting of 562 lines of code.³
- We show that use of Coccinelle improves the dependability of the backport development process, by guiding developers to designing backports in a more orthogonal way and by ensuring that backport changes are done consistently.

D. Overview

The rest of this paper is organized as follows. Section II provides background, including the relevant aspects of the Linux development model, the history of the backports project, and the use of Coccinelle. Section III then presents a tour of backporting strategies in more detail, based on a simple example. Next, Section IV expands this tour to illustrate a more complex case study, in which the changes required are determined by driver-specific information. Section V then highlights the benefits of our approach and considers some correctness and performance issues. Finally, we present related work, conclusions, and future directions. This paper emphasizes how each new backporting strategy has helped the backports project to grow and scale, and ultimately how each of these strategies has helped to shift the objective of the project away from simply backporting the Linux kernel, towards trying to automate the backport process as much as possible.

II. BACKGROUND

We first briefly review the starting points of our work: the Linux kernel development model, the Linux backports project, and the Coccinelle program transformation tool.

²A patch is a document indicating the lines of added and removed code, in a format generated by the Unix command `diff`. A patch can be automatically applied to a file using the Unix command `patch` [2].

³The added and deleted lines represent the cumulated effect of the patches generated from the semantic patches. This number is computed using the command `diffstat`. The number of lines in the semantic patches is computed using `wc -l` after removal of blank lines and comments.

Fig. 1: Four phases of Linux development

A. The Linux kernel development model

The need for backporting is intrinsically linked to the nature and frequency of changes in the Linux kernel. Thus, we begin by describing the Linux kernel development model, which controls when new code becomes available and when code changes start to affect users.

As illustrated in Fig. 1, a Linux kernel release is developed in four phases: the *merge window*, the *release candidate evaluation*, the *major release*, and the *maintenance of a stable kernel*. The merge window begins immediately after the previous major release, and lasts for 1-2 weeks. During this period, maintainers who have accumulated new features and device drivers for their subsystems request that Linus Torvalds collect and integrate their changes. This period culminates in the release of the first release candidate (rc1) making the complete set of merged changes available for general testing. The release of rc1 initiates the release candidate evaluation period. The purpose of this period is to find and fix regressions for new code merged since the last major kernel release. There may be 5-9 release candidates, one per week. New drivers may be integrated during the release candidate evaluation period. Next comes the major release, making public the new version. At this point, the merge window for the next release begins. In parallel with the merge window, and beyond, possibly for several years, the current release is maintained as a stable version. Stable releases only contain critical bug fixes; they never contain new features or drivers. Bug fixes are always sent upstream first, and must be merged into Linus Torvalds' tree before being cherry picked or ported to older maintained stable releases.

The Linux kernel development model introduces some delay across the release candidate evaluation period between when new features are developed and when they are integrated into a release candidate or major release. To compensate for this delay, the *linux-next*⁴ tree mimics the merge window each day, by pulling from each subsystem tree. Each day the tree is reset to the latest major kernel release and then each subsystem tree is pulled. The *linux-next* tree can thus be used to track the latest development efforts on all subsystems on a daily basis.

B. A brief history of the Linux kernel backports project

The Linux kernel backports project⁵ was started in 2007 by Luis R. Rodriguez while at Rutgers University, to help backport the 802.11 subsystem and a series of 802.11 device drivers to a series of older kernel releases.⁶ The project was

originally referred to as *compat-wireless*, reflecting the initial target. Over the years, the project grew to support more device drivers and subsystems. In April 2012, the project was folded under the Linux Foundation backports working group.⁷ The project now spearheads the Linux kernel backports effort. In April 2013, the project was renamed *backports*, to distinguish it from the Linux kernel compat layer, which addresses 64-bit and 32-bit compatibility. The backports project is now led by three core developers: Hauke Mehrtens, Johannes Berg, Luis Rodriguez and two co-maintainers: Mehrtens and Rodriguez. Git logs going back to September 2008 show 80 unique contributors. The project backports the subsystems Ethernet, Wireless, Bluetooth, NFC, ieee802154, Media, and Regulator. The backports project provides support for backporting both driver C code and configuration code. In this paper we focus on the backporting of driver C code.

The current goal of the backports project is to backport a slew of device drivers from the latest major kernel releases down to a series of supported stable kernel releases, at a minimum those listed on the main kernel website, kernel.org. Currently, 22 releases are supported. The project's master development branch always tracks *linux-next*, allowing it to track all the development trees. This ensures that at the end of each merge window, the state of the backports will be very close to the state of the first release candidate. At this point, the backports project creates a further branch that tracks the progress of the new release over the release candidate evaluation period, to the major release, and on to its lifetime as a stable kernel. The backports project thus makes three kinds of backports available: those derived from *linux-next*, those derived from the most recent release candidate if any, and those derived from recent stable kernels. A user may prefer a backport from a stable kernel to one from *linux-next* or from a release candidate, if one is available for the desired driver.

As shown in Fig. 2, as of June 2015, the backports project backports almost 800 drivers.⁸ These are kept up to date with *linux-next* and the recent stable kernels each day, and are guaranteed to at least compile correctly. Ensuring this each day typically requires 2-6 iterations of test, refinements, and compiles for all supported versions. For this development, the backports project uses a 32-core system with 236 GiB of RAM. Code generation and compile tests are all run in memory. As measured by GNU `time`, a full compilation test of a release across all 22 supported kernel revisions takes approximately 53 minutes of real time, 1440 minutes of user mode time and 219 minutes of kernel time. When the backports project began in 2007, it provided backporting support for drivers down to v2.6.25, first released in 2008. In order to scale to a wider range of drivers, however, the project now only supports kernels down to at most Linux v3.0, first released in 2011. The original motivation behind the project was to encourage silicon manufacturers to work upstream on the Linux kernel while providing them with a solution for backporting their drivers automatically down to older releases. The framework is designed *only* for Linux upstream drivers; the associated license enforces that proprietary drivers cannot and should not use this framework.

⁴[git://git.kernel.org/pub/scm/linux/kernel/git/next/linux-next.git](http://git.kernel.org/pub/scm/linux/kernel/git/next/linux-next.git)

⁵<https://backports.wiki.kernel.org>

⁶[git://git.kernel.org/pub/scm/linux/kernel/git/mcgrof/compat-wireless-2.6-old.git](http://git.kernel.org/pub/scm/linux/kernel/git/mcgrof/compat-wireless-2.6-old.git)

⁷http://lists.linuxfoundation.org/pipermail/lf_driver_backport/2012-August/001075.html

⁸The Linux v4.2 number is approximate, as Linux v4.2 is not yet released.

Fig. 2: Number of backported drivers, for kernels released since 2009

To better understand the usage of the code provided by the Linux kernel backports project, we conducted an informal survey of developers via the Linux kernel mailing list and the backports mailing list.⁹ Current users include the OpenWRT Linux distribution for embedded devices,¹⁰ and Google Fiber.¹¹ We are also aware of several major Linux enterprise distributions that are considering integration of backports for their future products. Other developers cited the need for backporting in their work, although they currently do backports in an ad hoc manner. Their experiences suggest that other subsystems may benefit from the methodology developed by the backports project. Between April 26 and May 29, 2015, the backports project website received an average of 4218 “GET” requests per day, with 13633 requests on May 28 alone. While we do not know who initiated these requests, nor how they use the resulting information, it is another sign of the overall interest in the work of the backports project.

C. Coccinelle

Coccinelle is a program matching and transformation engine for C code [3], [4]. It is released as open source, and is available in a number of Linux distributions. Coccinelle provides a scripting language, SmPL (*semantic patch language*), that allows patterns to be expressed as fragments of C code, and transformations to be expressed by annotating lines of code with `-`, for removal of the matched code, or `+`, for addition of the corresponding code. As such, SmPL specifications resemble patches [2]. Nevertheless, they are more robust than ordinary patches in that they are insensitive to comments and whitespace, and take into account aspects of the code semantics such as control flow and type information. Thus, we refer to SmPL specifications as *semantic patches*.

As a simple example of the use of Coccinelle, we consider the problem of replacing each call to a function `one` by a call to a function `two`, and adding a `NULL` argument. The semantic patch that makes this change is shown in Fig. 3.

```
1 @@
2 expression arg;
3 @@
4 - one(arg)
5 + two(arg, NULL)
```

Fig. 3: A simple semantic patch

This semantic patch consists of a single rule that makes the complete transformation. The rule has two parts: the declaration of *metavariables* that can match any term of a specified type, between the initial pair of `@@` (lines 1-3), followed by a *transformation specification* (lines 4-5). In this case, the only metavariable is `arg`, which is declared to match any expression. The transformation specification then removes the call to `one` with its argument, while at the same time binding the metavariable `arg` to this argument expression. It then constructs a new call to `two` with arguments the current binding of `arg` and `NULL`. This semantic patch can be applied to an entire code base, removing the tedious and error prone task of searching for and transforming each relevant call manually, and thus improve the dependability of the transformation process. A more detailed presentation of Coccinelle is available in previous work [3], [4] and at the Coccinelle website.¹²

Coccinelle was originally motivated by a study of how the Linux kernel evolves [5]. This study identified the problem of *collateral evolutions*, in which the interface of a library changes, and all clients of the library must be updated accordingly. Coccinelle was designed to help Linux developers make collateral evolutions faster and more reliably. The problem of implementing backports using Coccinelle is related to the problem of collateral evolutions. While collateral evolutions were envisioned as being needed to modernize software, backports must address changes in library interfaces to transport modern code to older versions.

The Linux backports project is widely used over a few industries, making its dependability critical. Thus, the project’s maintainers were initially concerned about the effect of replacing manually created and checked patches by an automatic tool. Coccinelle was chosen for initial experiments because it has already been extensively used on Linux kernel code, and because its patch-like notation implied that existing patch sets could be easily converted to Coccinelle semantic patch specifications. Conversely, the similarity between standard patches and Coccinelle semantic patches can help developers read and maintain the semantic patches over time, thus further improving the dependability of the backporting process.

III. BACKPORTING STRATEGIES

We now review two existing backporting strategies: i) the common strategy of providing version-specific implementations using `#ifdefs`, and ii) the strategy originally taken by the Linux backports project of factorizing these changes into a compatibility library. We then present our extension of the latter using Coccinelle.

⁹backports@vger.kernel.org, linux-kernel@vger.kernel.org, <http://lkml.kernel.org/r/CAB=NE6UUTmMMB0La3OF+hXpaZzdAwTs0QKBsTLsYrsSMFxrJg@mail.gmail.com>

¹⁰<https://openwrt.org/>, <http://lkml.kernel.org/t/556CA991.7030302@openwrt.org>

¹¹<https://fiber.google.com/about/>, <https://gfiber.googleusercontent.com/vendor/opensource/backports>

¹²<http://coccinelle.lip6.fr>

```

1 --- a/drivers/net/usb/usbnet.c
2 +++ b/drivers/net/usb/usbnet.c
3 @@ -1151,6 +1151,7 @@
4  }
5  EXPORT_SYMBOL_GPL(usbnet_disconnect);
6
7 +#if (LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,29))
8  static const struct net_device_ops usbnet_netdev_ops = {
9      .ndo_open          = usbnet_open,
10     .ndo_stop           = usbnet_stop,
11 @@ -1160,6 +1161,7 @@
12     .ndo_set_mac_address = eth_mac_addr,
13     .ndo_validate_addr   = eth_validate_addr,
14 };
15 #endif
16
17 /*-----*/
18
19 @@ -1229,7 +1231,15 @@
20     net->features |= NETIF_F_HIGHDMA;
21 #endif
22
23 +#if (LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,29))
24     net->netdev_ops = &usbnet_netdev_ops;
25 #else
26 + net->change_mtu = usbnet_change_mtu;
27 + net->hard_start_xmit = usbnet_start_xmit;
28 + net->open = usbnet_open;
29 + net->stop = usbnet_stop;
30 + net->tx_timeout = usbnet_tx_timeout;
31 #endif
32     net->watchdog_timeo = TX_TIMEOUT_JIFFIES;
33     net->ethtool_ops = &usbnet_ethtool_ops;

```

Fig. 4: Backporting the `net_device_ops` collateral evolution for the `usbnet` driver

A. Backporting with `#ifdefs`

Backporting a device driver typically consists of modifying the source code with `#ifdefs` to handle the different requirements of different kernel releases. This entails adding blocks of code that provide alternate implementations for various functionalities, for different ranges of kernel versions, according to which evolutions have occurred and which collateral evolutions must be performed to accommodate them.

As a running example, we consider an evolution that was introduced by Linux kernel commit d314774cf2 (“netdev: network device operations infrastructure”)¹³ and that was first merged upstream in Linux v2.6.29. This evolution moved a series of callback functions from the `net_device` data structure out into a new separate data structure of type `net_device_ops`. Backporting over this evolution, for Linux kernel versions before Linux v2.6.29, requires a collateral evolution that *undoes* this change. We refer to this collateral evolution as the `netdev_ops` collateral evolution. Specifically, the definition of the new `net_device_ops` structure and the initialization of the link from the `net_device` structure to this new structure via the `netdev_ops` field must be restricted, using an `#ifdef`, to the versions starting with Linux v2.6.29. Earlier versions must initialize the appropriate fields in the `net_device` structure itself, from among the callback functions that the modern driver puts in the new structure. Fig. 4 shows a patch that backports this single collateral evolution on one device driver.

Lines 7 and 23 of Fig. 4 introduce the `#ifdefs` that restrict

some code of the modern driver to be used only in Linux versions v2.6.29 and later. Lines 25-31 introduce the code to be used for earlier versions, placing the relevant callback functions from the modern code (lines 8-13) into the fields of the `net` structure. In each case, the callback functions are used by the driver support library of the kernel, which is not backported, and which thus finds the desired functions in the expected place, with no further code changes.

In general, every driver that initializes a `net_device` structure would require all of these changes. Creating these patches, and maintaining them as other collateral evolutions are needed, is complex, tedious, and error prone.

B. Backports via a compatibility library

Maintenance of patches is easy as long as the amount of changes being introduced is rather small. The `netdev_ops` collateral evolution, however, is an example of a collateral evolution that affects many network drivers, resulting in a large set of changes, that then have to be maintained in patch form. A better approach, proposed by the Linux backports project, consists of wrapping up the required changes into static inline or external helper functions and then using `#ifdefs` in these functions to adapt the code to each previous release.

This strategy is illustrated by the following code, which backports the `netdev_ops` collateral evolution for two device drivers. Now, the new `net_device_ops` structure used by the modern driver remains in the code as is. Instead, we replace the direct initialization of the `netdev_ops` field by a call to a single static inline function defined by the backports compat library, amounting to glue code. Now, only this function needs multiple lines of `#ifdef` code, performing the direct assignment for versions starting with Linux v2.6.29, and copying the fields from the new structure into the main `net_device` structure for older versions. Only one line of code is changed in each driver, in contrast to the 10 lines added to each driver by the previous approach.

```

1 --- a/drivers/net/usb/usbnet.c
2 +++ b/drivers/net/usb/usbnet.c
3 @@ -1446,7 +1446,7 @@ usbnet_probe (struct usb_interface *
4     udev
5     net->features |= NETIF_F_HIGHDMA;
6     #endif
7
8 - net->netdev_ops = &usbnet_netdev_ops;
9 + netdev_attach_ops(net, &usbnet_netdev_ops);
10    net->watchdog_timeo = TX_TIMEOUT_JIFFIES;
11    net->ethtool_ops = &usbnet_ethtool_ops;
12
13 --- a/drivers/net/wireless/ath/ath6kl/main.c
14 +++ b/drivers/net/wireless/ath/ath6kl/main.c
15 @@ -1289,7 +1289,7 @@ static const struct net_device_ops
16     ath6k
17
18 void init_netdev(struct net_device *dev)
19 {
20     dev->netdev_ops = &ath6kl_netdev_ops;
21 + netdev_attach_ops(dev, &ath6kl_netdev_ops);
22     dev->destructor = free_netdev;
23     dev->watchdog_timeo = ATH6KL_TX_TIMEOUT;

```

Between 2007 and 2013 the backports project exclusively followed this strategy to help reduce the amount of maintenance of patches. The backports library now has a large set of helper functions that help minimize the number and size of the patches required for each backported driver.

¹³git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=d314774cf2cd5dfcb39a00d37deec65d4c627927

C. The Coccinelle way to backport

The backports library approach reduces significantly the amount of code that must be modified in each driver. Still, backporting a new driver requires identifying the set of features that it uses, and comparing these features to those provided by the backports library to see where a collateral evolution to replace the existing code by glue code is needed. Our observation is that the required changes are often systematic, and thus, just as Coccinelle had been found to be useful in automating traditional (forward) collateral evolutions, we propose that Coccinelle can also be useful for the kinds of collateral evolutions required in backports. For example, the `netdev_ops` collateral evolution needed in each driver could be expressed as a Coccinelle semantic patch as follows:

```
1 @@
2 struct net_device *dev;
3 struct net_device_ops ops;
4 @@
5 - dev->netdev_ops = &ops;
6 + netdev_attach_ops(dev, &ops);
```

This semantic patch backports the `netdev_ops` collateral evolution for all networking device drivers. It is indeed no longer even necessary for the developer to identify whether a new device driver to backport uses this features; Coccinelle both finds and updates the relevant code automatically. Note that the semantic patch specifies the type of the expressions matching the `dev` and `ops` metavariables, to ensure that the transformation is performed only on structures of the appropriate type. Finally, this semantic patch amounts to only 6 lines of code to maintain, rather than 2 lines of code for each driver with the `#ifdef` approach.

IV. CASE STUDY

To test the limits of what can be backported using Coccinelle, we chose the most complex collateral evolution ever supported by the backports project as a test case. Specifically, we decided to try to backport threaded IRQ support, introduced in the v2.6.31 kernel. This backport requires modifications to a driver-specific structure, as well as to multiple driver functions. Note that this semantic patch is no longer part of the Linux kernel backports release, as the backports project currently supports only kernels down to Linux 3.0. Still, it illustrates the most complex backport case that we have encountered.

A. Backporting threaded IRQ support the old way

We first explain how the backports project provided backport support for threaded IRQ support prior to using Coccinelle. The first step was to extend the backports library with support for threaded IRQs, as shown in Fig. 5. This involved creating a new data type `compat_threaded_irq` (lines 1-11) to collect some extra information for each driver, and creating a set of helper functions to implement the threaded IRQ functionality (lines 13-69). The helper functions include `compat_request_threaded_irq` (lines 26-44), which initializes the fields of the `compat_threaded_irq` structure and then calls `request_irq`, and functions such as `compat_free_threaded_irq` (lines 46-51) and `compat_synchronize_threaded_irq` (lines 62-68) that call their unthreaded counterparts on information stored in

the `compat_threaded_irq` structure. The extension to the backports library amounts in all to 75 lines of code.¹⁴

```
1 #if LINUX_VERSION_CODE < KERNEL_VERSION(2,6,31)
2 struct compat_threaded_irq {
3     unsigned int irq;
4     irq_handler_t handler;
5     irq_handler_t thread_fn;
6     void *dev_id;
7     char wq_name[64];
8     struct workqueue_struct *wq;
9     struct work_struct work;
10 };
11 #endif
12
13 #if LINUX_VERSION_CODE < KERNEL_VERSION(2,6,31)
14 static inline
15 void compat_irq_work(struct work_struct *work)
16 {
17     ...
18 }
19
20 static inline
21 irqreturn_t compat_irq_dispatcher(int irq, void *dev_id)
22 {
23     ...
24 }
25
26 static inline
27 int compat_request_threaded_irq(
28     struct compat_threaded_irq *comp,
29     unsigned int irq,
30     irq_handler_t handler,
31     irq_handler_t thread_fn,
32     unsigned long flags,
33     const char *name,
34     void *dev_id)
35 {
36     comp->irq = irq;
37     comp->handler = handler;
38     comp->thread_fn = thread_fn;
39     comp->dev_id = dev_id;
40     INIT_WORK(&comp->work, compat_irq_work);
41     ...
42     return request_irq(irq, compat_irq_dispatcher, flags,
43         name, comp);
44 }
45
46 static inline
47 void compat_free_threaded_irq(
48     struct compat_threaded_irq *comp)
49 {
50     free_irq(comp->irq, comp);
51 }
52
53 static inline
54 void compat_destroy_threaded_irq(
55     struct compat_threaded_irq *comp)
56 {
57     if (comp->wq)
58         destroy_workqueue(comp->wq);
59     comp->wq = NULL;
60 }
61
62 static inline
63 void compat_synchronize_threaded_irq(
64     struct compat_threaded_irq *comp)
65 {
66     synchronize_irq(comp->irq);
67     cancel_work_sync(&comp->work);
68 }
69 #endif
```

Fig. 5: Extensions to the backports library to support request_threaded_irq

Each driver to backport that relies on threaded IRQs then

¹⁴Computed using David Wheeler's SLOCCount, <http://www.dwheeler.com/sloccount/>.

```

1 --- a/drivers/net/wireless/b43/b43.h
2 +++ b/drivers/net/wireless/b43/b43.h
3 @@ -805,6 +805,9 @@ enum {
4
5  /* Data structure for one wireless device (802.11 core) */
6  struct b43_wldev {
7  +#if LINUX_VERSION_CODE < KERNEL_VERSION(2,6,31)
8  +    struct compat_threaded_irq irq_compat;
9  +#endif
10     struct b43_bus_dev *dev;
11     struct b43_wl *wl;
12     /* a completion event structure needed if this call
13        is asynchronous */
14 --- a/drivers/net/wireless/b43/main.c
15 +++ b/drivers/net/wireless/b43/main.c
16 @@ -4243,8 +4243,17 @@ redo:
17     if (b43_bus_host_is_sdio(dev->dev)) {
18         b43_sdio_free_irq(dev);
19     } else {
20  +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
21         synchronize_irq(dev->dev->irq);
22  +#else
23 +    compat_synchronize_threaded_irq(&dev->irq_compat
24 +);
25  +#endif
26  +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
27         free_irq(dev->dev->irq, dev);
28  +#else
29         compat_free_threaded_irq(&dev->irq_compat);
30         compat_destroy_threaded_irq(&dev->irq_compat);
31  +#endif
32     }
33     mutex_lock(&wl->mutex);
34     dev = wl->current_dev;
35 @@ -4290,9 +4299,17 @@ static int b43_wireless_core_start(
36     goto out;
37 }
38 } else {
39  +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
40     err = request_threaded_irq(dev->dev->irq,
41         b43_interrupt_handler,
42         b43_interrupt_thread_handler,
43         IRQF_SHARED, KBUILD_MODNAME, dev);
44  +#else
45     err = compat_request_threaded_irq(
46         &dev->irq_compat, dev->dev->irq,
47         b43_interrupt_handler,
48         b43_interrupt_thread_handler,
49         IRQF_SHARED, KBUILD_MODNAME, dev);
50  +#endif
51     if (err) {
52         b43err(dev->wl, "Cannot request IRQ-%d\n",
53             dev->dev->irq);

```

Fig. 6: Backporting the b43 driver

needs to be modified to make use of the new helper functions. Fig. 6 shows the modifications for the b43 driver. Lines 1-13 update a header file to extend the driver’s private `b43_wldev` structure type with a field containing the `compat` structure when the kernel version is lower than the first one that supports threaded irqs, *i.e.*, Linux v2.6.31. Lines 14-52 replace each threaded IRQ operation with its `compat` version, again for kernels for which threaded irqs are not already supported.

The changes shown in Fig. 6 amount to 20 lines of added code, over half of which are `#if`defs, and only apply to a single driver. Note that the `#if`defs cannot be factorized into the backports library, because they involve code that references an added field of a driver-specific structure. As of the linux-next of April 2, 2015, 170 files contain at least one call to `request_threaded_irq`, and of these 16 are

in subsystems supported by the Linux backports project.¹⁵ Backporting all of the 170 files that use threaded IRQs to Linux versions prior to v2.6.31 would require developing and maintaining over 3000 lines of patch code.

B. Backporting threaded IRQ support with Coccinelle

Fig. 7 shows a Coccinelle semantic patch that automates these changes. Most rules are fairly trivial: replace one call with another, with the new call using the backport data structure among its arguments. A typical example is illustrated in the first rule (lines 1-24), where a call to `request_threaded_irq` is replaced by a call to `compat_request_threaded_irq`. The new call takes the same arguments as the old one, with the addition of the first argument (line 17), which is the `compat_threaded_irq` structure.

A challenge in implementing this backport is where to store the `compat_threaded_irq` structure. As a running kernel may include multiple instances of a device, this structure cannot simply be a global variable of the device driver. For the b43 driver, we placed this structure into the driver’s existing `b43_wldev` structure. To generalize this approach, we need to find a suitable location for this structure in each driver to which the semantic patch may be applied.

The need to support multiple instances of a data structure at run time is a common problem in device driver development, and the Linux kernel proposes a standard solution, the use of a *private* data structure. An instance of this private structure is created when a device is initialized, and then the driver infrastructure makes this structure available to each driver callback function, much like the implicit “this” argument found in object-oriented languages. Normally, each driver defines a specific private-structure type, containing the information that is specific to its operation. We exploit this private structure to store the `compat_threaded_irq` structure.

To use the private structure to store the `compat_threaded_irq` structure, we must address two issues. First, for each driver, we must find the type of the private structure and extend the corresponding type declaration with a field for the `compat_threaded_irq` structure. Second, we must find the name of the current instance of the private structure at each point where a reference to the `compat_threaded_irq` structure is needed for the backport process.

To address the first issue, we exploit the fact that Coccinelle collects type information when analyzing the source code, and makes it possible to manipulate this type information via metavariables in a semantic patch. Fortunately, device drivers typically already pass their private structure as the last argument to `request_threaded_irq`, as the information contained in the private structure is typically also needed by the interrupt handler, which is installed by this function. By matching this reference to the private structure, Coccinelle makes it possible to obtain its type. Concretely, line 5 of Fig. 7 declares a type metavariable `T`, which is then used in describing the type of metavariable `private`. Matching `private` against the code in the last argument of `request_threaded_irq` has the side effect of storing the type of the matched code in

¹⁵1 in ethernet, 7 in wireless, 0 in bluetooth, 2 in nfc, 0 in ieee802145, 4 in media, 2 in regulator.


```

1 @ threaded_irq @
2 identifier ret;
3 expression irq, irq_handler, irq_thread_handler, flags,
4         name;
5 type T;
6 T *private;
7 @@
8
9 +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
10 ret = request_threaded_irq(irq,
11         irq_handler,
12         irq_thread_handler,
13         flags,
14         name,
15         private);
16 +#else
17 +ret = compat_request_threaded_irq(&private->irq_compat,
18         irq,
19         irq_handler,
20         irq_thread_handler,
21         flags,
22         name,
23         private);
24 +#endif
25
26 @ sync_irq depends on threaded_irq @
27 expression irq;
28 type threaded_irq.T;
29 T *threaded_irq.private;
30 @@
31
32 +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
33 synchronize_irq(irq);
34 +#else
35 +compat_synchronize_irq(&private->irq_compat);
36 +#endif
37
38 @ free depends on threaded_irq @
39 expression irq, dev;
40 type threaded_irq.T;
41 T *threaded_irq.private;
42 @@
43
44 +#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,31)
45 free_irq(irq, dev);
46 +#else
47 +compat_free_threaded_irq(&private->irq_compat);
48 +compat_destroy_threaded_irq(&dev->irq_compat);
49 +#endif
50
51 @ modify_private_header depends on threaded_irq @
52 type threaded_irq.T;
53 @@
54
55 T {
56 +#if LINUX_VERSION_CODE < KERNEL_VERSION(2,6,31)
57 +     struct compat_threaded_irq irq_compat;
58 +#endif
59 ...
60 };

```

Fig. 7: Backporting threaded IRQs with Coccinelle

T, where it can be used by subsequent rules. In the fourth rule (lines 51-60), T, referenced as `threaded_irq.T`, is used to match and extend the definition of the private structure, adding a new field `irq_compat` to the beginning of the private structure when the kernel version is less than v2.6.31.

To address the second issue, we exploit the fact that, within a given driver, the Linux developers typically give the private structure the same name, in every function in which it is used. Thus, we simply inherit the term matched by the metavariable `private` defined in the rule `threaded_irq`, and use that term in the added calls in the `sync_irq` and `free` rules. This solution is not safe, but it is pragmatic, in that it simplifies

the transformation and exploits properties of the Linux coding style. In the `free_irq` case (lines 38-49), we could also use the second argument to `free_irq`, which by definition of the IRQ API should point to the same structure as the last argument to `request_threaded_irq`.¹⁶ This value is not immediately available in the `sync_irq` rule (lines 26-36), but we could extend the rule to match a neighboring expression of the right type, for a safer solution.

V. DISCUSSION

We now analyze the benefits of using Coccinelle for backporting, and then consider some correctness and performance issues raised by the use of Coccinelle.

A. Benefits of using Coccinelle for backporting

Reimplementing the threaded IRQ backport using Coccinelle revealed that the original manual backport was inconsistent. Specifically, in the manual backport, the `compat` structure, `compat_thread_irq`, was integrated into different kinds of structures in different drivers. Backporting is intrinsically risky, because the older code may not respect the invariants required by the backported code. Backporting the code in a consistent way reduces the set of issues that can arise. Doing so also makes the results easier for developers to understand. Reimplementing the threaded IRQ backport using Coccinelle also revealed that the existing threaded IRQ patch series also backported another collateral evolution, related to the management of IRQ flags. Isolating each collateral evolution in a separate patch benefits the backports project, by making it easier to understand how to backport other drivers, which may need only one of the changes. Using Coccinelle not only makes the need for this split apparent, it also makes it easier to manage the resulting set of changes. While two sets of changes are now needed, each is performed by a single semantic patch that can be applied to many files, rather than having to implement and record each of the changes individually.

More generally, we define two metrics of efficiency, *development efficiency* and *maintenance efficiency*. For development efficiency, we start with the number of insertions and deletions that a semantic patch generates, ignoring context information, as reported by `diffstat`, and take the ratio of this number with the size of the semantic patch, exclusive of comments and whitespace. The number of insertions and deletions represents the number of manual changes required when modifying the code. Development efficiency thus represents the initial coding savings induced by using semantic patches. For maintenance efficiency, we compute the same ratio, but this time consider the complete size of the patch, not only the insertions and deletions, but also all the metadata information contained within the patch generated by the semantic patch, including file names, file offsets, and (unmodified) context lines; all of this metadata must also be kept up to date so that the patch command can apply the patch to the relevant files. A development (resp., maintenance) efficiency value of 1 means the semantic patch has the same number of lines as the changes (resp., lines) in the patch series it replaces. A development (resp., maintenance) efficiency value of 2 means the semantic

¹⁶Actually, unintentionally, in the second call that is added by this rule, this strategy is used.

Fig. 8: Development and maintenance efficiency for the 22 semantic patches used in Linux v4.2 (projected). The semantic patches are ordered by their age in the backports project.

patch is producing twice as many changes (resp., patch lines) as the number of lines in the semantic patch.

Fig. 8 shows the development and maintenance efficiency of each semantic patch planned for Linux v4.2. As maintenance efficiency takes into account all of the same patch lines as development efficiency, plus the metadata, the maintenance efficiency is always greater than the development efficiency. A number of the semantic patches have a development efficiency of less than one. 12 of our 22 semantic patches affect only one file, and the need to declare metavariables in a semantic patch increases the semantic patch size, as compared to the number of lines in a single change. Still all of our 22 semantic patches have a maintenance efficiency greater than or equal to 1. Any variance in the maintenance efficiency, whether an increase or decrease from one release to another, implies that if semantic patches were not used manual work would have had to be done to account for new code additions or removals from a patch. Developers would have to manually make such changes for each new release of Linux. In contrast, a semantic patch can adapt to small changes in the source code that are not relevant to the backport, and thus the same semantic patch can be used for multiple versions. Furthermore, some existing semantic patches may also be found to be applicable when backporting support is added for new drivers, thus increasing the ease and dependability of the backport process.

B. Correctness of Coccinelle and the semantic patches

For the traditional uses of Coccinelle, for bug finding and collateral evolutions [3], [6], a developer runs Coccinelle once on a code base, checks and possibly adjusts the results, and submits a patch upstream for review by kernel maintainers. The patch is integrated into the Linux kernel just like a manually generated patch, and Coccinelle is no longer involved. The use case for backports is, however, rather different. Here, the developer typically starts with a collection of patches, reflecting the result of backporting a number of drivers by hand. The developer then generalizes the existing patches into a semantic patch, which is then applied every day, as linux-next and the various stable kernels evolve. In this context, manually studying each result is not practical, and would likely not be reliable. Thus, it is necessary to account for the possibility of errors in the automatically generated result, either in the

semantic patch definition or in Coccinelle itself. Indeed, some improvements to Coccinelle were required to enable its use for backporting, such as improving the support for adding `#ifdefs` around complete function definitions. In this section, we examine some of the issues that affect, positively and negatively, the dependability of the semantic patch application process, and thus the dependability of our backporting approach.

Coccinelle has been designed according to pragmatic goals, balancing ease of use and expressiveness with the need for correctness. As such, it is primarily based on matching of syntax, augmented with control-flow and type information. The only correctness check that is performed is to ensure that at the top-level, terms are replaced by terms of the same kind, ensuring that syntactic well-formedness is preserved. Thus, Coccinelle’s analysis engine is unaware of semantic issues such as variable values and interprocedural effects, and may potentially generate incorrect code as a result. In practice, however, due to the nature of the glue code involved in backports, transformations dependent on variable values and interprocedural effects rarely arise. Indeed, most of the rules depend only on the names of functions, structure fields, and types, all of which are global and constant. Among the 22 semantic patches currently used by the backports project, only three contain a rule that involves matching and then non-locally reusing names of local variables. Aside from these three rules, there should be essentially no risk of making a semantically incorrect transformation. The three rules that raise some risk have furthermore not yet posed any problem in practice.

Besides the danger of making a semantically incorrect transformation, there is also the risk of not making a transformation where one is needed. Several factors work together to help mitigate this risk. First, the drivers to which backporting is applied are already part of the mainline Linux kernel. They have undergone a rigorous review process before being accepted into the kernel, and this review process has a tendency to uniformize the coding style, choice of API functions, etc. These properties were indeed the basis of the initial design of Coccinelle [3], which was found to be able to completely express 93% of 62 representative collateral evolutions from Linux 2.5 and 2.6, involving over 5800 device driver files. Since a backport is typically just the inverse of such a collateral evolution, we expect to achieve a similar success rate, or even better due to the factorizing of the complex parts of the backports into the backports library. Finally, in practice the changes in backported code are mostly motivated by compiler errors identified when trying to use new code or enabling new device drivers in an old release of the Linux kernel. If a needed semantic patch does not apply to a given driver, the resulting driver is likely to trigger a compiler error, and the problem will be quickly detected.

The effects of transformations of collateral evolutions expressed with semantic patches are test compiled prior to a backports release – typically daily. Each release is then tested by various backports users. Thus, incorrect changes are likely to be detected quickly, and to be addressed by changes in the semantic patch. We can thus use the number of changes in the semantic patches over time as a proxy for their dependability. The 22 semantic patches currently used by the backports project range in age from 1 year and 3 months (5 semantic patches) to 1-4 months (17 semantic patches). In this

time, only 5 semantic patches, all but one originating from 1 year and 3 months ago, have been revised. In two cases, the changes involve removing two `#ifdefs` and `#endifs` per semantic patch, in one case the change involves adding an `#ifdef` and the corresponding `#endif`, in one case, the change involves refactoring the semantic patch to use typed expressions to improve its robustness, and the remaining case fixed a type name. These changes resulted in a total of 7 inserted lines and 21 removed lines, out of a total of 562 lines of semantic patch code. This low rate of change, to rules that are applied every day to evolving code, further testifies to the dependability of the Coccinelle-based backporting process. In contrast, the 70 standard patches currently maintained by the backports patches have on average undergone 3.61 changes after their initial creation, and five of these patches have changed 15 or more times.

While we argue that the semantic patches currently maintained by the backports project are dependable, developing a semantic patch amounts to a programming task, and as for other programming tasks, the developer may not produce a completely correct specification immediately. To help in understanding a backporting problem and debugging the semantic patch, the developer can manually create some standard patches that address the specific issue. Indeed, some may be available already, if there has been a prior effort to backport the code. We have developed a simple tool that applies the manually developed patches and the semantic patches to separate copies of the Linux kernel source code, and then checks that the results are equal, to provide a check on the correctness of the semantic patch.

The latter checking methodology only applies when the source code to which the semantic patch is applied is the same as the source code for which the standard patch has been designed. It could be possible to gain further confidence in the correctness of a semantic patch via regression testing. However, in our case, the relevant input, *i.e.*, the current version of version of a driver in a Linux kernel release or release candidate, in `linux-next`, or in a stable tree, is a moving target, changing weekly or even daily. Thus, the result of backporting a driver at one point in time is not likely to be identical to the result of backporting the most recent version of the same driver at a later time. Nevertheless, it could be possible to design a regression test system that checks only the parts of the code to which the semantic patches apply. We leave this to future work.

C. Performance of Coccinelle

Each day, the Linux kernel backports project generates and compile tests backport releases for all supported kernels, of which there are currently 19. Generation for all of them takes around 1.5 minutes on our 32-core machine, and compile testing of the 19 resulting patched kernels takes 44 minutes of real time, comprising 1234 minutes of user time and 147 minutes of system time. Given the long compilation time, it is important to minimize the time for generating backports. Several solutions were explored to avoid the use of Coccinelle overwhelming the backporting time. Figs. 9 and 10 relate the backport generation time for each version, according to the method used at the time of the release of that version, to the

Fig. 9: Generation time compared to the number of patches, the number of semantic patches, and the number of supported kernel versions in the backports releases for various Linux kernel versions

Fig. 10: Number of backported drivers, for recent kernels

number of supported kernels, patches, and resulting backported drivers for recent kernels.

In the first two versions shown, v3.11 and v3.12, the backports project did not use Coccinelle, resulting in a high number of standard patches, that had to be maintained individually. The backport generation time, however, is low, below 50 seconds, because patches are applied efficiently, in a single-threaded manner, on a line-by-line basis, without parsing the code. In version v3.13, three semantic patches were introduced. Coccinelle is designed to apply only one semantic patch at a time. The first strategy taken to apply multiple semantic patches was to concatenate them, and apply the resulting semantic patch in a single thread. This approach saves the cost of starting up Coccinelle for each semantic patch, but the need to parse the semantic patch and the source code entails a high backport generation time. This approach furthermore does not exploit the parallelism inherent in applying a single specification to the hundreds of drivers supported by the backports project. Coccinelle supports static parallelism via an external script that starts up n instances of Coccinelle and causes each of them to process $1/n$ of the files in the code base. With this optimization, semantic patch support can be less expensive than the sequential application of the corresponding patches. For example, for v3.16.2, the number of supported drivers increases significantly (Fig. 10), but the generation time actually goes down slightly.

In the last two versions considered, Linux v4.1 and v4.2,

the semantic patch application time has risen considerably, with a similar increase in the number of semantic patches supported, from 5 in v4.0 to 22 currently in the upcoming v4.2. Further optimizations may be possible. One possible optimization is to use index information, precollected by a tool such as *glimpse*,¹⁷ to identify the subset of files that may be relevant to a given semantic patch. While *Coccinelle* already supports the use of *glimpse*, library incompatibility issues have prevented its use with the backports project. Instead, the backports project uses *Coccinelle*'s internal support for a form of *grep* that scans files for tokens known to be essential to the application of the semantic patch. Another possible optimization is to reduce the need to parse header files, either by sharing the parsed code between the application of different semantic patches or by ignoring header files when the semantic patch does not require the information, such as type information, that they may contain. Finally, dynamic scheduling of parallel threads has recently become possible within *Coccinelle* via the OCaml library *Parmap* [7]. The use of dynamic scheduling may reduce the risk of load imbalances.

D. Semantic patch design issues

Normally, the developer can write the semantic patch rules in any order, as long as the information required to perform a transformation is available. In some cases, however, obtaining reasonable performance requires taking into account some details of how *Coccinelle* applies semantic patches, which may not be obvious for a kernel developer. Specifically, *Coccinelle* applies the rules of a semantic patch in order, from first to last, with any changes specified by a rule being performed as soon as the rule is matched. This can lead to extra matching and undesired modifications if a semantic patch starts with rules that perform generic transformations. Consider the following rule that begins a semantic patch for backporting PCI drivers:

```
1 @ simple_dev_pm depends on module_pci @
2 identifier ops, pci_suspend, pci_resume;
3 declarer name SIMPLE_DEV_PM_OPS;
4 declarer name compat_pci_suspend;
5 declarer name compat_pci_resume;
6 @@
7 +compat_pci_suspend(pci_suspend);
8 +compat_pci_resume(pci_resume);
9 SIMPLE_DEV_PM_OPS(ops, pci_suspend, pci_resume);
```

This rule collects some information from each `SIMPLE_DEV_PM_OPS` structure, and introduces some PCI-specific calls from the backports library. The rule applies to every `SIMPLE_DEV_PM_OPS` declaration, and is not specific to PCI drivers in any way. Not only does it transform code that should not be transformed, it also can potentially have a significant performance impact. In the linux-next files backported as of June 26, 2015, 32 files contain a `SIMPLE_DEV_PM_OPS` declaration, while only 18 of these are PCI drivers. All of these extra files will be parsed and (unnecessarily) transformed if the semantic patch is written in this way.

A solution is to add an initial rule that matches against some other more specific term that must be present if any transformation is needed. Subsequent rules can then depend on the success of matching this rule. Essentially, this rule acts as a “needle in a haystack” to find the source files where the

transformation is actually relevant. In this particular case, we observe that PCI drivers contain a `MODULE_DEVICE_TABLE` declaration with `pci` as the first argument. We thus add the following rule at the beginning of the semantic patch:

```
1 @ module_pci @
2 declarer name MODULE_DEVICE_TABLE;
3 identifier pci_ids;
4 @@
5 MODULE_DEVICE_TABLE(pci, pci_ids);
```

The `SIMPLE_DEV_PM_OPS` rule can now be declared to depend on the rule `module_pci`, ensuring that it is only applied to PCI drivers:

```
1 @ simple_dev_pm depends on module_pci @
2 identifier ops, pci_suspend, pci_resume;
3 declarer name SIMPLE_DEV_PM_OPS;
4 declarer name compat_pci_suspend;
5 declarer name compat_pci_resume;
6 @@
7 +compat_pci_suspend(pci_suspend);
8 +compat_pci_resume(pci_resume);
9 SIMPLE_DEV_PM_OPS(ops, pci_suspend, pci_resume);
```

The final rule to perform the backport uses metavariables that are defined by the previous one, and thus by that rule's dependence will also only be applied to PCI drivers:

```
1 @@
2 identifier backport_driver;
3 expression pm_ops;
4 fresh identifier backports_pci_suspend = simple_dev_pm.
    pci_suspend ## "_compat";
5 fresh identifier backports_pci_resume = simple_dev_pm.
    pci_resume ## "_compat";
6 @@
7 struct pci_driver backport_driver = {
8 +#if (LINUX_VERSION_CODE >= KERNEL_VERSION(2,6,29))
9     .driver.pm = pm_ops,
10 +#elif defined(CONFIG_PM_SLEEP)
11 +     .suspend      = backports_pci_suspend,
12 +     .resume       = backports_pci_resume,
13 +#endif
14 };
```

VI. RELATED WORK

The specific problem of backporting has not received much attention in the research community. Backporting is, however, related to more general issues of change management, as arise when merging trees in a source code management system and when integrating changes developed for one branch of a software product line into another branch.

Uquillas Gómez et al. [8] propose visualization tools to aid the developer in integrating a patch developed for one branch of a software project into another branch of the software project. They focus on individual changes and on helping the developer to identify semantic issues that may affect the correctness of the change in the new context. Other work on change impact analysis includes that of Gallagher and Lyle [9], who use program slicing [10] to collect information about the impact of a change, and the tool Chianti, which identifies change impact using test cases [11]. The work on change impact is complementary with ours. In our case, the correct backport is already identified, and we are concerned with expressing it in a concise and robust way. In the future, we could combine change impact analysis with our approach, to further check the correctness of the backported code.

Fiuczynski et al. faced the challenge of keeping an externally maintained patchset up to date with the evolutions in

¹⁷<http://webglimpse.net/>. As a side effect of this work, the first author convinced the developers of *glimpse* to release *glimpse* as open source.

the Linux kernel [12]. They proposed a preliminary solution based on aspect-oriented programming [13] to re-express these patches in a more generic and robust way. To the best of our knowledge, this tool remained in a prototype stage.

Integrated development environments, such as Eclipse, are able to perform certain classes of refactorings [14], *i.e.*, semantics-preserving code transformations, automatically, and to record sequences of refactorings for subsequent replay on other code. The transformations that we require in the context of backporting, however, are not general purpose, but instead rely on the specific semantics of the functions of the backports compatibility library. Thus, the general-purpose refactorings provided by IDEs are typically not sufficient for the kinds of transformations that are required.

VII. CONCLUSIONS AND FUTURE WORK

The Linux backports project currently, as of commit 065a5d39 (“backports-update-manager: bump 4.1-rc1 to 4.1-rc8”)¹⁸, supports 22 semantic patches totaling 562 lines of SmPL code. These semantic patches affect 63 files on linux-next next-20150612 (preparation for v4.2). The corresponding standard patch they generate contains 2558 lines of code, including change lines and context information, implying that the SmPL patches provide a code maintenance savings of 47.40%. Still, some glue code is implemented by direct modification to the driver code. Using a semantic patch is appropriate when the changes are complex, are relevant to many drivers, and are susceptible to be affected by other evolutions in the code, but it may be overkill when a specific change is required in only one place, *i.e.*, when the maintenance efficiency is near 1. Overall, the use of Coccinelle has contributed to the dependability of the backport process. Indeed, another developer on the backports project recently stated “All the patches that broke often in the early days are now using coccinelle or are removed because they were only needed for the older kernel versions.”¹⁹

Our work on backports raises a number of directions for future work. One direction would be to reduce the need for glue code by integrating upstream the needed static inline functions for accessing and updating key data structures. Patches to address this have been submitted and have now started to be accepted, at least on the networking subsystem, specifically to help reduce the amount of work to backport the ieee802154 subsystem.²⁰ One such change was merged as part of the v3.18-rc1 release.

Another direction would be to infer semantic patches. For many of our backports, we have a collection of manually written patches that make the same change. Backporting could be further streamlined by inferring semantic patches from these examples. Preliminary work has indeed been done on the automatic inference of change specifications [15], [16]. Alternatively, we observe that our additions of glue code amount to (the inverses of) collateral evolutions. If a library change is accompanied by a semantic patch, to ease updating the library’s clients, then it might be possible to systematically

invert this semantic patch for subsequent backporting. Another direction would be to infer the glue code itself.

Finally, it is always a concern that a code change may break semantic invariants. We leave as future work to investigate whether change impact analysis, as described above, can be relevant here. Backporting is also relevant to other kinds of infrastructure software, where users may rely on older versions, but need to use more recent modules. We leave the exploration of whether it is beneficial to use Coccinelle to express the kinds of changes needed in such software to future work.

A. Acknowledgements

The original work on backports was funded by Rutgers University. Backports using Coccinelle was initially supported by funding from Inria and the IRILL. Ongoing research in this area is also supported by SUSE Labs. Finally, we would also like to thank the ongoing contributors to the backports project, in particular Johannes Berg and Hauke Mehrtens.

REFERENCES

- [1] G. Kroah-Hartman, “The Linux kernel driver interface,” Linux 3.17: Documentation/stable_api_nonsense.txt, section: “What to do”.
- [2] D. MacKenzie, P. Eggert, and R. Stallman, *Comparing and Merging Files With Gnu Diff and Patch*. Network Theory Ltd, Jan. 2003, Unified Format section, http://www.gnu.org/software/diffutils/manual/html_node/Unified-Format.html.
- [3] Y. Padiou, J. Lawall, R. R. Hansen, and G. Muller, “Documenting and automating collateral evolutions in Linux device drivers,” in *EuroSys* 2008. Glasgow, Scotland: ACM, Mar. 2008, pp. 247–260.
- [4] J. Brunel, D. Doligez, R. R. Hansen, J. Lawall, and G. Muller, “A foundation for flow-based program matching using temporal logic and model checking,” in *POPL*, Jan. 2009, pp. 114–126.
- [5] Y. Padiou, J. L. Lawall, and G. Muller, “Understanding collateral evolution in Linux device drivers,” in *EuroSys*, Apr. 2006, pp. 59–71.
- [6] J. L. Lawall, J. Brunel, N. Palix, R. R. Hansen, H. Stuart, and G. Muller, “WYSIWIB: exploiting fine-grained program structure in a scriptable API-usage protocol-finding process,” *Software: Practice and Experience*, vol. 43, no. 1, pp. 67–92, Jan. 2013.
- [7] M. Danelutto and R. Di Cosmo, “A “minimal disruption” skeleton experiment: Seamless map & reduce embedding in OCaml,” *Procedia Computer Science*, vol. 9, pp. 1837–1846, 2012.
- [8] V. Uquillas Gómez, S. Ducasse, and T. D’Hondt, “Visually supporting source code changes integration: The Torch dashboard,” in *WCRE*, Beverly, MA, USA, Oct. 2010, pp. 55–64.
- [9] K. B. Gallagher and J. R. Lyle, “Using program slicing in software maintenance,” *Transactions on Software Engineering*, vol. 17, no. 18, pp. 751–761, Aug. 1991.
- [10] M. Weiser, “Program slicing,” in *ICSE*, 1981, pp. 439–449.
- [11] X. Ren, F. Shah, F. Tip, B. Ryder, and O. Chesley, “Chianti: A tool for change impact analysis of Java programs,” in *OOPSLA*, Vancouver, BC, Canada, Oct. 2004, pp. 432–448.
- [12] M. Fiuczynski, R. Grimm, Y. Coady, and D. Walker, “Patch (1) considered harmful,” in *10th Workshop on Hot Topics in Operating Systems (HotOS X)*, Santa Fe, NM, Jun. 2005.
- [13] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold, “An overview of AspectJ,” in *ECOOP*, ser. LNCS, no. 2072, Budapest, Hungary, Jun. 2001, pp. 327–353.
- [14] M. Fowler, *Refactoring: Improving the Design of Existing Code*. Addison Wesley, 1999.
- [15] J. Andersen and J. L. Lawall, “Generic patch inference,” *Automated Software Engineering*, vol. 17, no. 2, pp. 119–148, Jun. 2010.
- [16] N. Meng, M. Kim, and K. S. McKinley, “LASE: locating and applying systematic edits by learning from examples,” in *ICSE*, 2013, pp. 502–511.

¹⁸git.kernel.org/cgit/linux/kernel/git/backports/backports.git/commit/?id=065a5d394ff78dffbe32a78227f5544d77f779d

¹⁹Hauke Mehrtens, private email of October 23, 2014.

²⁰lkml.kernel.org/r/1397784176-15809-2-git-send-email-mcgruf@do-not-panic.com