

# Automatic OS Kernel TCB Reduction by Leveraging Compile-Time Configurability

Reinhard Tartler<sup>1</sup>, Anil Kurmus<sup>2</sup>,  
Bernhard Heinloth<sup>1</sup>, Valentin Rothberg<sup>1</sup>, Andreas Ruprecht<sup>1</sup>, Daniela Dorneanu<sup>2</sup>,  
Rüdiger Kapitza<sup>3</sup>, Wolfgang Schröder-Preikschat<sup>1</sup>, and Daniel Lohmann<sup>1</sup>

<sup>1</sup>*Friedrich-Alexander University Erlangen-Nürnberg*

<sup>2</sup>*IBM Research - Zurich*

<sup>3</sup>*TU Braunschweig*

## Abstract

The Linux kernel can be a threat to the dependability of systems because of its sheer size. A simple approach to produce smaller kernels is to manually configure the Linux kernel. However, the more than 11,000 configuration options available in recent Linux versions render this a demanding task. We report on designing and implementing the first automated generation of a workload-tailored kernel configuration and discuss the security gains such an approach offers in terms of reduction of the Trusted Computing Base (TCB) size. Our results show that the approach prevents the inclusion of 10% of functions known to be vulnerable in the past.

## 1 Introduction

The Linux kernel is a commonly attacked target. In 2011, 148 Common Vulnerabilities and Exposures (CVE)<sup>1</sup> entries for Linux have been recoded, and this number is expected to grow every year. This is a serious problem for system administrators who rely on a distribution-maintained kernel for the daily operation of their systems. On the Linux distributor side, kernel maintainers can make only very few assumptions on the kernel configuration for their users: Without a specific use case, the only option is to enable every available configuration option to maximize the functionality. The ever-growing kernel code size, caused by the addition of new features, such as drivers, file systems and so on, indicates that the risk of undetected vulnerabilities will constantly increase in the foreseeable future.

If the intended use of a system is known at kernel compilation time, an effective approach to reduce the kernel's attack surface is to configure the kernel to not compile unneeded functionality. However, finding a fitting configuration requires extensive technical expertise about currently more than 11,000 Linux configuration options,

<sup>1</sup><http://cve.mitre.org/>

and needs to be repeated at each kernel update. Therefore, maintaining such a custom-configured kernel entails considerable maintenance and engineering costs.

This paper presents a tool-assisted approach to automatically determine a kernel configuration that enables only kernel functionalities that are actually necessary in a given scenario. We quantify the security gains in terms of reduction of the Trusted Computing Base (TCB) size. The evaluation section (Section 3) focuses on an appliance-like virtual machine that runs a web server similar to those used to power large distributed web services in the cloud. Our approach exhibits promising security improvements for this use case: Compared with a default distribution kernel, 10% of the kernel functions (i.e., 17 out of 179), for which in total 31 vulnerabilities have been reported, are removed from the tailored kernel.

The remainder of this paper is structured as follows: Section 2 presents the design and implementation of the first automated workload-specific kernel-build generation tool. Section 3 evaluates the usability of such an approach in a real-world scenario. Security benefits of the tailored Linux kernel are discussed in Section 4. Section 5 presents the related work. The paper concludes in Section 6.

## 2 Kernel-Configuration Tailoring

The goal of our approach is to compile a Linux kernel with a configuration that has only those features enabled which are necessary for a given use case. This section shows the fundamental steps of our approach to tailor such a kernel. The six necessary steps are depicted in Figure 1.

➊ **Enable tracing.** The first step is to prepare the kernel so that it records which parts of the kernel code are executed at run time. We use the Linux-provided `ftrace` feature, which is enabled with the `KCONFIG` configuration option `CONFIG_FTRACE`. Enabling this configuration option modifies the Linux build process to include profiling




Figure 1: Workflow of the approach

code that can be evaluated at runtime.

In addition, our approach requires a kernel built with debugging information so that any function addresses in the code segment can be correlated to functions and thus source file locations in the source code. For Linux, this is configured with the KCONFIG configuration option `CONFIG_DEBUG_INFO`.

**2 Run workload.** In this step, the system administrator runs the targeted application after enabling `ftrace`. The `ftrace` feature now records all addresses in the text segment that have been instrumented. For Linux, this covers most code, except for a small amount of critical code such as interrupt handling, context switches and the tracing feature itself.

To avoid overloading the system with often accessed kernel functions, `ftrace`'s own ignore list is dynamically being filled with functions when they are used. This prevents such functions from appearing more than once in the output file of `ftrace`. We use a small wrapper script for `ftrace` to set the correct configuration before starting the trace, as well as to add functions to the ignore list while tracing and to parse the output file, printing only addresses that have not yet been encountered.

**3 Correlation to source lines.** A system service translates the raw address offsets to source line locations using the `ADDR2LINE` tool from the `binutils` tool suite. This identifies the source files and the `#ifdef` blocks that are actually being executed during the tracing phase. Technically, the tool stores its result to a text file with source-file names and line numbers on each line.

**4 Establishment of the propositional formula.** This step translates the source-file locations into a propositional formula. The propositional variables of this formula are the *variation points* the Linux configuration tool KCONFIG controls during the compilation process. This means that every C Preprocessor (CPP) block, KCONFIG item and source file can appear as a propositional variable in the resulting formula. This formula is constructed with the variability constraints that have been extracted from `#ifdef` blocks, KCONFIG feature descriptions and Linux Makefiles. The extractors we use have been developed, described and evaluated in previous work [5, 21, 22]. The resulting formula holds for every KCONFIG configuration that enables all source lines simultaneously.

**5 Derivation of a tailored kernel configuration.** A SAT checker proves the satisfiability of this formula and returns one concrete configuration that fulfills all these constraints. Note that finding an optimal solution to this problem is an NP-hard problem and was not the focus of our work. Instead, we rely on heuristics and configurable search strategies in the SAT checker to obtain a sufficiently small configuration.

As the resulting kernel configuration will contain some additional unwanted code, such as the tracing functionality itself, whitelists and blacklists are employed, allowing the user to specify additional constraints in order to force the selection (or deselection) of certain KCONFIG features. This results in additional constraints being conjugated to the formula just before invoking the SAT checker.

**6 Compiling the kernel.** The resulting solution to the propositional formula, obtained as described above, can only cover KCONFIG features of code that has been traced. As the KCONFIG feature descriptions declare non-trivial dependency constraints [25], special care must be taken to ensure that as many KCONFIG features as possible are not selected while still fulfilling all dependency constraints. We therefore use the KCONFIG tool itself to process this feature selection to a KCONFIG configuration that is both consistent and selects as few features as possible.

### 3 Practical Application

We evaluate the usefulness of our approach by setting up a Linux, Apache, MySQL and PHP (LAMP)-based web presence in a manner that is suited for deployment in a cloud environment. The system serves static webpages, the collaboration platform `DOKUWIKI` [7] and the message board system `PHPBB3` [19] as an example for typical real-world applications. We use the distribution-provided packages from the Debian distribution without further specific configuration changes or optimization. Evaluation results are summarized in Table 1.

#### 3.1 Kernel Tailoring

To derive a minimized kernel configuration, the first step consists of compiling a tracing-enabled Linux ker-

nel. We use the standard Linux kernel source and configuration from the Debian distribution (version 2.6.32-41squeeze2) as a template for our tracing kernel (Step ❶ in Figure 1). On this kernel, we enable the features `CONFIG_FTRACE` and `CONFIG_DEBUG_INFO` to include the `ftrace` tracing infrastructure and compile with debugging symbols. As our current prototype is not able to resolve functions from loadable kernel modules (LKMs) yet, we disable module support in the kernel configuration, which causes all compiled code to be loaded into the system at boot time.

Furthermore, a number of drivers cause compilation and linking errors when not compiled as LKMs. Most of these issues stem from drivers in the `staging`<sup>2</sup> area. Also, when trying to boot this kernel, we observe kernel panics during the initialization of a range of watchdog drivers. As these drivers turn out to be unnecessary for this application scenario, we turn off the `KCONFIG` options `CONFIG_STAGING` and `CONFIG_WATCHDOG`. These configuration changes account for the difference in size and features between the kernel shipped with Debian (~42 MB of code in the `text` segment) and the intermediary kernel that is used for collecting traces (~36 MB of code in the `text` segment).

With this intermediary tracing kernel, the system is tested against a test workload that covers all required functionality. We use the Skipfish [24] security analysis tool to systematically access all functionality of the appliance in an automated manner. This corresponds to Step ❷ in Figure 1 and results in a total of 5,377 observed kernel functions.

These traced kernel functions correlate to 4,686 different source lines in 379 source files (Step ❸). We use a modified version of the `UNDERTAKER` tool [22] to establish the propositional formula (Step ❹) and to derive a solution for it (Step ❺). To avoid unwanted functionality enabled in the resulting kernel, such as the `ftrace` infrastructure itself and LKM support, the `UNDERTAKER` tool obeys a blacklist that consists of the `KCONFIG` options `CONFIG_FTRACE` and `CONFIG_MODULES`. Also, we add eight additional,<sup>3</sup> use-case-agnostic `KCONFIG` items to the whitelist in Step ❻ to enable features that are used by the initialization startup scripts, which run before the system-wide tracing process starts. These steps take 69 sec on a commodity 2.8 GHz quad-core workstation with 4 GB of RAM.

<sup>2</sup>The `staging` area contains unfinished and incomplete drivers that are included as a technology preview.

<sup>3</sup>Specifically: `CONFIG_ACPI`, `CONFIG_UNIX`, `CONFIG_DEVTMPFS`, `CONFIG_DEVTMPFS_MOUNT`, `CONFIG_SERIAL_8250_CONSOLE` and `CONFIG_NOTIFY_USER`, `CONFIG_PM`

Kernel Shipped by Debian	
Loaded Code	5,465,602 Bytes
Total Loadable Code	42,188,538 Bytes
Loaded Kernel Modules	29
<code>KCONFIG</code> options set to <code>y</code>	1,093
<code>KCONFIG</code> options set to <code>m</code>	2,299
Functions with CVE entries	179
Intermediary kernel used for tracing	
Loaded Code	36,341,888 Bytes
Total Loadable Code	36,341,888 Bytes
Loaded Kernel Modules	0
<code>KCONFIG</code> options set to <code>y</code>	3,298
<code>KCONFIG</code> options set to <code>m</code>	0
Functions with CVE entries	207
Resulting application-tailored kernel	
Loaded Code	3,990,153 Bytes
Total Loadable Code	3,990,153 Bytes
Loaded Kernel Modules	0
<code>KCONFIG</code> options set to <code>y</code>	379
<code>KCONFIG</code> options set to <code>m</code>	0
Functions with CVE entries	162

Table 1: Results of the experiment at a glance. The code sizes were obtained with the `SIZE` tool from the `BINUTILS` suite by adding the sizes of the text segments of the bootable kernel image and all loadable `.ko` files.




Figure 2: Comparison of reply rates of the web server with the tailored kernel and the standard distribution kernel.

## 3.2 Evaluation

To ensure the functionality of the appliance, we run the Skipfish [24] security scan again on the system with the tailored kernel, and compare the results with the previous run on the tracing kernel. The comparison of these two reports indicates no differences in the number of vulnerabilities or other issues.

The performance is tested with the `httperf` tool [18]. The tool accesses a static website continuously, at a constant number of requests per second in each run. We did two setups of the same test scenario, both times using the same system, but once booted with the Debian standard kernel, and once with our tailored kernel. The data shows that our tailored kernel achieves a performance very similar to that of the original kernel provided by the distribution.

## 4 Discussion

After the presentation of a practical use case for our approach, this section now evaluates the security benefits. For this, we present an applicable security model to determine the TCB, and discuss security improvements in terms of TCB reduction.

**Security Model.** In the context of the web service presented, we assume both local and remote malicious attackers that target the kernel. However, we do not consider attackers that have physical access to the machine nor attacks that directly target hardware and firmware vulnerabilities.

The security goal is to prevent an attacker from gaining full control with arbitrary code execution in kernel mode, information leakage (e.g., recover uninitialized kernel memory content) to breach confidentiality, and denial-of-service attacks by crashing the kernel to reduce the availability of the system.

**TCB sizes.** Following the literature [13], we define the TCB as “the subset of components that need to be trusted to fulfill the security goals given in the security model”. Therefore, in the security model above, the TCB is solely composed of the kernel, including all LKM loaded during normal operation.

We apply three different metrics to measure the TCB reduction: a) the compiled code (text segment) size of the kernel, b) the total number of features that are enabled in `KCONFIG` and c) the number of functions compiled into the kernel for which there has been an CVE entry in the past 10 years. More precisely, through a semi-automated process, we map a subset of 197 out of 873 CVE entries to vulnerabilities in 215 unique functions in the kernel, and use this dataset. The results for all three kernels used in the experiment in Section 3 are shown in Table 1.

**Results.** The data shows that the Linux kernel shipped by Debian loads 5.5 MB of program code into the memory for the virtual machine in the scenario described in Section 3. Compared with the code size of 4 MB for our tailored kernel, the total TCB size is reduced by 27%.

The number of features enabled is also reduced significantly, from 3,392 (with 1,093 features compiled statically into the kernel and 2,299 as LKM) to 379. The omission of functionality to load further LKMs constitutes an additional security benefit.

Finally, for each function in the TCB, we record the number of known vulnerabilities that have been reported in the past 10 years. When comparing the default distribution kernel to the tailored kernel, we observe a reduction of 10% of functions for which vulnerabilities have been

reported in the past. However, this number is a lower-bound estimate, as the Linux kernel supports on-demand insertion of LKM, resulting in a higher initial TCB size, and therefore higher TCB reduction.

**Sampling bias.** Compared with the code size reduction results above, the CVE reduction numbers may seem lower than expected. We hypothesize that this impression can be attributed to sampling bias: code that is used more often is also audited more often, and better care is taken in documenting the vulnerabilities of such functions. A comparison of the average number of CVEs in kernel functions that are loaded and used (9.8%) with the average number of CVEs in kernel functions that are not used (3.7%) supports this hypothesis. Previous studies [3] have also shown that code in the `drivers/` sub-directory of the kernel, which is known to contain a significant number of rarely-used code, on average contains significantly more bugs than any other parts of the kernel tree. Consequently, it is likely that unused features provided by the kernel still contain a significant amount of relatively easy-to-find vulnerabilities. This further confirms the importance of reducing the TCB size as presented in this work.

**Unexpected impacts.** The presented approach in this work could in turn cause a reduction of the security of the system – a drawback that is common to many security software but is often overlooked. Reviewing the process described in Section 2 (Step 6), we cannot rule out that for some application scenarios, performance-critical or security features might be removed from the base kernel. Possible reasons for this include that a) the feature was not triggered during the system-wide trace, b) the functionality has been excluded from the instrumentation with `ftrace` (e.g., for performance reasons), or c) the configuration options influence the resulting kernel in non-functional ways (e.g., different compilation flags, etc.). Although we were not able to find any results confirming this in this experiment — for example, we have verified that the `CONFIG_CC_STACKPROTECTOR` configuration option, which toggles the inclusion of the GCC flag for adding a stack frame canary, remains enabled, in future work we intend to further evaluate potential adverse impacts.

**Applicability and Incomplete Traces.** The presented approach relies on the assumption that the use-case of the system is clearly defined. Thanks to this a priori knowledge, it is possible to determine what kernel functionalities the application requires and therefore, what kernel configuration options have to be enabled. Still, the lack of guarantees that a trace of a given scenario is sufficient

and captures all required functionality may raise concerns. However, this does not invalidate the approach. Our approach works best for service providers that instantiate fairly homogeneous services, for which the the chance to capture all necessary functionality is highest. Unfortunately, even the best tracing mechanism could still not guarantee catching all possible cases. Such a guarantee could only be provided with formal analysis and verification – a method that stands mathematical proof. However, for many scenarios organizations and system operators are likely to accept less costly approaches – such as the one presented in this paper. With the increasing importance of compute clouds, where customers employ virtual machines for very dedicated services such as the web server presented in Section 3, we expect that our approach can be easily applied to further use cases that are commonly deployed in the cloud.

Fortunately, incomplete traces do not lead to system crashes. With `KCONFIG`, changing the static configuration of the Linux kernel statically restricts the provided functionality to applications that interact with the kernel via well-defined interfaces. These interfaces, and the careful maintenance of kernel developers, allows applications to check for the presence of kernel features that are required for correct operation. In the absence of kernel bugs, applications must not be able to crash the kernel. On the other hand, applications that crash because of missing kernel features fail to provide proper error handling – which alone makes them unsuitable for security sensitive environments. In summary, crashes that occur because of incomplete traces always stem from easy-to-address software bugs for the scenario presented in this paper.

**Usability.** Most of the steps presented in Section 2 require no domain specific knowledge of Linux internals. We therefore expect that they can be conducted in a straightforward manner by system administrators without specific experience in Linux kernel development. The system administrator, however, continues to use a code base that continuously receives maintenance in form of bug fixes and security updates from the Linux distributor. We therefore are confident that our approach to automatically tailor a kernel configuration for specific use-cases is both practical and feasible to implement in real-world scenarios.

**Extensibility.** The experiment in Section 3 shows that the resulting kernel requires eight additional `KCONFIG` options for proper operation. Alternatively to adding these features to the whitelist with distribution-specific knowledge, starting the application tracer at the start of the boot process would also capture the missing functionality. However, in this way we demonstrate the ability to specify wanted or unwanted `KCONFIG` options independently of

the tracing. This allows our approach to be assisted in the future by methods to determine kernel features that tracers such as `ftrace` cannot observe at all.

## 5 Related work

As we show below, this work relates to two research areas.

**Kernel specialization.** Several researchers have suggested approaches to tailor the Linux kernel, although security is usually not a goal, but improvements in code size or execution speed are: Lee et al. [14] manually modify the source code (e.g., by removing unnecessary system calls) based on a static analysis of the applications and the kernel. Chanet et al. [2], in contrast, propose a method based on link-time binary rewriting, but also employ static analysis techniques to infer and specialize the set of system calls to be used. Both approaches, however, do not leverage any of the built-in configurability of Linux to reduce unneeded code. Moreover, our approach is completely automated.

TCB reduction has always been a major design goal for micro-kernels [1, 15], which in turn facilitates a formal verification of the kernel [10] or its implementation in type-lafe languages, such as OCaml [16].

**Kernel attack surface reduction.** The security model used in this paper is commonly used when building *sandboxing* or *isolation* solutions, in which each process must be contained within a particular security domain, such as [4, 9, 17], which are all based on the Linux Security Module (LSM) framework [23]. The idea of directly restricting the system call interface on a per-process basis has been previously explored as well, e.g., by Fraser, Badger, and Feldman [6], Ko et al. [11], or Provos [20], although not with specific focus on reducing the kernel’s attack surface. `Seccomp` [8] directly tackles this issue by allowing processes to be sandboxed at the system call interface. `Ktrim` [12] is a current research project which goes beyond simply limiting the system call interface, and explores the possibility of finer-granularity kernel attack surface reduction by restricting individual functions (or sets of functions) inside the kernel. In contrast, this work focuses on compile-time removal of functionality from the kernel at a system-wide level instead of a runtime removal at a per-application level.

## 6 Conclusion and Future Work

This paper presents an approach for automatically tailoring a Linux kernel configuration to a given use case. The result is a Linux kernel in which unnecessary functionality is removed at compile-time, hence significantly reducing

TCB size. The reduction can be quantified with 27% less code loaded and at least 10% fewer kernel functions which were previously vulnerable to attacks.

While the current prototype shows promising results, we intend to improve on the usability and applicability to additional use-cases. For instance, the current prototype unconditionally disables module loading support. As this may be undesirable in some cases, we intend to improve the handling of LKMs, as well as to remove the need for an intermediary tracing kernel.

## Acknowledgments

This research has been partially supported by the TClouds project<sup>4</sup> funded by the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreement number ICT-257243.

<sup>4</sup><http://www.tclouds-project.eu>

## References

- [1] Mike Accetta, Robert Baron, David Golub, Richard Rashid, Avadis Tevanian, and Michael Young. “MACH: A New Kernel Foundation for UNIX Development”. In: *USENIX Summer Conference*. USENIX, 1986.
- [2] Dominique Chagnet, Bjorn De Sutter, Bruno De Bus, Ludo Van Put, and Koen De Bosschere. “System-wide Compaction and Specialization of the Linux Kernel”. In: *2005 ACM SIGPLAN/SIGBED Conf. on Languages, Compilers and Tools for Embedded Systems (LCTES ’05)*. ACM, 2005. DOI: 10.1145/1065910.1065925.
- [3] Andy Chou, Junfeng Yang, Benjamin Chelf, Seth Hallem, and Dawson Engler. “An empirical study of operating systems errors”. In: *18th ACM Symp. on OS Principles (SOSP ’01)*. ACM, 2001. DOI: 10.1145/502034.502042.
- [4] Kees Cook. *Yama LSM*. 2010. URL: <http://lwn.net/Articles/393012/> (visited on 06/04/2012).
- [5] Christian Dietrich, Reinhard Tartler, Wolfgang Schröder-Preikschat, and Daniel Lohmann. “A Robust Approach for Variability Extraction from the Linux Build System”. In: *16th Software Product Line Conf. (SPLC ’12)*. (To appear). ACM, 2012.
- [6] Timothy Fraser, Lee Badger, and Mark Feldman. “Hardening COTS software with generic software wrappers”. In: *20th Symp. on Security and Privacy*. IEEE, 1999. DOI: 10.1109/SECPR1.1999.766713.
- [7] Andreas Gohr. *DokuWiki*. URL: <http://dokuwiki.org> (visited on 06/03/2012).
- [8] *Google Seccomp Sandbox for Linux*. URL: <http://code.google.com/p/seccompsandbox/wiki/overview> (visited on 06/05/2012).
- [9] Toshiharu Harada, Takashi Horie, and Kazuo Tanaka. “Task Oriented Management Obviates Your Onus on Linux”. In: *Japan Linux Conference* (2004).
- [10] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. “seL4: formal verification of an OS kernel”. In: *22nd ACM Symp. on OS Principles (SOSP ’09)*. ACM, 2009. DOI: 10.1145/1629575.1629596.
- [11] Calvin Ko, Timothy Fraser, Lee Badger, and Douglas Kilpatrick. “Detecting and countering system intrusions using software wrappers”. In: *9th Conf. on USENIX Security Symposium (SSYM ’00)*. USENIX, 2000.
- [12] Anil Kurmus, Alessandro Sorniotti, and Rüdiger Kapitza. “Attack surface reduction for commodity OS kernels: trimmed garden plants may attract less bugs”. In: *4th Eur. Workshop on system security (EUROSEC ’11)*. ACM, 2011. DOI: 10.1145/1972551.1972557.
- [13] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. “Authentication in distributed systems: theory and practice”. In: *ACM Trans. Comp. Syst.* 10.4 (1992). DOI: 10.1145/138873.138874.
- [14] C.T. Lee, J.M. Lin, Z.W. Hong, and W.T. Lee. “An Application-Oriented Linux Kernel Customization for Embedded Systems”. In: *Journal of information science and engineering* 20.6 (2004).
- [15] Jochen Liedtke. “On  $\mu$ -Kernel Construction”. In: *15th ACM Symp. on OS Principles (SOSP ’95)*. ACM OSR. ACM, 1995. DOI: 10.1145/224057.224075.
- [16] A. Madhavapeddy, R. Mortier, R. Sohan, T. Gazagnaire, S. Hand, T. Deegan, D. McAuley, and J. Crowcroft. “Turning Down the LAMP: Software Specialisation for the Cloud”. In: *2nd USENIX Conf. on hot topics in cloud computing (HOTCLOUD ’10)*. USENIX, 2010.
- [17] Frank Mayer, Karl MacMillan, and David Caplan. *SELinux By Example: Using Security Enhanced Linux*. Prentice Hall, 2006.
- [18] David Mosberger and Tai Jin. “htperf. A tool for measuring web server performance”. In: *SIGMETRICS Performance Evaluation Review* 26.3 (1998). DOI: 10.1145/306225.306235.
- [19] *phpBB. Free and Open Source Forum Software*. URL: [www.phpbb.com](http://www.phpbb.com) (visited on 06/03/2012).
- [20] Niels Provos. “Improving host security with system call policies”. In: *12th Conf. on USENIX Security Symposium (SSYM ’03)*. Vol. 12. USENIX, 2003.
- [21] Julio Sincero, Reinhard Tartler, Daniel Lohmann, and Wolfgang Schröder-Preikschat. “Efficient Extraction and Analysis of Preprocessor-Based Variability”. In: *9th Int. Conf. on Generative Programming and Component Engineering (GPCE ’10)*. ACM, 2010. DOI: 10.1145/1868294.1868300.
- [22] Reinhard Tartler, Daniel Lohmann, Julio Sincero, and Wolfgang Schröder-Preikschat. “Feature Consistency in Compile-Time-Configurable System Software: Facing the Linux 10,000 Feature Problem”. In: *ACM SIGOPS/EuroSys Eur. Conf. on Computer Systems 2011 (EuroSys ’11)*. ACM, 2011. DOI: 10.1145/1966445.1966451.
- [23] Chris Wright, Crispin Cowan, James Morris, Stephen Smalley, and Greg Kroah-Hartman. “Linux Security Module Framework”. In: *Ottawa Linux Symposium*. 2002.
- [24] Michal Zalewski, Niels Heinen, and Sebastian Roschke. *skipfish. Web application security scanner*. URL: <http://code.google.com/p/skipfish/> (visited on 06/03/2012).
- [25] Christoph Zengler and Wolfgang Kuchlin. “Encoding the Linux Kernel Configuration in Propositional Logic”. In: *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010) Workshop on Configuration 2010*. 2010.