Verifying the Reliability of Operating System-Level Information Flow Control Systems in Linux

Laurent Georget* Mathieu Jaume† Guillaume Piolle* Frédéric Tronel* Valérie Viet Triem Tong*

*EPC CIDRE SUPELEC/INRIA/CNRS/University of Rennes 1 †Sorbonne Universités, UPMC, CNRS, LIP6 UMR 7606
Rennes, France Paris, France

Email: laurent.georget@irisa.fr

Abstract—Information Flow Control at Operating System (OS) level features interesting properties and have been an active topic of research for years. However, no implementation can work reliably if there does not exist a way to correctly and precisely track all information flows occurring in the system. The existing implementations for Linux are based on the Linux Security Modules (LSM) framework which implements *hooks* at specific points in code where any security mechanism may interpose a security decision in the execution. However, previous works on the verification of LSM only addressed access control and no work has raised the question of the reliability of information flow control systems built on LSM. In this work, we present a compiler-assisted and reproducible static analysis on the Linux kernel to verify that the LSM hooks are correctly placed with respect to operations generating information flows so that LSM-based information flow monitors can properly track all information flows. Our results highlight flaws in LSM that we propose to solve, thus improving the suitability of this framework for the implementation of information flow monitors.

Keywords—Information flow control; LSM; Static analysis

I. INTRODUCTION

Information flow tracking is a security mechanism designed to monitor how sensitive information spreads in a system. Using this knowledge, information flow control can be applied to put limits on the dissemination of a piece of sensitive data once it is out of its *container of information* and let high level policies such as “my banking information shall not be sent to the Internet” or “information from accounting files from 2014 and 2015 shall not be mixed” be enforced easily. Another very active application domain of the information flow control (IFC) is the analysis of malwares [1], [2], [3] because, as obfuscated as its binary code may be, the information flows performed by a program always reveal its behavior.

We define an explicit information flow as the copy, usually partial, of the content of one container of information to another. The scope of this paper is the IFC at the Operating System (OS) level. For an OS, a container of information can be an active entity executing code (a process), a storage place (a file) or a channel of communication between active entities (a network socket, a shared memory segment, a pipe). For example, a process writing the content of a buffer in memory to a file generates an information flow from the process to the file. Information flow tracking consists in attaching a piece of meta-information on each container, called a *label*, that remembers

where the information it contains comes from. Each time there is an information flow in the system, this should make the information flow tracker update the destination label according to the source label. Although mandatory access control models can also enforce information flow policies, we focus here on IFC performed with information flow *tracking*.

The flows are not usually performed directly by the processes but by the OS privileged software, the kernel, for security reasons. To write a file, or to start a new process, for example, a process has to ask the kernel to perform the desired task on its behalf by a *system call*, an entry point function implemented in the kernel. The kernel can deny any such request, if it lacks necessary resources or if the process lacks the required clearance. Therefore, the system call barrier is an ideal candidate to draw a clear boundary between containers of information and to define the granularity of the IFC. For instance, as all threads of execution in a single process share a single memory space, and can freely exchange information and work on the same data without requiring intervention from the kernel, we ought to consider the entire process, including all its threads, as a single container of information. We restrict ourselves to *explicit* information flows such as information storing or inter-process communications through means designed for this purpose. We thus exclude covert channels and information leakage due to some event or operation *not* occurring in the system.

Several approaches have been proposed to implement information flow tracking in Linux-based OSes. For instance, Blare [4], Laminar [5], and Weir [6] implement information flow tracking inside the Linux kernel. They do so thanks to the Linux Security Modules framework (LSM) [7], [8]. This framework was designed to allow any security mechanism be implemented efficiently and easily in the Linux kernel. Before security-sensitive operations (such as writing to a file, creating a new process, etc.), the kernel hands the execution out to the installed security module to let it take a security decision. The points in the code where security modules may interpose itself in the execution are called *hooks*. Although the formal validation of the position of the hooks has already been explored in details [9], [10], [11], [12], these works have not considered the information flow tracking use case. The main difference is that access control only needs to deal with granting or denying access to an object to a subject based on

a predetermined and mostly static policy and classification of objects and subjects according to their sensitivity and need-to-know. Information flow control takes into account the past interactions of a subject or object before taking a security decision. For example, it is possible for a top-secret subject to write into a public file only if said subject has not read secret information beforehand. The analysis of our results in Section V gives more insight on why a correct placement of hooks for the access control may not be correct for information flow monitoring purpose. To the best of our knowledge, our work is the first to formally assess the correctness of LSM hooks positioning for information flow tracking purposes.

Our work aims at verifying a necessary condition on the placement of hooks: that any execution path that generates an information flow goes through a LSM hook before the flow is performed. This is a necessary condition for the implementation of correct information flow trackers. Our main contributions are: (1) a formal model of the Linux code, produced by the GCC compiler itself, (2) a reproducible static analysis that for each system call, and each execution path reaching a point where an information flow is done, verifies that the path is either impossible or covered by a LSM hook (3) a study of the flaws in the design of LSM exhibited by our results and the corresponding patch to improve LSM's suitability for the implementation of information flow trackers. As a side note, we also make a case for compiler-assisted static analyses. Indeed, we are not able to define entirely the source language (a dialect of C specific to the compiler which is used, including parts written in assembly) but we can make working hypotheses on it and rely on the compiler to provide us with the knowledge that is hard to statically correctly decide such as the aliasing information (the C language supports arbitrary pointer arithmetic). Due to space restrictions, we do not expose all the formalism but details and proofs are available on the project website at <https://kayrebt.gforge.inria.fr>.

After giving in section II some background about our research, we present our analysis in the remaining sections. In Section III, we describe the formal model we use to represent the system calls. We want to formally verify that all execution paths which generate information flows but which do not contain a LSM hook are impossible paths. To do so, we equip the formal model with a semantics and design a static analysis we prove sound in Section IV. In section V, we present the GCC plugin we implemented to run our analysis and the results of our experimentations. Along with these results, we propose a revision of LSM with some hooks added, and some existing hooks relocated, to better adapt it to the control flow purpose. Finally, we discuss related work in section VI and conclude in Section VII with the further work we envision.

II. BACKGROUND AND OBJECTIVES

A. Linux Security Modules

Since version 2.6, Linux supports a framework to implement security extensions for the kernel called Linux Security Modules (LSM) [7]. This framework provides a set of *hooks* strategically placed in the kernel code associated with fields

in internal data structures for exclusive use by security extensions. The hooks are functions which can be used by security extensions to (1), allocate, free, and maintain the security state of various internal data structures having a dedicated security field, and (2), implement security checks at specific points of execution, based on the security state and a policy. Security modules have a chance to apply security restrictions anywhere a hook is present, but only at these places. LSM's original design is the access control and this has dictated the placement of hooks in the code. It is thus necessary to verify the correctness of this placement for the purpose of information flow tracking to ensure that information flow trackers such as Blare, Laminar and Weir can operate properly.

B. Complete Mediation Property

Our main goal is to verify the property we call "Complete Mediation". For any execution path in the kernel starting with a system call and leading to an information flow, there is at least one LSM hook in this path which is reached before the flow is performed. The rationale here is that if a path generating a flow but not going through any LSM hook exists, then a malicious program can exploit it to perform illegal information flows without triggering any alarm because the information flow monitor can only react when one of the hooks is reached. Identifying all the paths that lead to information flows requires solving two common problems in static analysis. First of all, the number of execution paths is infinite because of loops in the code so we must identify a finite subset on which running the analysis is sufficient to draw a conclusion for all paths. Secondly, many execution paths that appear in the control flow graph of a program cannot actually be taken. These paths are often called *infeasible paths* in the literature [13]. It is important to identify them, otherwise we might declare that the information flow monitor does not correctly detect some possible flows while in fact, there is no way to perform them.

For our analysis, we consider the list of system calls presented in Table I, which covers all the system calls monitored by either Blare, Laminar or Weir. Flow control requires knowing precisely when an information flow starts and when it stops. Otherwise, it is unable to maintain a correct and precise representation of all flows currently taking place in the system at any given time. This can result in either (1) wrong security decisions, compromising the confidentiality or integrity of sensitive data, or (2) overly conservative ones, making the system unusable. Access control on the other hand typically requires monitoring only the beginning of the flow (it would be inconsistent to authorize a process to gain access to a container but to forbid it to lose that access). Since LSM was designed with access control in mind, some hooks are missing to perform information flow tracking. Section V explore these points more in detail.

III. MODELING THE SYSTEM CALLS CAUSING FLOWS

A. Control Flow Graphs

The analysis we propose relies on the C compiler from the Gnu Compilers Collection [14], used to compile the Linux

kernel. The model we designed to represent system calls and their execution paths does not describe directly the C source code but instead an internal representation created by GCC in a language called GIMPLE [15]. Each system call is represented by a control flow graph (CFG), where paths, as defined by the classical graph theory, model execution paths in the program [16]. Our analysis examines one system call at a time. To produce a single graph for an entire system call and all the functions it calls, we force the inlining of the latter into the former. This does not change the semantics of the code but reduces our analysis to the intra-procedural case. In these CFGs, we mark two kinds of nodes: the nodes corresponding to LSM hooks and the nodes corresponding to operations generating flows. The former can be identified automatically and we identified manually the latter in the CFGs (cf. Section III-B).

In our CFGs, contrarily to most approaches, a node is not a basic block but a single GIMPLE instruction. Edges can bear guards, to indicate conditional jumps between instructions. Our analysis does not need to deal with all expressions and variables of the language. For example, we do not have to handle floating-point values because their use is explicitly prohibited in the Linux code. We chose not to handle variables representing structures or unions when they involve pointer arithmetic. Powerful static analysis are allowed by generic framework such as Blast [17] but dealing with complex types is not necessary for our specific needs. We do not handle global or volatile variables either since they can have an arbitrary value at any point in the execution, regardless of what is done in the function. Ignoring some variables does not hinder the soundness of our approach: less impossible paths might be detected as such but we never declare as impossible a possible path. Our results show that the analysis is precise enough to conclude in all cases.

To handle the pointer variables, we need to partition variables along two different criteria: (1) the set of pointers $\mathcal{Vars}^{ptr}$, i.e. variables whose value is the address of another variable, versus the set of integer variables $\mathcal{Vars}^{\mathbb{Z}}$; and (2) the set of variables whose address is taken at one point in the program $\mathcal{Vars}^{mem}$ versus the variables that are not addressable (not *aliasable*) and do not participate in side effects $\mathcal{Vars}^{temp}$ (i.e. their value can only be set by direct assignment). The segregation between the variables is directly extracted from the compiler, which maintains precise typing information about all expressions in the code.

In addition to variables, several kinds of expressions have to be handled: the integer constants $\mathbb{Z}$, the pointers dereferences $\{*p \mid p \in \mathcal{Vars}^{ptr}\}$, the address-taking operations $\{\&v \mid v \in \mathcal{Vars}^{mem}\}$, and finally the expressions of unknown value. As a matter of fact, all expressions leading to computations fall into the latter category, as well as unanalyzed variables. Finally, our analysis relies on the *points-to* oracle maintained by GCC, which gives for any pointer the set of variables it can possibly points to.

A path in the CFG is said to be impossible when any execution that would follow it would enter in an impossible state.

For example, a path including two conditional branchings with incompatible conditions would require a boolean expression to be both true and false at the same time.

B. Rationale and Discussion of our Approach

Using the compiler itself to extract the knowledge we need from the kernel has several advantages. First of all, we have the guarantee that the analysis is bound to the exact code that is compiled in. Due to conditional compilation, the code of a given function can vary greatly depending on the compilation options. Moreover, it is important to note that GCC’s extensions to the C language are explicitly permitted. Relying on GCC is therefore not only useful but practically mandatory to correctly capture the true semantics of the code. Finally, even if the representation we extract during the compilation process and use as a model of the code does not reflect what the developers *intended to mean*, it is anyway what the code *does mean* when executed.

Another decision we made was declaring where information flows are actually performed inside the system calls. This is not trivial because the low-level data structures manipulated in the kernel code sometimes only relates loosely to the abstractions we call “files” or “sockets”. For instance, writing to a file on disk actually means writing to a cache in memory, which is periodically flushed to the disk by a dedicated kernel process. It is natural to consider that when the writing process returns from `write` system call, the information flow is done because any process reading the file would find the information just written there but in case of a power outage for example, the information could be lost if it had not been flushed to disk yet. We used several heuristics to decide objectively that the flow was performed at some point in the code: the fact that all checks were passed, the fact that all locks on structures were taken, and finally the fact that the operation was “committed” by the kernel by a log message for example.

IV. STATIC ANALYSIS ON PATHS

Our objective is, for any CFG representing a system call generating an information flow, to enumerate all the paths of this CFG reaching the point where the information flow is done and for each one of these paths, to verify that it goes through a LSM hook or that it is impossible. If there exists a path which is possible but not covered by a LSM hook, then this is a flaw in the placement of hooks. In this section, we describe the main lines of the static analysis and the main results we have proved onto it.

A. Verifying the Complete Mediation

We consider the set $\mathcal{Paths}$ of all paths in a CFG. Among this paths, we distinguish two particular subsets: (1) the set $\mathcal{Paths}_{flows}$ of paths starting at the initial node of the CFG and ending at one node generating an information flow; and (2) the set $\mathcal{Paths}_{LSM}$ of paths having a node corresponding to a LSM hook. The placement of the hooks would be obviously correct if we could prove that $\mathcal{Paths}_{flows} \subseteq \mathcal{Paths}_{LSM}$. However, as

we will see, this is not the case. $\mathbf{P}_f = \text{Paths}_{flows} \setminus \text{Paths}_{LSM}$ represents the set of paths that may be problematic.

As explained earlier, some paths in Paths_{flows} are actually impossible, and therefore even if there are no LSM hooks in them, they are not actually problematic. Recall that an impossible path is a path that does not correspond to a possible execution. The objective of our analysis is thus to verify that $\mathbf{P}_f \subseteq \mathbf{I}$ where $\mathbf{I} \subseteq \text{Paths}$ is the set of impossible paths in the CFG. However, due to the presence of loops, the set $\mathbf{P}_f$ is infinite and we cannot run our analysis on all paths in $\mathbf{P}_f$. To tackle this problem, we define an equivalence relation on paths in $\mathbf{P}_f$: two paths are equivalent if they are identical up to the cycles they contain. Our analysis run only on acyclic paths (of which there is a finite number) and we prove that our analysis is conservative, i.e. that it will not report a path as impossible if at least one equivalent path is possible.

Of course, in any CFG there is a finite number of acyclic paths, so it is possible (and sufficient) to verify them all. We state the complete mediation property:

Property 1 (Complete Mediation):

Complete mediation holds iff: $\mathbf{P}_f \subseteq \mathbf{I}$, i.e. all the execution paths that perform an information flow and are not controlled by the information flow monitor since they do not contain a LSM hook are impossible according to the static analysis.

Figure 1. Sets involved in the analysis

B. Configurations and Satisfiability

The analysis of a path requires maintaining an abstraction of the possible state of the execution at each node. This state, which we call an *abstract configuration*, is twofold: the path configuration contains information about the path followed so far by the execution in the graph and the variables configuration contains constraints about the possible values of variables. Our analysis is said to be *path-sensitive* because it leverages the knowledge of the execution path taken. Actually, due to the properties of the graphs and the needs of our analysis, it is only necessary to remember the edge taken to reach the last node with strictly more than one predecessor node (i.e. the node at the beginning of the current basic block).

A variables configuration written (C, P) is made up of a *set of constraints* C on the integer variables and a mapping P for the pointers. C is a set of tuples of the form (x, R, y) where x is a variable, $R \in \{=, \neq, <, >, \leq, \geq\}$ a relation operator, and y either a variable or an integer constant. Such a constraint

(x, R, y) expresses the fact that the predicate $\bar{x} R \bar{y}$, where $\bar{v}$ stands for the value of v , holds. The pointer mapping P records the value of pointers, if known. We have $P : \text{Vars}^{ptr} \rightarrow \text{Vars}^{mem} \cup \{\top\}$. With $p \in \text{Vars}^{ptr}$, if $P(p) = v \in \text{Vars}^{mem}$ then we know that p points to v , otherwise, we have no more information than the points-to oracle. Indeed, since our analysis is path-sensitive, it has sometimes more information about the value of a pointer than the path-insensitive points-to analysis made by the compiler. The configurations are updated by following the execution path: assignments add constraints on variables or, on the contrary, remove some. For example, when a variable is assigned the return value of a function, the analysis remove all constraints about the variable since the return value of a function is statically undecidable in the general case. An edge always adds as a constraint the condition it bears, if any.

Our problem of identifying impossible paths is thus reduced to a problem of satisfiability: if there does not exist a valuation for the variables that satisfies all the constraints simultaneously, then we consider that we have reached an impossible state of execution. This corresponds to incompatible constraints and means that the followed path does not correspond to a possible execution. If the set of constraints is satisfiable, or at least if the analysis cannot prove it is unsatisfiable, then we declare the path possible.

C. Concrete Configurations

In order to prove the soundness of the static analysis, we have to describe the concrete execution. We define the *concrete configuration* as a valuation of the variables of the program, at a given point in its execution. First of all, we consider the memory as a set of memory cells, each one holding a single value, which may be either an integer or the identifier of another memory cell (an “address”). A concrete configuration has three components: (1) a mapping from memory cells to values, (2) a mapping from non-addressable variables to values, and (3) a mapping from addressable variables to memory cells. The first mapping is the memory state of the program. The second one is used to deal with the particular case of non-addressable variables which might not live in memory. Although they appear in the CFG built by the compiler, they have a single value during their lifespan and are not susceptible to be modified via an indirect assignment such as a pointer for example, so it is not necessary to account for them in the memory state. In fact, although they appear in the CFG built by the compiler, they might very well be optimized out in the resulting executable program.

D. Concrete and Abstract Semantics

We aim at proving the soundness of our static analysis: if a path is declared as impossible by the static analysis, then it is indeed impossible, although the analysis might not detect all impossible paths. To prove this property, we suppose the existence of a concrete semantics of the CFG, the one computed by the compiler. Since the C language, and the compiler’s internal representations, including the CFG, have

only informal semantics, we only make minimal, common hypotheses on the concrete semantics, necessary to our analysis, without defining it entirely. This way, we eliminate the risk of diverging from the compiler’s actual semantics. The concrete semantics is a transition relation $\rightarrow \subseteq \Theta \times \mathcal{V} \times \mathcal{C} \times \Theta$ where Θ is the set of concrete configurations, $\mathcal{V}$ the set of nodes and $\mathcal{C} = (\text{Vars} \times \{=, \neq, <, >, \leq, \geq\} \times (\text{Vars} \cup \mathbb{Z})) \cup \{\text{true}\}$ is the set of constraints on edges. true is a special constraint which is always satisfied (to model unconditional edges).

We then proceed to completely define an abstract semantics. The abstract configuration is a transition relation $\rightarrow \subseteq \mathcal{K} \times \mathcal{V} \times \mathcal{C} \times \mathcal{K}$ where $\mathcal{K}$ is the set of abstract configurations. The entire definition can be found in the extended appendix of the paper, available on the project’s website. We extend naturally these relations over paths.

To express the soundness of our analysis, we define a satisfiability relation between concrete and abstract configurations. We say that a concrete configuration satisfies an abstract one if the valuation of the variables given by the concrete configuration satisfies all the constraints in the abstract configuration. Noting θ a concrete configuration, and k an abstract one, we write $\theta \models k$ when θ satisfies k and $\models k$ the fact that k is satisfiable, i.e. there exists some θ such that $\theta \models k$.

The soundness of our analysis is stated as follows.

Proposition 1 (Soundness):

For each path p in a CFG, for all concrete configurations θ_1 and θ_2 , and all abstract configurations k_1 and k_2 such that $\theta_1 \rightarrow_p^* \theta_2$ and $k_1 \rightarrow_p^* k_2$, we have $\theta_1 \models k_1 \implies \theta_2 \models k_2$.

We define the set of executable paths as the set of all paths p such that there exists two concrete configuration θ_1 and θ_2 such that $\theta_1 \rightarrow_p \theta_2$. Remember that a transition may not exist from one configuration to another if, along the path p , there is an edge bearing a condition which is not verified by the memory state. The set of impossible paths is defined as the set of paths p for which given any two abstract configurations k_1 and k_2 , if $k_1 \rightarrow_p^* k_2$, then $\not\models k_2$. In other words, there is no satisfiable configuration that can result from the analysis of path p . The following proposition is an immediate consequence of the previous one.

Proposition 2 (Executable vs. impossible paths): The sets of executable paths and impossible paths are disjoint.

E. Handling Loops

The static analysis, as presented so far, handles all paths of finite length. However, in a CFG, there are an infinity of such paths, because of the presence of loops. Fortunately, loops have a special syntax in CFGs. Specifically, loops must start with one join node with exactly two predecessors: one before the loop, and one inside. Loops are also disambiguated by the compiler: a unique node cannot be the beginning of two different loops. Furthermore, for any node, we can use the compiler as an oracle to tell what is the most-outer loop it is part of, if any, and whether it is the beginning of a loop. We make the hypothesis that the number of iterations of each loop is unbounded but finite, since it would actually be a bug for a system call not to terminate. We deal with loops in our analysis

by computing a loop abstractor, i.e. a configuration such that if it is unsatisfiable, then all abstract configurations that could possibly result from the analysis of this loop are unsatisfiable. In other words, we abstract in a single equivalence class all the paths composed of a finite, albeit arbitrarily large, number of iterations of the loop. Our analysis therefore trades off precision for termination by reducing the analysis of any loop to the computation of its abstractor.

V. IMPLEMENTATION AND RESULTS

The static analysis is run during the compilation itself. To this effect, we developed a plugin for GCC version 4.8. This plugin is inserted into the compilation process, when the code is in GIMPLE form. We chose to place the plugin as late as possible in the compilation process, just before GCC abandons the graph intermediate representation, in order to benefit the most from optimizations and to work on a code as simple as possible. This allows us to do the assumptions we presented in the previous sections, such as the particular properties of loops. The representation we dump is broken down to very elementary pieces. The code we handle is in three-addresses mode, which has the effect that all complex boolean conditions that could exist in the original code base are already broken down in the appropriate number of binary decision nodes and branchings. We also leverage as much as possible the compiler. For example, we rely on it to identify the loops in the CFG, to know the precise typing information of variables, and to run the points-to analysis on pointers. The latter is already implemented because it is used for several optimizations passes by GCC. Finally, the careful use of inlining and the fact that GCC applies some optimizations on the code actually limits the path explosion. Our tools, including the plugin for the static analysis and also a plugin able to dump CFGs, are available at <http://kayrebt.gforge.inria.fr>.

Prior to running the analysis, we place a special annotation on places in the code of system calls where information flows are performed. When the analysis is run on a system call, for each annotated places, a subgraph of the CFG built by GCC is considered: the subgraph of paths starting from the node representing the entry of the system call and going to the information flow node that do not pass through any LSM hook. These paths are the set $\mathbf{P}_f$. For each path, we start with an empty abstract configuration and we update the configuration as we go along the path as per the transition rules of the abstract semantics. The satisfiability of the abstract configuration is tested by Yices [18], a SAT-solver equipped with a decidable subset of the classic theory of integers. If the constraint solver declares the set of constraints unsatisfiable then we declare the path impossible.

We did our analysis on the Linux kernel version 4.3, released on November 2nd 2015, compiled with the default configuration options for architecture x86_64. However, the same analysis can be reproduced with the same tools on any version, any platform. Most of the code we analyze is part of high-level submodules, which are not likely to change

greatly from one version to the next, and are not architecture-dependent. Nonetheless, the list of available system calls depends on the configuration options, therefore, the options available for performing information flows can vary from a Linux system to another. The results of our analysis are presented in Table I, where system calls are grouped by the LSM hooks they share. In a majority of cases, either no paths evading the hooks are found or they are all impossible. In these cases, LSM hooks are placed correctly to detect the information flows generated by these system calls. In some other cases, though, some system calls we have identified as generating flows have no LSM hooks at all.

a) *System calls vmsplice, splice, and tee*: These system calls take advantage of the implementation of the pipes as memory buffers to perform efficient copies to or from pipes. No hook is triggered when the flow occurs between two pipes.

b) *mq_timedreceive, mq_timedsend*: These system calls are used to, respectively, receive and send a message through a POSIX message queue. They behave the same as *msgrcv* and *msgsnd* for System V message queues, which do have hooks. We got some elements of explanations on the LSM kernel development mailing-list from SELinux developer Stephen Smalley [19]. First, these system calls are more recent than LSM so they were not included in the original LSM design. Secondly, the difference between the two message queues APIs is that the POSIX one is based on the virtual filesystem. For the purpose of access control, it already benefits from the security hooks in the *open* system call. However, and this also applies to *splice*, in the system calls that do the actual information flow, there are no hooks because pipes and message queues are not thought as being subject to the same problem of policy change and access control revalidation as regular files, and therefore it feels less necessary to place a hook before each individual “read” and “write” operation as for files. These considerations are valid only when considering access control. Flow tracking requires monitoring each individual information flow.

c) *process_vm_readv, process_vm_writev*: In these system calls, two paths are possible. However, they are designed as such and correspond to flows from a process to itself (possibly from one thread to another). In our model, a process as a whole is a container of information because flows between threads are possible without system calls anyway.

d) *recvmsg, sendmsg*: The case of these system calls illustrates why a framework built for access control might be insufficient for flow control. The system call *recvmsg* allows a process to receive multiple messages from a socket with a single system call. A process may use it to communicate with another one. A LSM hook is present in this system call, so that the communication is mediated. However, the hook is positioned such that the security function is called only once, before the first message is received. A malicious process could exploit this fact to send sensitive information acquired (via another thread) between the first and the subsequent messages. This is a problem relevant only to flow control and not to access control. With access control, the dynamic content of

a container of information has no influence over the security decision, therefore it is sufficient to do the check only once per system call.

VI. RELATED WORK

To the best of our knowledge, no previous work has aimed at assessing the adequacy of LSM for information flow tracking. However, several works have focused on the correctness of the placement of the LSM hooks for its original purpose: access control. Zhang et al. performed a static analysis with CQUAL [20] to check if all accesses to internal data structures are correctly intercepted by the appropriate LSM hooks. This is different from our approach because we consider more abstract structures, containers of information, whose link with the real data structures of the kernel is not always direct. This work was extended [9] with a consistency analysis on the relationships between the security-sensitive operations performed by the kernel and LSM hooks. In this work, Zhang and Jaeger identified all the paths that can lead to an access to a member of an internal data structure at a granularity finer from that of system calls. For example, with their analysis, they have been able to detect that the *f_owner.pid* member of *struct file* can be modified via an operation which, as a side-effect, resets this field to 0 without calling the hook designed to protect this field before. Again, this approach is too low-level for us because we deal with more abstract objects. Nevertheless, our future work will adapt from this approach the notion of *consistency*. In this current work, we have only checked that any flow is monitored by *some* hook. Later works by Ganapathy, Jaeger, and Jha [10] and Muthukumaran, Jaeger, and Ganapathy [12] focused on automatic placement of hooks. This approach is promising but not practical in our case since we do not use the LSM framework for its primary purpose but we want to modify it as little as possible nevertheless. Bringing IFC in Linux through LSM is an active topic and adding hooks or moving some of them has been necessary for *Blare* and *Laminar* [4], [21], even if the placement of LSM hooks have already been carefully analyzed. This encourages us to verify the correctness of hooks placement for information flow tracking purposes.

VII. PERSPECTIVES AND CONCLUSION

In this paper, we have presented an approach to verify the suitability of LSM for the implementation of information flow trackers. We have shown that the compilation process can be leveraged to build a formal model of the source code to run single-purpose static analysis. Indeed, the compiler already needs to build a control flow graph, and maintains precious knowledge such as a points-to analysis to compile and optimize the code, and all this can be reused for more static analysis on the code. The clearest advantage of this method is the prevention of any risk of making errors in the semantics of the code. We designed and proved a reproducible static analysis on the position of the LSM hooks to track information flows. This analysis is simple, specifically adapted for its purpose and fast. We then examined the results and

Table I
PRESENCE OF PATHS ESCAPING ALL LSM HOOKS IN THE SYSTEM CALLS

read pread64	msgrcv	shmdt
readv preadv	msgsnd	vmsplice
writew pwritev	clone fork vfork	tee
write pwrite64	execve execveat	munmap
sendfile sendfile64	send sendto sendmsg	mq_timedsend
shmat	recv recvfrom recvmmsg	mq_timedreceive
mmap		

(a) System calls in which all paths are mediated by at least one LSM hook

(b) System calls producing information flows but having no LSM hooks

process_vm_readv	two cases of possible paths:
process_vm_writev	a thread reading another thread's memory in the same process and a thread reading its own memory
splice	no hooks for the pipe to pipe information flow no paths for the file to pipe information flow no paths for the pipe to file information flow
sendmmsg (same hooks as send)	the LSM hook is hit only once, for the first message, for the other messages sent with the same system call to the same destination, the hook is not called again
recvmmsg (same hooks as recv)	the LSM hook is hit only once, for the first message, for the other messages read with the same system call, the hook is not called again
ptrace	no paths for the attachment, but no hooks for the detachment

(c) System calls where not all paths are mediated

concluded that it was possible, after some modifications, to use LSM to track individual information flows. Finally, we have written a patch for the Linux kernel, version 4.3, to place the LSM hooks where we suggest. We validated the new placement by redoing the static analysis. The patch can be found at <http://kayrebt.gforge.inria.fr/pathexaminer.html>. This work validates the approach of developing information flow trackers on top of LSM and provides a strong basis for the formal assessment of their correctness.

We have shown that the LSM frameworks must be extended to cover all execution paths generating information flows on containers of information. This is a necessary condition to correctly implement information flow control but we plan also to consider further issues relative to the way the hooks are used. For example, when simultaneous information flows involving the same containers are considered, it is important that the tags are propagated in the same order than the flows are performed. These problems are already under investigation by the authors.

REFERENCES

- [1] A. Abraham, R. Andriatsimandefitra, A. Brunelat, J.-F. Lalande, and V. Viet Triem Tong, "GrodDroid: a gorilla for triggering malicious behaviors," IEEE, Oct. 2015.
- [2] H. Yin, D. Song, M. Egele, C. Kruegel, and E. Kirda, "Panorama: Capturing System-wide Information Flow for Malware Detection and Analysis," in *ACM Conference on Computer and Communications Security*. ACM, 2007.
- [3] L. K. Yan and H. Yin, "DroidScope: Seamlessly Reconstructing the OS and Dalvik Semantic Views for Dynamic Android Malware Analysis," in *USENIX Conference on Security Symposium*, ser. Security'12. USENIX Association, 2012.
- [4] L. Georges, V. V. T. Tong, and L. Mé, "Blare Tools: A Policy-Based Intrusion Detection System Automatically Set by the Security Policy," in *International Symposium on Recent Advances in Intrusion Detection (RAID 2009)*, 2009.
- [5] I. Roy, D. E. Porter, M. D. Bond, K. S. McKinley, and E. Witchel, "Laminar: Practical Fine-grained Decentralized Information Flow Control," *SIGPLAN Not.*, vol. 44, no. 6, 2009.
- [6] A. Nadkarni, B. Andow, W. Enck, and S. Jha, "Practical DIFC enforcement on Android," in *USENIX Security Symposium*, 2016.
- [7] C. Wright, C. Cowan, S. Smalley, J. Morris, and G. Kroah-Hartman, "Linux Security Modules: General Security Support for the Linux Kernel," in *USENIX Security Symposium*. USENIX Association, 2002.
- [8] P. Loscocco and S. Smalley, "Integrating Flexible Support for Security Policies into the Linux Operating System," in *FREENIX Track: 2001 USENIX Annual Technical Conference (FREENIX '01)*, 2001.
- [9] T. Jaeger, A. Edwards, and X. Zhang, "Consistency analysis of authorization hook placement in the Linux security modules framework," *ACM Trans. Inf. Syst. Secur.*, vol. 7, no. 2, 2004.
- [10] V. Ganapathy, T. Jaeger, and S. Jha, "Automatic Placement of Authorization Hooks in the Linux Security Modules Framework," in *ACM Conference on Computer and Communications Security*. ACM, 2005.
- [11] J. D. Guttman, A. L. Herzog, J. D. Ramsdell, and C. W. Skorupka, "Verifying information flow goals in security-enhanced Linux," *J. Comput. Secur.*, vol. 13, no. 1, 2005.
- [12] D. Muthukumar, T. Jaeger, and V. Ganapathy, "Leveraging 'choice' to automate authorization hook placement," *ACM Conference on Computer and Communications Security*, 2012.
- [13] R. Bodik, R. Gupta, and M. L. Soffa, "Refining Data Flow Information Using Infeasible Paths," *SIGSOFT Software Engineering Notes*, vol. 22, no. 6, Nov. 1997.
- [14] R. M. Stallman and the GCC developer community, "Using the GNU Compiler Collection (GCC)," *Tech. Rep.*, 2013. [Online]. <https://gcc.gnu.org/onlinedocs/gcc-4.8.4/gcc/>
- [15] J. Merrill, "GENERIC and GIMPLE: A new tree representation for entire functions," in *GCC Developers Summit*, 2003.
- [16] F. E. Allen, "Control Flow Analysis," in *Symposium on Compiler Optimization*. ACM, 1970.
- [17] D. Beyer, T. A. Henzinger, R. Jhala, and R. Majumdar, "The software model checker Blast," *International Journal on Software Tools for Technology Transfer*, vol. 9, no. 5, 2007.
- [18] B. Dutertre and L. de Moura, "The Yices SMT solver," *SRI International*, *Tech. Rep.*, 2006.
- [19] S. Smalley, "[PATCH 0/2] Add missing LSM hooks in mq_timed{send, receive} and splice, Email on the Linux Security Modules development mailing list," 2016. [Online]. <http://thread.gmane.org/gmane.linux.kernel.lsm/28737>
- [20] X. Zhang, A. Edwards, and T. Jaeger, "Using CQUAL for Static Analysis of Authorization Hook Placement," in *USENIX Security Symposium*. USENIX Association, 2002.
- [21] D. E. Porter, M. D. Bond, I. Roy, K. S. McKinley, and E. Witchel, "Practical Fine-Grained Information Flow Control Using Laminar," *ACM Transactions on Programming Languages and Systems*, vol. 37, no. 1, Nov. 2014.

APPENDIX

A Control Flow Graph is a pair $(\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the set of nodes in the graph, each one representing one instruction in the code, and $\mathcal{E}$ is the set of edges between nodes, representing the *may-follow* relation between instructions. The set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{C} \times \mathcal{V}$ is a relation between a *source node* and a *destination node*, labelled by a constraint. The constraint on an edge gives a condition on the variables of the program that must be satisfied at this point for the control flow go from the source node to the destination node. From all the edges going out from a given node, there is always exactly one constraint which is true. The set of constraints $\mathcal{C}$ is defined as

$$\begin{aligned} \mathcal{C} = & \{\mathbf{true}\} \\ & \cup \mathcal{Vars} \times \{=, \neq, <, >, \leq, \geq\} \times \mathcal{Vars} \\ & \cup \mathcal{Vars} \times \{=, \neq, <, >, \leq, \geq\} \times \mathbb{Z} \end{aligned}$$

There are five kinds of nodes. $\mathcal{V} = \mathcal{V}^{assign} \uplus \mathcal{V}^{mem} \uplus \mathcal{V}^{join} \uplus \mathcal{V}^\varphi \uplus \mathcal{V}^{call}$.

- $\mathcal{V}^{assign}$ is the set of direct assignment nodes: the left-hand side variable is assigned the value of the right-hand side expression,
- $\mathcal{V}^{mem}$ is the set of assignment-through-pointer nodes: the value of the right-hand side expression is assigned to the memory location identified by the value of the left-hand side pointer
- $\mathcal{V}^{join}$ is the set of join nodes: not an assignment node but has the particularity of having several incoming edges,
- $\mathcal{V}^\varphi$ is the set of phi nodes: the left-hand side variable is assigned the value of one of the right-hand side parameters, depending on the edge taken to reach the last previous join node, this kind of nodes is generated by the compiler when the value of a variable depends on the program control flow,
- $\mathcal{V}^{call}$ is the set of function call nodes: a function is called and its result is assigned to the left-hand side variable, if any,

Assembly code nodes (representing a piece of assembly code written in the C source file) are represented as function call nodes because they have similar properties in our semantics.

A. Definition of the Path Configuration

We define a path configuration as a pair $(n, v) \in \{m \in \mathbb{N} \mid m \geq 1\} \times \mathcal{V}$ where n is the index of the edge taken to reach the last join node visited and v is the last node visited. This is the information needed to interpret the Phi nodes. The path configuration is updated by the relation $\hookrightarrow \subseteq (\mathbb{N} \times \mathcal{V}) \times (\mathcal{V} \times \mathcal{C}) \times (\mathbb{N} \times \mathcal{V})$. We write: $(n_1, v_1) \xrightarrow{v, c} (n_2, v_2)$ to actually mean $((n_1, v_1), (v, c), (n_2, v_2)) \in \hookrightarrow$, which means that starting from a path configuration (n_1, v_1) , and going through the node v and then the edge following v bearing constraint c , we get the configuration (n_2, v_2) . $\hookrightarrow$ is defined as:

$$\begin{aligned} (n_1, v_1) & \xrightarrow{v, c} (n, v) & \text{if } v \in \mathcal{V}^{join} \text{ and } \text{pred}(v, n) = v_1 \\ (n_1, v_1) & \xrightarrow{v, c} (n_1, v) & \text{if } v \in \mathcal{V} \setminus \mathcal{V}^{join} \end{aligned}$$

B. Properties Satisfied by the Concrete Semantics

1) Definition of the Valuation of an Expression:

$$\forall x \in \mathcal{Vars} \quad \theta(x) = \begin{cases} \gamma(x) & \text{if } x \in \mathcal{Vars}^{temp} \\ \sigma(\alpha(x)) & \text{if } x \in \mathcal{Vars}^{mem} \end{cases}$$

$$\begin{aligned} \llbracket k \rrbracket^\theta &= k & \text{for all } k \in \mathbb{Z} & \quad \llbracket x \rrbracket^\theta = \theta(x) & \text{for all } x \in \mathcal{Vars} \\ \llbracket *p \rrbracket^\theta &= \sigma(\theta(p)) & \text{for all } p \in \mathcal{Vars}^{ptr} & \quad \llbracket \&x \rrbracket^\theta = \alpha(x) & \text{for all } x \in \mathcal{Vars}^{mem} \end{aligned}$$

2) *Rules*: The concrete executions are assumed to respect a transition relation $\rightarrow \subseteq \theta \times \theta$ between configurations. We assume the following minimal set of rules:

$$\begin{array}{lll} \mathcal{V}^{assign} & \theta = (\gamma, \sigma, \alpha) & \xrightarrow{x=e, c} (\gamma[x \leftarrow \llbracket e \rrbracket^\theta], \sigma, \alpha) \quad x \in \mathcal{Vars}^{temp} \\ & \theta = (\gamma, \sigma, \alpha) & \xrightarrow{x=e, c} (\gamma, \sigma[\alpha(x) \leftarrow \llbracket e \rrbracket^\theta], \alpha) \quad x \in \mathcal{Vars}^{mem} \\ & \theta & \xrightarrow{x=e, c} \theta \quad x \in \overline{\mathcal{Vars}} \\ \mathcal{V}^{mem} & \theta = (\gamma, \sigma, \alpha) & \xrightarrow{*p=e, c} (\gamma, \sigma[\llbracket p \rrbracket^\theta \leftarrow \llbracket e \rrbracket^\theta], \alpha) \\ \mathcal{V}^\varphi & \theta = (\gamma, \sigma, \alpha) & \xrightarrow{x=\text{PHI}(e_1, \dots, e_n), c} (\gamma[x \leftarrow \llbracket e_{n_{path}} \rrbracket^\theta], \sigma, \alpha) \\ \mathcal{V}^{call} & (\gamma_1, \sigma_1, \alpha) & \xrightarrow{x=f(e_1, \dots, e_n), c} (\gamma_2, \sigma_2, \alpha) \quad \text{with } \forall z \in \mathcal{Vars} \setminus \{x\} \quad \gamma_1(z) = \gamma_2(z) \\ & (\gamma_1, \sigma_1, \alpha) & \xrightarrow{f(e_1, \dots, e_n), c} (\gamma_2, \sigma_2, \alpha) \\ v \in \mathcal{V}^{join} \cup \{v_{init}\} & \theta & \xrightarrow{v, c} \theta \end{array}$$

3) Other Property Satisfied by the Concrete Semantics:

$$\forall \theta_1, \theta_2 \in \Theta \forall v \in \mathcal{V} \forall c \in \mathcal{C} \theta_1 \xrightarrow{v, c} \theta_2 \Rightarrow \theta_2 \models c \quad (1)$$

4) Satisfiability of a Constraint: We write $\theta \models c$ when a concrete configuration θ satisfies a constraint $c \in \mathcal{C}$:

$$\theta \models c \Leftrightarrow \begin{pmatrix} c = \text{true} \\ \vee c = (x, R, y) \in \mathcal{Vars} \times \mathcal{Rel} \times \mathcal{Vars} \wedge \theta(x) R \theta(y) \\ \vee c = (x, R, k) \in \mathcal{Vars} \times \mathcal{Rel} \times \mathbb{Z} \wedge \theta(x) R k \end{pmatrix} \quad (2)$$

where R stands for the usual interpretation of the relation symbols in $\mathcal{Rel} = \{=, \neq, \leq, \geq, <, >\}$.

C. Definition of the Static Semantics

1) Definition of Function reset:

$$\begin{aligned} \text{reset}(C, x) &= C \setminus \{\langle x, \cdot, \cdot \rangle, \langle \cdot, \cdot, x \rangle\} \\ \text{reset}(C, X) &= C \setminus (\{\langle x, \cdot, \cdot \rangle \mid x \in X\} \cup \{\langle \cdot, \cdot, x \rangle \mid x \in X\}) \end{aligned}$$

2) Rules: The following tables gives the semantics of each type of node.

$v \in \mathcal{V}^{assign}$		
$(C, P) \xrightarrow{x=e, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, e \rangle, c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, e \in \mathcal{Vars}^{\mathbb{Z}} \cup \mathbb{Z}$
$(C, P) \xrightarrow{x=*p, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, P(p) \rangle, c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, p \in \mathcal{Vars}^{ptr}, P(p) \neq \top$
$(C, P) \xrightarrow{x=*p, c}$	$(\text{reset}(C, x) \cup \{c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, p \in \mathcal{Vars}^{ptr}, P(p) = \top$
$(C, P) \xrightarrow{x=e, c}$	$(\text{reset}(C, x) \cup \{c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, e \in \mathbb{Q}$
$(C, P) \xrightarrow{x=e, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, e \rangle, c\}, P[x \leftarrow P(e)])$	$x \in \mathcal{Vars}^{ptr}, e \in \mathcal{Vars}^{ptr}$
$(C, P) \xrightarrow{x=*p, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, P(p) \rangle, c\}, P[x \leftarrow P(P(p))])$	$x \in \mathcal{Vars}^{ptr}, p \in \mathcal{Vars}^{ptr}, P(p) \neq \top$
$(C, P) \xrightarrow{x=*p, c}$	$(\text{reset}(C, x) \cup \{c\}, P[x \leftarrow \top])$	$x \in \mathcal{Vars}^{ptr}, p \in \mathcal{Vars}^{ptr}, P(p) = \top$
$(C, P) \xrightarrow{x=\&y, c}$	$(\text{reset}(C, x) \cup \{c\}, P[x \leftarrow y])$	$x \in \mathcal{Vars}^{ptr}, y \in \mathcal{Vars}^{mem}$
$(C, P) \xrightarrow{x=e, c}$	$(\text{reset}(C, x) \cup \{c\}, P[x \leftarrow \top])$	$x \in \mathcal{Vars}^{ptr}, e \in \mathbb{Q}$
$(C, P) \xrightarrow{x=e, c}$	$(C \cup \{c\}, P)$	$x \in \overline{\mathcal{Vars}}$
$v \in \mathcal{V}^{mem}$		
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, P(p)) \cup \{\langle P(p), =, e \rangle, c\}, P)$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}, e \in \mathcal{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P(p) \neq \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, \mathcal{O}(p)) \cup \{c\}, P)$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}, e \in \mathcal{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P(p) = \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, P(p)) \cup \{\langle P(p), =, e \rangle, c\}, P[P(p) \leftarrow P(e)])$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{ptr}, e \in \mathcal{Vars}^{ptr}, P(p) \neq \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, \mathcal{O}(p)) \cup \{c\}, P[q \leftarrow \top \mid q \in \mathcal{O}(p)])$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}, e \in \mathcal{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P(p) = \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, P(p)) \cup \{c\}, P)$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}, e \in \mathbb{Q}, P(p) \neq \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, \mathcal{O}(p)) \cup \{c\}, P)$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}, e \in \mathbb{Q}, P(p) = \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, P(p)) \cup \{c\}, P[P(p) \leftarrow \top])$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{ptr}, e \in \mathbb{Q}, P(p) \neq \top$
$(C, P) \xrightarrow{*p=e, c}$	$(\text{reset}(C, \mathcal{O}(p)) \cup \{c\}, P[q \leftarrow \top \mid q \in \mathcal{O}(p)])$	$\mathcal{O}(p) \subseteq \mathcal{Vars}^{ptr}, e \in \mathbb{Q}, P(p) = \top$
$v \in \mathcal{V}^{\varphi}$		
$(C, P) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, e_{n_{path}} \rangle, c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, e_{n_{path}} \in \mathbb{Z} \cup \mathcal{Vars}^{\mathbb{Z}}$
$(C, P) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$	$(\text{reset}(C, x) \cup \{c\}, P)$	$x \in \mathcal{Vars}^{\mathbb{Z}}, e_{n_{path}} \in \mathbb{Q}$
$(C, P) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$	$(\text{reset}(C, x) \cup \{\langle x, =, e_{n_{path}} \rangle, c\}, P[x \leftarrow P(e_{n_{path}})])$	$x \in \mathcal{Vars}^{ptr}, e_{n_{path}} \in \mathcal{Vars}^{ptr}$
$(C, P) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$	$(\text{reset}(C, x) \cup \{c\}, P[x \leftarrow y])$	$x \in \mathcal{Vars}^{ptr}, e_{n_{path}} = \&y, y \in \mathcal{Vars}^{mem}$
$(C, P) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$	$(\text{reset}(C, x) \cup \{c\}, P[x \leftarrow \top])$	$x \in \mathcal{Vars}^{ptr}, e_{n_{path}} \in \mathbb{Q}$

$$\begin{array}{lcl}
& v \in \mathcal{V}^{call} & \\
(C, P) & \xrightarrow{x=f(e_1, \dots, e_n), c} & (\text{reset}(C, \{x\} \cup \mathcal{Vars}^{mem}) \cup \{c\}, P[q \leftarrow \top \mid q \in \mathcal{Vars}^{mem} \cap \mathcal{Vars}^{ptr}]) \quad x \in \mathcal{Vars}^{\mathbb{Z}} \\
(C, P) & \xrightarrow{x=f(e_1, \dots, e_n), c} & (\text{reset}(C, \{x\} \cup \mathcal{Vars}^{mem}) \cup \{c\}, P[q \leftarrow \top \mid q \in \{x\} \cup (\mathcal{Vars}^{mem} \cap \mathcal{Vars}^{ptr})]) \quad x \in \mathcal{Vars}^{ptr} \\
(C, P) & \xrightarrow{x=f(e_1, \dots, e_n), c} & (\text{reset}(C, \mathcal{Vars}^{mem}) \cup \{c\}, P[q \leftarrow \top \mid q \in \mathcal{Vars}^{mem} \cap \mathcal{Vars}^{ptr}]) \quad x \in \overline{\mathcal{Vars}} \\
(C, P) & \xrightarrow{f(e_1, \dots, e_n), c} & (\text{reset}(C, \mathcal{Vars}^{mem}) \cup \{c\}, P[q \leftarrow \top \mid q \in \mathcal{Vars}^{mem} \cap \mathcal{Vars}^{ptr}]) \\
(C, P) & \xrightarrow{v, c} & (C \cup \{c\}, P) \quad v \in \mathcal{V}^{join} \cup \{v_{init}\}
\end{array}$$

D. Satisfiability of an Abstract Configuration

Definition 1: Satisfiability of an abstract configuration A concrete configuration θ satisfies an abstract one k (written $\theta \models k$), if, and only if:

$$\begin{array}{l}
\forall \theta = (\gamma, \sigma, \alpha) \in \Theta \\
\forall k = (C, P) \in K
\end{array}
\quad \theta \models k \Leftrightarrow \left(\begin{array}{l} \theta \models C \\ \wedge \forall x \in \mathcal{Vars}^{ptr} P(x) \neq \top \Rightarrow \alpha(P(x)) = \theta(x) \end{array} \right)$$

θ must satisfy every constraint and the values of the pointers known in k must be the same in the concrete configuration. The satisfiability of an abstract configuration $k \in K$ is written $\models k$ and is defined by: $\forall k \in K (\models k) \Leftrightarrow (\exists \theta \in \Theta \theta \models k)$.

1) *Extensions of the Relations over Paths:* The relations defined so far handle only individual components of paths. We extend them naturally to entire paths.

Definition 2: Transitive closure of $\rightarrow$ and $\twoheadrightarrow$ Given a path in the CFG $p = (v_1, c_1), \dots, (v_n, c_n) \in (\mathcal{V} \times \mathcal{C}) \times \dots \times (\mathcal{V} \times \mathcal{C})$, we define the transitive closure of $\rightarrow$ (respectively $\twoheadrightarrow$) along p noted $\rightarrow_p^*$ (respectively $\twoheadrightarrow_p^*$) as follows.

$$\begin{array}{l}
\bullet \theta \rightarrow_p^* \theta' \Leftrightarrow \exists \theta_1, \dots, \theta_{n-1} \theta \xrightarrow{v_1, c_1} \theta_1 \xrightarrow{v_2, c_2} \dots \xrightarrow{v_{n-1}, c_{n-1}} \theta_{n-1} \xrightarrow{v_n, c_n} \theta' \\
\bullet k \twoheadrightarrow_p^* k' \Leftrightarrow \exists k_1, \dots, k_{n-1} k \twoheadrightarrow k_1 \twoheadrightarrow k_2 \twoheadrightarrow \dots \twoheadrightarrow k_{n-1} \twoheadrightarrow k' \\
\text{where } \forall \mathbf{k} \in \{k, k_1, \dots, k_{n-1}\} \models \mathbf{k}
\end{array}$$

Note that we may have $\not\models k'$, the last configuration may be impossible.

E. Proof that the $\twoheadrightarrow$ Relation is Left-total

First, we state formally the left-totalness:

Property 1 (Left-totalness of $\twoheadrightarrow$):

$$\forall v \in \mathcal{V} \forall c \in \mathcal{C} \forall k \in K \exists k' \in K k \xrightarrow{v, c} k'$$

Proof: By case analysis, we prove easily that, for each kind of node, there is a corresponding definition of $\twoheadrightarrow$. For example, the assignment nodes are defined as:

$$\left\{ \begin{array}{l} x = e \quad \text{with } x \in \mathcal{Vars}^{\mathbb{Z}} \text{ and } e \in \mathbb{Z} \cup \mathcal{Vars}^{\mathbb{Z}} \cup \{ *p \mid p \in \mathcal{Vars}^{ptr} \wedge \mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}} \} \cup \textcircled{?} \\ x = e \quad \text{with } x \in \mathcal{Vars}^{ptr} \text{ and } e \in \mathcal{Vars}^{ptr} \cup \{ *p \mid p \in \mathcal{Vars}^{ptr} \wedge \mathcal{O}(p) \subseteq \mathcal{Vars}^{ptr} \} \\ \quad \cup \{ \&y \mid y \in \mathcal{Vars}^{mem} \} \cup \textcircled{?} \\ x = e \quad \text{with } x \in \overline{\mathcal{Vars}} \end{array} \right.$$

Hence, there must be a rule for:

- $x \in \mathcal{Vars}^{\mathbb{Z}}, e \in \mathbb{Z} \cup \mathcal{Vars}^{\mathbb{Z}}$
- $x \in \mathcal{Vars}^{\mathbb{Z}}, e = *p$ with $p \in \mathcal{Vars}^{ptr}$ and $\mathcal{O}(p) \subseteq \mathcal{Vars}^{\mathbb{Z}}$
  - $P(p) = \top$
  - or $P(p) \neq \top$
- $x \in \mathcal{Vars}^{\mathbb{Z}}, e \in \textcircled{?}$
- $x \in \mathcal{Vars}^{ptr}, e \in \mathcal{Vars}^{ptr}$
- $x \in \mathcal{Vars}^{ptr}, e = *p$ with $p \in \mathcal{Vars}^{ptr}$ and $\mathcal{O}(p) \subseteq \mathcal{Vars}^{ptr}$
  - $P(p) = \top$
  - or $P(p) \neq \top$
- $x \in \mathcal{Vars}^{ptr}, e = \&y$ with $y \in \mathcal{Vars}^{mem}$
- $x \in \mathcal{Vars}^{ptr}, e \in \textcircled{?}$

- $x \in \overline{\text{Vars}}$

And the semantics gives precisely a rule for each one of these cases, in this order. ■

The proposition is stated as:

Proposition 1 (Soundness): Given a path p in the CFG, we claim:

$$\forall \theta_1, \theta_2 \in \Theta \forall k_1, k_2 \in K \quad (\theta_1 \rightarrow_p^* \theta_2 \wedge k_1 \rightarrow_p^* k_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2$$

We prove the following lemma first:

lemma 1 (One-step soundness): Given a path p in the CFG, we claim:

$$\forall \theta_1, \theta_2 \in \Theta \forall k_1, k_2 \in K \forall v \in \mathcal{V} \forall c \in \mathcal{C} \quad (\theta_1 \xrightarrow{v,c} \theta_2 \wedge k_1 \xrightarrow{v,c} k_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2$$

Proof: We assume the existence of $\theta_1, \theta_2 \in \Theta$, $k_1, k_2 \in K$ and $v \in \mathcal{V}, c \in \mathcal{C}$ such that $\theta_1 \xrightarrow{v,c} \theta_2$, $k_1 \xrightarrow{v,c} k_2$, and $\theta_1 \Vdash k_1$. We suppose that $\theta_2 \not\Vdash k_2$.

In the following, we suppose that:

$$\begin{aligned} \theta_1 &= (\gamma_1, \sigma_1, \alpha_1) & k_1 &= (C_1, P_1) \\ \theta_2 &= (\gamma_2, \sigma_2, \alpha_2) & k_2 &= (C_2, P_2) \end{aligned}$$

Suppose that there exist $v \in \mathcal{V}$ and $c \in \mathcal{C}$ such that:

$$\begin{cases} \theta_1 \xrightarrow{v,c} \theta_2 \\ k_1 \xrightarrow{v,c} k_2 \\ \theta_1 \Vdash k_1 \\ \theta_2 \not\Vdash k_2 \end{cases}$$

We can thus deduce:

$$\begin{aligned} \exists c' \in C_2 \quad \theta_2 \not\Vdash c' \\ \forall \exists p \in \text{Vars}^{ptr} \quad P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p) \end{aligned}$$

We prove the lemma by case analysis on the relation $k_1 \xrightarrow{v,c} k_2$, showing a contradiction in the hypotheses in all cases.

1. **Case** $(C_1, P_1) \xrightarrow{x=e, c} (\text{reset}(C_1, x) \cup \{\langle x, =, e \rangle, c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, e \in \text{Vars}^{\mathbb{Z}} \cup \mathbb{Z}$.

In this case, we have

$$\forall y \in \text{Vars} \quad \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

- 1.1. **Case** $\exists c' \in \text{reset}(C_1, x) \quad \theta_2 \not\Vdash c'$.

By case analysis on c' .

- 1.1.1. **Case** $c' = \langle a, R, b \rangle$ **with** $a \in \text{Vars} \setminus x, b \in (\text{Vars} \cup \mathbb{Z}) \setminus x, R \in \text{Rel}$.

In this case, we have:

$$\begin{cases} \theta_2(a) = \theta_1(a) \\ \theta_2(b) = \theta_1(b) \\ \theta_1(a) R \theta_2(b) \end{cases}$$

from which we deduce $\theta_2(a) R \theta_2(b)$, and thus $\theta_2 \models c'$.

- 1.1.2. **Case** $c' = \text{true}$.

In this case, we have $\theta_2 \models c'$ by definition.

- 1.2. **Case** $\theta_2 \not\models c$.

By the property 1 p. 9, we have $\theta_2 \models c$.

- 1.3. **Case** $\exists p \in \text{Vars}^{ptr} \quad P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p)$.

We have $\alpha_2(P_2(p)) = \alpha_1(P_1(p))$ for all p because $P_2 = P_1$ and $\alpha_2 = \alpha_1$ as well $\theta_2(p) = \theta_1(p)$ because $p \neq x$ ($p \in \text{Vars}^{ptr}$ but $x \in \text{Vars}^{\mathbb{Z}} = \text{Vars} \setminus \text{Vars}^{ptr}$). Hence, the hypothesis $\forall p \in \text{Vars}^{ptr} \quad P_1(p) \neq \top \Rightarrow \alpha_1(P_1(p)) = \theta_1(p)$ entails $\forall p \in \text{Vars}^{ptr} \quad P_2(p) \neq \top \Rightarrow \alpha_2(P_2(p)) = \theta_2(p)$.

- 1.4. **Case** $\theta_2 \not\models \langle x, =, e \rangle$.

We have $\theta_2(x) = \theta_1(e) = \theta_2(e)$ by hypothesis, so $\theta_2 \models \langle x, =, e \rangle$.

2. **Case** $(C_1, P_1) \xrightarrow{x=*p, c} (\text{reset}(C_1, x) \cup \{\langle x, =, P_1(p) \rangle, c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, p \in \text{Vars}^{ptr}, P_1(p) \neq \top$.

In this case, we have

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \sigma_1(\theta_1(p)) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{and } \alpha_2 = \alpha_1$$

- 2.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.
 As above

- 2.2. **Case** $\theta_2 \not\models c$.
 As above

- 2.3. **Case** $\exists p \in \text{Vars}^{ptr} \ P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p)$.
 As above

- 2.4. **Case** $\theta_2 \not\models \langle x, =, P_1(p) \rangle$.
 We can deduce the following:

$$\begin{aligned} \theta_2(x) &= \sigma_1(\theta_1(p)) \\ &= \sigma_1(\alpha_1(P_1(p))) && \text{because } \theta_1 \models k_1 \\ &= \theta_1(P_1(p)) && \text{by definition of } \theta_1 \\ &= \theta_2(P_1(p)) && \text{by hypothesis if } P_1(p) \neq x, \\ &&& \text{trivially otherwise} \end{aligned}$$

and thus, $\theta_2 \models \langle x, =, P_1(p) \rangle$.

3. **Case** $(C_1, P_1) \xrightarrow{x=*p, c} (\text{reset}(C_1, x) \cup \{c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, p \in \text{Vars}^{ptr}, P_1(p) = \top$.

In this case, we have

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \sigma_1(\theta_1(p)) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{and } \alpha_2 = \alpha_1$$

- 3.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.
 As above

- 3.2. **Case** $\theta_2 \not\models c$.
 As above

- 3.3. **Case** $\exists p \in \text{Vars}^{ptr} \ P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p)$.
 As above

4. **Case** $(C_1, P_1) \xrightarrow{x=e, c} (\text{reset}(C_1, x) \cup \{c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, e \in \text{?}$.

In this case, we have

$$\forall y \in \text{Vars} \ \theta_2(y) = \theta_1(y) \text{ if } x \neq y \quad \text{and } \alpha_2 = \alpha_1$$

- 4.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.
 As above

- 4.2. **Case** $\theta_2 \not\models c$.
 As above

- 4.3. **Case** $\exists p \in \text{Vars}^{ptr} \ P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p)$.
 As above

5. **Case** $(C_1, P_1) \xrightarrow{x=e, c} (\text{reset}(C_1, x) \cup \{\langle x, =, e \rangle, c\}, P_1[x \leftarrow P_1(e)])$
 $x \in \mathcal{Vars}^{ptr}, e \in \mathcal{Vars}^{ptr}$.

In this case, we have

$$\forall y \in \mathcal{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{and } \alpha_2 = \alpha_1$$

- 5.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.

As above

- 5.2. **Case** $\theta_2 \not\models c$.

As above

- 5.3. **Case** $\exists p \in \mathcal{Vars}^{ptr} \ P_2(p) \neq \top \wedge \alpha_2(P_2(p)) \neq \theta_2(p)$.

Let us suppose that such a p exists. There are two cases here, either $p = x$ or $p \neq x$.

- 5.3.1. **Case** $p = x$.

In this case, $P_2(p) = P_2(x) = P_1(e)$. Assuming $P_1(e) \neq \top$, we have:

$$\begin{aligned} \alpha_2(P_2(p)) &= \alpha_2(P_2(x)) = \alpha_1(P_2(x)) \\ &= \alpha_1(P_1(e)) && \text{by hypothesis} \\ &= \theta_1(e) && \text{because } \theta_1 \Vdash k_1 \\ &= \theta_2(x) = \theta_2(p) && \text{by hypothesis} \end{aligned}$$

- 5.3.2. **Case** $p \neq x$.

In this case, we have:

$$\begin{aligned} \alpha_2(P_2(p)) &= \alpha_2(P_1(p)) = \alpha_1(P_1(p)) \\ &= \theta_1(p) && \text{because } \theta_1 \Vdash k_1 \\ &= \theta_2(p) && \text{by hypothesis, because } p \neq x \end{aligned}$$

- 5.4. **Case** $\theta_2 \not\models \langle x, =, e \rangle$.

We have $\theta_2(x) = \theta_1(e) = \theta_2(e)$, so $\theta_2 \models \langle x, =, e \rangle$.

6. **Case** $(C_1, P_1) \xrightarrow{x=*p, c} (\text{reset}(C_1, x) \cup \{\langle x, =, P_1(p) \rangle, c\}, P_1[x \leftarrow P_1(P_1(p))])$
 $x \in \mathcal{Vars}^{ptr}, p \in \mathcal{Vars}^{ptr}, P_1(p) \neq \top$.

In this case, we have

$$\forall y \in \mathcal{Vars} \ \theta_2(y) = \begin{cases} \sigma_1(\theta_1(p)) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{and } \alpha_2 = \alpha_1$$

- 6.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.

As above

- 6.2. **Case** $\theta_2 \not\models c$.

As above

- 6.3. **Case** $\exists p' \in \mathcal{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

Let us suppose that such a p' exists. There are two cases here, either $p' = x$ or $p' \neq x$.

- 6.3.1. **Case** $p' = x$.

In this case, $P_2(p') = P_2(x) = P_1(P_1(p))$. Assuming $P_1(P_1(p)) \neq \top$, we have:

$$\begin{aligned}
\alpha_2(P_2(p')) &= \alpha_2(P_2(x)) \\
&= \alpha_1(P_2(x)) \\
&= \alpha_1(P_1(P_1(p))) && \text{by hypothesis} \\
&= \theta_1(P_1(p)) && \text{because } \theta_1 \Vdash k_1 \\
\theta_2(p') &= \theta_2(x) \\
&= \sigma_1(\theta_1(p)) && \text{by hypothesis} \\
&= \sigma_1(\alpha_1(P_1(p))) && \text{because } \theta_1 \Vdash k_1 \\
&= \theta_1(P_1(p)) && P_1(p) \in \mathcal{Vars}^{mem} \text{ by definition}
\end{aligned}$$

6.3.2. Case $p \neq x$.

In this case, we have:

$$\begin{aligned}
\alpha_2(P_2(p')) &= \alpha_2(P_1(p')) \\
&= \alpha_1(P_1(p')) \\
&= \theta_1(p') && \text{because } \theta_1 \Vdash k_1 \\
&= \theta_2(p') && \text{by hypothesis, because } p' \neq x
\end{aligned}$$

6.4. Case $\theta_2 \not\models \langle x, =, P_1(p) \rangle$.

We have

$$\begin{aligned}
\theta_2(x) &= \sigma_1(\theta_1(p)) \\
&= \sigma_1(\alpha_1(P_1(p))) \\
&= \theta_1(P_1(p)) \\
&= \theta_2(P_1(p))
\end{aligned}$$

so $\theta_2 \models \langle x, =, P_1(p) \rangle$.

**7. Case $(C_1, P_1) \xrightarrow{x=*p, c} (\text{reset}(C_1, x) \cup \{c\}, P_1[x \leftarrow \top])$
 $x \in \mathcal{Vars}^{ptr}, p \in \mathcal{Vars}^{ptr}, P_1(p) = \top$.**

In this case, we have:

$$\forall y \in \mathcal{Vars} \ \theta_2(y) = \begin{cases} \sigma_1(\theta_1(p)) & \text{if } x = y \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

7.1. Case $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.

As above

7.2. Case $\theta_2 \not\models c$.

As above

7.3. Case $\exists p' \in \mathcal{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

Let us suppose that such a p' exists. There are two cases here, either $p' = x$ or $p' \neq x$.

7.3.1. Case $p' = x$.

We have $P_2(p) = P_2(x) = \top$, in contradiction with the hypothesis.

7.3.2. Case $p' \neq x$.

As above

**8. Case $(C_1, P_1) \xrightarrow{x=\&y, c} (\text{reset}(C_1, x) \cup \{c\}, P_1[x \leftarrow y])$
 $x \in \mathcal{Vars}^{ptr}, y \in \mathcal{Vars}^{mem}$.**

In this case, we have:

$$\forall z \in \mathcal{Vars} \ \theta_2(z) = \begin{cases} \alpha_1(y) & \text{if } x = z \\ \theta_1(z) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

8.1. **Case** $\exists c' \in \text{reset}(C_1, x) \ \theta_2 \not\models c'$.

As above

8.2. **Case** $\theta_2 \not\models c$.

As above

8.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

Let us suppose that such a p' exists. There are two cases here, either $p' = x$ or $p' \neq x$.

8.3.1. **Case** $p' = x$.

We have

$$\begin{aligned} \alpha_2(P_2(p')) &= \alpha_1(P_2(p')) \\ &= \alpha_1(P_2(x)) \\ &= \alpha_1(y) \\ \theta_2(p') &= \theta_2(x) \\ &= \alpha_1(y) \end{aligned}$$

so $\alpha_2(P_2(p')) = \theta_2(p')$.

8.3.2. **Case** $p' \neq x$.

As above

9. **Case** $(C_1, P_1) \xrightarrow{x=e, c} (\text{reset}(C_1, x) \cup \{c\}, P_1[x \leftarrow \top])$
 $x \in \text{Vars}^{ptr}, e \in \textcircled{?}$.

As case 7.

10. **Case** $(C_1, P_1) \xrightarrow{x=e, c} (C_1 \cup \{c\}, P_1)$
 $x \in \overline{\text{Vars}}$.

In this case, we have $\theta_2 = \theta_1$. The only non-trivial case to prove false is $\theta_2 \not\models c$. It is indeed false by property 1 p. 9.

11. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, P_1(p)) \cup \{P_1(p), =, e\}, P_1)$
 $O(p) \subseteq \text{Vars}^{\mathbb{Z}}, e \in \text{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P_1(p) \neq \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

11.1. **Case** $\exists c' \in \text{reset}(C_1, P_1(p)) \ \theta_2 \not\models c'$.

By case analysis on c' .

11.1.1. **Case** $c' = \langle a, R, b \rangle$ **with** $a \in \text{Vars} \setminus \{P_1(p)\}, b \in (\text{Vars} \cup \mathbb{Z}) \setminus \{P_1(p)\},$
 $R \in \text{Rel}$.

We know that $\forall y \in \text{Vars}^{mem} \ \theta_2(y) = \theta_1(y)$ if $\theta_1(p) \neq \alpha_1(y)$. We also know that $\theta_1(p) = \alpha_1(P_1(p))$ since $\theta_1 \models k_1$ and $P_1(p) \neq \top$. Memory locations are unique ($\forall y_1, y_2 \in \text{Vars}^{mem} \ \alpha_1(y_1) = \alpha_1(y_2) \Rightarrow y_1 = y_2$) so, since $a \neq P_1(p)$ and $b \neq P_1(p)$, we have:

$$\begin{cases} a \notin \text{Vars}^{mem} \vee \alpha_1(a) \neq \alpha_1(P_1(p)) \\ b \notin \text{Vars}^{mem} \vee \alpha_1(b) \neq \alpha_1(P_1(p)) \end{cases} \Rightarrow \begin{cases} \theta_2(a) = \theta_1(a) \\ \theta_2(b) = \theta_1(b) \\ \theta_1(a) R \theta_2(b) \end{cases}$$

from which we deduce $\theta_2(a) R \theta_2(b)$, and thus $\theta_2 \models c'$.

11.1.2. **Case** $c' = \text{true}$.

In this case, we have $\theta_2 \models c'$ by definition.

11.2. **Case** $\theta_2 \not\models c$.

As above

11.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

As above.

11.4. **Case** $\theta_2 \not\models \langle P_1(p), =, e \rangle$.

We have $\forall y \in \text{Vars} \ \theta_2(y) = \theta_2(e)$ if $\theta_1(p) = \alpha_1(y)$. But we know that $\theta_1(p) = \alpha_1(P_1(p))$ since $\theta_1 \Vdash k_1$ and $P_1(p) \neq \top$. So, we conclude that:

$$\theta_2(P_1(p)) = \theta_2(e)$$

from which we conclude $\theta_2 \models \langle P_1(p), =, e \rangle$.

12. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, \mathcal{O}(p)) \cup \{c\}, P_1)$
 $\mathcal{O}(p) \subseteq \text{Vars}^{\mathbb{Z}}, e \in \text{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P_1(p) = \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

12.1. **Case** $\exists c' \in \text{reset}(C_1, \mathcal{O}(p)) \ \theta_2 \not\models c'$.

By case analysis on c' .

12.1.1. **Case** $c' = \langle a, R, b \rangle$ **with** $a \in \text{Vars} \setminus \mathcal{O}(p), b \in (\text{Vars} \cup \mathbb{Z}) \setminus \mathcal{O}(p), R \in \text{Rel}$.

In this case, we have $\theta_2(a) = \theta_1(a)$ since $\forall y \in \text{Vars} \ \theta_2(y) = \theta_1(y)$ if $y \in \mathcal{O}(p)$ and $a \notin \mathcal{O}(p)$. Similarly we have $\theta_2(b) = \theta_1(b)$ and we also have by hypothesis $\theta_1(a) R \theta_1(b)$. Thus, we deduce $\theta_2(a) R \theta_2(b)$, and finally $\theta_2 \models c'$.

12.1.2. **Case** $c' = \text{true}$.

In this case, we have $\theta_2 \models c'$ by definition.

12.2. **Case** $\theta_2 \not\models c$.

As above.

12.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

As above.

13. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, P_1(p)) \cup \{\langle P_1(p), =, e \rangle, c\}, P_1[P_1(p) \leftarrow P_1(e)])$
 $\mathcal{O}(p) \subseteq \text{Vars}^{ptr}, e \in \text{Vars}^{ptr}, P_1(p) \neq \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases} \quad \text{with } \alpha_2 = \alpha_1$$

13.1. **Case** $\exists c' \in \text{reset}(C_1, P_1(p)) \ \theta_2 \not\models c'$.

As above.

13.2. **Case** $\theta_2 \not\models c$.

As above.

13.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

Two cases: $p' = P_1(p)$ or $p' \neq P_1(p)$.

13.3.1. **Case** $p' = P_1(p)$.

$$\begin{aligned}
\alpha_2(P_2(p')) &= \alpha_1(P_2(p')) \\
&= \alpha_1(P_2(P_1(p))) \\
&= \alpha_1(P_1(e)) \\
&= \theta_1(e) \\
&= \theta_2(P_1(p)) \\
&= \theta_2(p')
\end{aligned}$$

13.3.2. **Case** $p' \neq P_1(p)$.

$$\begin{aligned}
\alpha_2(P_2(p')) &= \alpha_1(P_2(p')) \\
&= \alpha_1(P_1(p')) \\
&= \theta_1(p') \\
&= \theta_2(p')
\end{aligned}$$

13.4. **Case** $\theta_2 \not\models \langle P_1(p), =, e \rangle$.

As in case 11.

14. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, \mathcal{O}(p)) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \mathcal{O}(p)])$
 $\mathcal{O}(p) \subseteq \text{Vars}^{\mathbb{Z}}, e \in \text{Vars}^{\mathbb{Z}} \cup \mathbb{Z}, P_1(p) = \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases}$$

with $\alpha_2 = \alpha_1$

14.1. **Case** $\exists c' \in \text{reset}(C_1, \mathcal{O}(p)) \ \theta_2 \not\models c'$.

As in case 12.

14.2. **Case** $\theta_2 \not\models c$.

As above.

14.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

$P_2(p') \neq \top$ implies that $p' \notin \mathcal{O}(p)$, which implies in turn implies that $\theta_1(p) \neq \alpha_1(p')$. Hence, we have:

$$\begin{aligned}
\alpha_2(P_2(p')) &= \alpha_1(P_2(p')) \\
&= \alpha_1(P_1(p')) && \text{since } p' \notin \mathcal{O}(p) \\
&= \theta_1(p') \\
&= \theta_2(p') && \text{since } \theta_1(p) \neq \alpha_1(p')
\end{aligned}$$

15. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, P_1(p)) \cup \{c\}, P_1)$
 $\mathcal{O}(p) \subseteq \text{Vars}^{\mathbb{Z}}, e \in \text{?}, P_1(p) \neq \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases}$$

with $\alpha_2 = \alpha_1$

15.1. **Case** $\exists c' \in \text{reset}(C_1, \mathcal{O}(p)) \ \theta_2 \not\models c'$.

As in case 11.

15.2. **Case** $\theta_2 \not\models c$.

As above.

15.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.
As in case 11.

16. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, \mathcal{O}(p)) \cup \{c\}, P_1)$
 $\mathcal{O}(p) \subseteq \text{Vars}^{\mathbb{Z}}, e \in \textcircled{?}, P_1(p) = \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases}$$

with $\alpha_2 = \alpha_1$

16.1. **Case** $\exists c' \in \text{reset}(C_1, \mathcal{O}(p)) \ \theta_2 \not\models c'$.
As in case 12.

16.2. **Case** $\theta_2 \not\models c$.
As above.

16.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.
As in case 11.

17. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, P_1(p)) \cup \{c\}, P_1[P_1(p) \leftarrow \top])$
 $\mathcal{O}(p) \subseteq \text{Vars}^{ptr}, e \in \textcircled{?}, P_1(p) \neq \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases}$$

with $\alpha_2 = \alpha_1$

17.1. **Case** $\exists c' \in \text{reset}(C_1, P_1(p)) \ \theta_2 \not\models c'$.
As in case 11.

17.2. **Case** $\theta_2 \not\models c$.
As above.

17.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.
The only subcase is $p' \neq P_1(p)$ since $P_2(p') \neq \top$ and $P_2(P_1(p)) = \top$. This case is solved as case 13.3.2.

18. **Case** $(C_1, P_1) \xrightarrow{*p=e, c} (\text{reset}(C_1, \mathcal{O}(p)) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \mathcal{O}(p)])$
 $\mathcal{O}(p) \subseteq \text{Vars}^{ptr}, e \in \textcircled{?}, P_1(p) = \top$.

In this case, we have:

$$\forall y \in \text{Vars} \ \theta_2(y) = \begin{cases} \theta_1(e) & \text{if } y \in \text{Vars}^{mem} \text{ and } \theta_1(p) = \alpha_1(y) \\ \theta_1(y) & \text{otherwise} \end{cases}$$

with $\alpha_2 = \alpha_1$

18.1. **Case** $\exists c' \in \text{reset}(C_1, \mathcal{O}(p)) \ \theta_2 \not\models c'$.
As in case 11.

18.2. **Case** $\theta_2 \not\models c$.
As above.

18.3. **Case** $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.
As in case 14.3.

19. **Case** $(C_1, P_1) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c} (\text{reset}(C_1, x) \cup \{\langle x, =, e_{n_{path}} \rangle, c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, e_{n_{path}} \in \mathbb{Z} \cup \text{Vars}^{\mathbb{Z}}$.

This case is identical to case 1. with $e = e_{n_{path}}$.

20. **Case** $(C_1, P_1) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c} (\text{reset}(C_1, x) \cup \{c\}, P_1)$
 $x \in \text{Vars}^{\mathbb{Z}}, e_{n_{path}} \in \text{?}$.

This case is identical to case 4. with $e = e_{n_{path}}$.

21. **Case** $(C_1, P_1) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c}$
 $(\text{reset}(C_0, x) \cup \{\langle x, =, e_{n_{path}} \rangle, c\}, P_1[x \leftarrow P_1(e_{n_{path}})])$
 $x \in \text{Vars}^{ptr}, e_{n_{path}} \in \text{Vars}^{ptr}$.

This case is identical to case 5. with $e = e_{n_{path}}$.

22. **Case** $(C_1, P_1) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c} (\text{reset}(C_1, x) \cup \{c\}, P_1[x \leftarrow y])$
 $x \in \text{Vars}^{ptr}, e_{n_{path}} = \& y, y \in \text{Vars}^{mem}$.

This case is identical to case 8. with $e = e_{n_{path}}$.

23. **Case** $(C_1, P_1) \xrightarrow{x=\text{PHI}\langle e_1, \dots, e_n \rangle, c} (\text{reset}(C_1, x) \cup \{c\}, P_1[x \leftarrow \top])$
 $x \in \text{Vars}^{ptr}, e_{n_{path}} \in \text{?}$.

This case is identical to case 9. with $e = e_{n_{path}}$.

24. **Case** $(C_1, P_1) \xrightarrow{x=f(e_1, \dots, e_n), c} (\text{reset}(C_1, \{x\} \cup \text{Vars}^{mem}) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \text{Vars}^{mem} \cap \text{Vars}^{ptr}])$
 $x \in \text{Vars}^{\mathbb{Z}}$.

In this case, we have:

$$\forall y \in \text{Vars}^{temp} \setminus \{x\} \theta_2(y) = \theta_1(y) \text{ and } \alpha_2 = \alpha_1$$

We have the following cases.

- 24.1. **Case** $\exists c' \in \text{reset}(C_1, \{x\} \cup \text{Vars}^{mem}) \theta_2 \models c'$.

By case analysis on c' .

- 24.1.1. **Case** $c' = \langle a, R, b \rangle$ **with** $a \in \text{Vars} \setminus (\{x\} \cup \text{Vars}^{mem}), b \in (\text{Vars} \cup \mathbb{Z}) \setminus (\{x\} \cup \text{Vars}^{mem}), R \in \text{Rel}$.

In this case, we have $\theta_2(a) = \theta_1(a)$ since $\forall y \in \text{Vars} \theta_2(y) = \theta_1(y)$ if $y \in \text{Vars}^{temp} \setminus \{x\}$ and $a \notin \text{Vars}^{mem}$ and $a \neq x$. Similarly we have $\theta_2(b) = \theta_1(b)$ and we also have by hypothesis $\theta_1(a) R \theta_1(b)$. Thus, we deduce $\theta_2(a) R \theta_2(b)$, and finally $\theta_2 \models c'$.

- 24.1.2. **Case** $c' = \text{true}$.

In this case, we have $\theta_2 \models c'$ by definition.

- 24.2. **Case** $\theta_2 \not\models c$.

As above.

- 24.3. **Case** $\exists p' \in \text{Vars}^{ptr} P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$.

We have $p' \in \text{Vars}^{temp}$ since $P_1(q) = \top$ for all $q \in \text{Vars}^{mem} \cap \text{Vars}^{ptr}$. We also have $p' \neq x$ since $x \notin \text{Vars}^{ptr}$. So finally, we have:

$$\begin{aligned} \alpha_2(P_2(p')) &= \alpha_1(P_2(p')) \\ &= \alpha_1(P_1(p')) \\ &= \theta_1(p') &= \theta_2(p') \end{aligned}$$

25. **Case** $(C_1, P_1) \xrightarrow{x=f(e_1, \dots, e_n), c} (\text{reset}(C_1, \{x\} \cup \text{Vars}^{mem}) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \{x\} \cup (\text{Vars}^{mem} \cap \text{Vars}^{ptr})])$
 $x \in \text{Vars}^{ptr}$.

As case 24. with the difference that for the subcase $\exists p' \in \text{Vars}^{ptr} \ P_2(p') \neq \top \wedge \alpha_2(P_2(p')) \neq \theta_2(p')$, we have $p' \neq x$ by hypothesis directly.

26. **Case** $(C_1, P_1) \xrightarrow{x=f(e_1, \dots, e_n), c} (\text{reset}(C_1, \text{Vars}^{mem}) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \text{Vars}^{mem} \cap \text{Vars}^{ptr}])$
 $x \in \overline{\text{Vars}}.$

As case 24. with the difference that for the subcase $\exists c' \in \text{reset}(C_1, \text{Vars}^{mem}) \ \theta_2 \models c'$, in the subsubcase $c' = \langle a, R, b \rangle$, we have $a \neq x$ and $b \neq x$ by hypothesis since $a \in \text{Vars}$ and $b \in \text{Vars}$.

27. **Case** $(C_1, P_1) \xrightarrow{f(e_1, \dots, e_n), c} (\text{reset}(C_1, \text{Vars}^{mem}) \cup \{c\}, P_1[q \leftarrow \top \mid q \in \text{Vars}^{mem} \cap \text{Vars}^{ptr}]).$

As case 26. with the difference that $\forall y \in \text{Vars}^{temp} \ \theta_2(y) = \theta_1(y)$, which actually makes the proof slightly simpler.

28. **Case** $(C_1, P_1) \xrightarrow{v \in \mathcal{V}^{join} \cup \{v_{init}\}, c} (C_1, P_1).$

In this case, we have $\forall y \in \text{Vars} \ \theta_2(y) = \theta_1(y)$. So, trivially, as $\theta_2 = \theta_1$ and $k_2 = k_1$, we have $\theta_1 \Vdash k_1 \Rightarrow \theta_2 \Vdash k_2$.

Finally, we have proven that it was impossible to exhibit a constraint $c' \in k_2$ such that $\theta_2 \not\models c'$ or a pointer $p' \in \text{Vars}^{ptr}$ such that $P_2(p') \neq \top$ and $\alpha_2(P_2(p')) \neq \theta_2(p')$. We thus conclude that:

$$\forall v \in \mathcal{V} \ \forall c \in \mathcal{C}; (k_1 \xrightarrow{v, c} k_2 \wedge \theta_1 \xrightarrow{v, c} \theta_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2$$

■

F. Proof of the proposition

We now prove the proposition:

$$\begin{aligned} & \forall p \in ((\mathcal{V} \times \mathcal{C}) \times \dots \times (\mathcal{V} \times \mathcal{C})) \text{ a path of finite length} \\ & \forall k_1, k_2 \in K \ \forall \theta_1, \theta_2 \in \Theta \\ & (k_1 \rightarrow_p^* k_2 \wedge \theta_1 \rightarrow_p^* \theta_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2 \end{aligned}$$

Proof: By induction on the length of the path p .

1. Case Length of 1.

In this case, there exists $v \in \mathcal{V}$ and $c \in \mathcal{C}$ such that $p = ((v, c))$ and we are left with proving that

$$(k_1 \xrightarrow{v, c} k_2 \wedge \theta_1 \xrightarrow{v, c} \theta_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2$$

which is exactly the result of the previous lemma.

2. Case We assume the result true for a length of n .

3. Case Length of $n + 1$.

In this case, we can write the path p as a path p' of length n followed by a last pair $(v, c) \in \mathcal{V} \times \mathcal{C}$. We must thus prove:

$$\forall k' \in K \ \forall \theta \in \Theta \ (k_1 \rightarrow_p^* k' \xrightarrow{v, c} k_2 \wedge \theta_1 \rightarrow_p^* \theta' \xrightarrow{v, c} \theta_2 \wedge \theta_1 \Vdash k_1) \Rightarrow \theta_2 \Vdash k_2$$

By the induction hypothesis, we have that $\theta' \Vdash k'$, and by the lemma, we have thus $\theta_2 \Vdash k_2$.

■

The proposition is stated as follows:

Proposition 2 (No execution paths match impossible paths): Given $\mathbf{P}$ the set of paths in some CFG:

$$\text{Impossib}(\mathbf{P}) \cap \text{Exec}(\mathbf{P}) = \emptyset$$

Proof: Suppose there exists $p \in \text{Unsatisf}(\mathbf{P}) \cap \text{Exec}(\mathbf{P})$. Since $p \in \text{Exec}(\mathbf{P})$, there exists $\theta_1, \theta_2 \in \Theta$ such that $p \in \text{Exec}(\mathbf{P})$ and $\theta_1 \rightarrow_p^* \theta_2$. Let $k_1 \in K$ be an abstract configuration such that $\theta_1 \Vdash k_1$ (such a configuration always exists since it suffices to consider (C, P) where $C = \{\text{true}\}$ and $P = x \mapsto \diamond$ which clearly satisfies this property). Since $\rightarrow$ is a left-total relation, there exists an abstract configuration $k_2 \in K$ such that $k_1 \rightarrow_p^* k_2$ and therefore, by soundness property, we get $\theta_2 \Vdash k_2$. However, since $p \in \text{Impossib}(\mathbf{P})$ we have $\not\models k_2$, which leads to a contradiction. ■

The proposition is stated as follows:

Proposition 3 (Correctness of the loop abstractor): We consider a loop $l = (\mathcal{V}_l, v_l, \mathcal{T}_l)$. We call p a path in l from v_l to v_l (so, a finite number of iterations of the loop). We claim:

$$\forall k \in K (\exists k_p \in K k \xrightarrow{*}_p k_p \wedge \theta \Vdash k_p) \Rightarrow \theta \Vdash k^l$$

To prove it, we first introduce an order relation on abstract configuration and the property that this order is compatible with the satisfiability:

Definition 3: Order relation $\preceq$ on abstract configurations We define an order relation $\preceq \subseteq K \times K$ on abstract configurations by:

$$(C, P) \preceq (C', P') \Leftrightarrow C \subseteq C' \wedge (\forall p \in \text{Vars}^{ptr} P(p) = \top \vee P(p) = P'(p))$$

Proposition 4 ($\preceq$ vs. $\Vdash$):

$$\forall \theta \in \Theta \forall k, k' \in K k \preceq k' \Rightarrow (\theta \Vdash k' \Rightarrow \theta \Vdash k)$$

Proof: Suppose that we have $k \preceq k'$, $\theta \Vdash k'$ but $\theta \not\Vdash k$. Then, assuming $\theta = (\gamma, \sigma, \alpha)$, $k = (C, P)$ and $k' = (C', P')$, we know that

$$(\exists c \in C \theta \not\models c) \vee (\exists p \in \text{Vars}^{ptr} P(p) \neq \top \wedge \alpha(p) \neq \theta(p))$$

By case analysis:

- $\exists c \in C \theta \not\models c$. If $c \in C$, then $c \in C'$ since $C \subseteq C'$ by hypothesis. But we know that $\theta \Vdash C'$, so $\theta \models c$, hence a contradiction.
- $\exists p \in \text{Vars}^{ptr} P(p) \neq \top \wedge \alpha(p) \neq \theta(p)$. If $P(p) \neq \top$ then $P'(p) = P(p)$ by hypothesis. But we know that $\forall p \in \text{Vars}^{ptr} P'(p) \neq \top \Rightarrow \alpha(p) = \theta(p)$, so we have $P(p) \neq \top \wedge \alpha(p) = \theta(p)$, a contradiction with the hypotheses. ■

To prove the proposition, we consider, without loss of generality, a loop $l = (\mathcal{V}_l, v_l, \mathcal{T}_l)$. We now introduce a technical lemma on the abstract configurations reachable inside the loop.

lemma 2 (The loop abstractor is smaller than any reachable configuration):

$$\forall (v, c, v') \in \mathcal{T}_l \forall k, k' \in K k \xrightarrow{v, c} k' \Rightarrow k^l \preceq k'$$

Proof: We suppose that $k = (C, P)$, $k^l = (C^l, P^l)$ and $k' = (C', P')$. We must prove:

- 1) $C^l \subseteq C'$
- 2) $\forall p \in \text{Vars}^{ptr} P^l(p) = \top \vee P^l(p) = P'(p)$
- 1) $C \subseteq C'$. By case analysis on the nodes v , and by the definition of the semantic rules, it is clear that the only constraints in C and not in C' are constraints removed by reset and these are only constraints concerning the left-hand side variable in an assignment, the variable equal to $P(p)$ or $\mathcal{O}(p)$ in case of an assignment through a pointer p , or Vars^{mem} in case of a function call. Furthermore, we note that $P(p) \in \text{Vars}^{mem}$ and $\mathcal{O}(p) \subseteq \text{Vars}^{mem}$ when $p \in \text{Vars}^{ptr}$. By the definition of C^l , these constraints in C do not appear in C^l . As $C^l \subseteq C$, we conclude therefore that $C^l \subseteq C'$.
- 2) $\forall p \in \text{Vars}^{ptr} P^l(p) = \top \vee P^l(p) = P'(p)$ All pointers p such that $P(p) \neq \top \wedge P'(p) \neq P(p)$ are pointers assigned directly or pointers in Vars^{mem} assigned through a pointer. But in that case, by definition, we have $P^l(p) = \top$ so we can conclude that $\forall p \in \text{Vars}^{ptr} P^l(p) = \top \vee P^l(p) = P'(p)$. ■

We can now prove a generalized version of the proposition.

lemma 3: Let p a path in $\mathcal{T}_l$. We claim:

$$\forall k, k' \in K \forall \theta \in \Theta (k \xrightarrow{*}_p k' \wedge \theta \Vdash k') \Rightarrow \theta \Vdash k^l$$

Proof: Since $\preceq$ is a relation order, it is transitive and by induction on the length of path p , we can conclude from lemma 2 that $k \xrightarrow{*}_p k' \Rightarrow k^l \preceq k'$. From there, by proposition 4, we have $\forall \theta \in \Theta \theta \Vdash k' \Rightarrow \theta \Vdash k^l$. ■

This final lemma leads trivially to the proposition as stated.

G. Proof that $\equiv$ Is an Equivalence Relation

For all p_1, p_2 two paths, we have $p_1 \equiv p_2 \Leftrightarrow p_1^\uparrow = p_2^\uparrow$. Since the equality is reflexive, transitive and symmetric, so is the loop-equivalence $\equiv$.

H. Existence and Unicity of the Normal Form

Suppose that $\mathbf{P}^{v_{flow}, V_{LSM}}$ is not empty. Then, let us consider (without loss of generality) S , a loop-equivalence class in $\mathbf{P}^{v_{flow}, V_{LSM}}$. S is non-empty by construction.

Existence Let $p \in S$. We have $p^\uparrow \in S$ by definition, and $(p^\uparrow)^\uparrow = p^\uparrow$ by definition of $\bullet^\uparrow$. Thus, the normal form $p^\uparrow$ exists in S .

Uniqueness Let $p_1, p_2 \in S$ such that $p_1^\uparrow = p_1$ and $p_2^\uparrow = p_2$. By definition of S , we have $p_1^\uparrow = p_2^\uparrow$ and thus, $p_1 = p_2$.

I. Proof that Analyzing the Normal Forms is Sufficient

Let us consider $S \in \mathbf{P}_{\text{flow}, V_{LSM}}^{\text{flow}, V_{LSM}}$ a loop-equivalence class and $p \in S$ a path of S . We want to prove that $[S] \in \text{Impossib}(\mathbf{P})$ implies $p \in \text{Impossib}(\mathbf{P})$.

We first prove that the abstract semantics' rules preserve the order on abstract configurations.

lemma 4 (Preservation of the Order on Abstract Configurations):

$$\forall k_1, k'_1, k_2, k'_2 \in K \forall v \in \mathcal{V} \forall c \in \mathcal{C} \left(\left\{ \begin{array}{l} k_1 \xrightarrow{v,c} k_2 \\ k'_1 \xrightarrow{v,c} k'_2 \end{array} \right. \wedge k'_1 \preceq k_1 \right) \Rightarrow k'_2 \preceq k_2$$

Proof: We suppose that $k_1 = (C_1, P_1)$, $k_2 = (C_2, P_2)$, $k'_1 = (C'_1, P'_1)$, $k'_2 = (C'_2, P'_2)$, and $k'_1 \preceq k_1$.

Let $q \in \text{Vars}^{ptr}$ a pointer. By the definition of the semantic rules, at least one of the four following cases applies.

$$1. \text{ Case } \begin{cases} P'_2(q) = P'_1(q) \\ P_2(q) = P_1(q) \end{cases}.$$

This is the case when the transition (v, c) does not change the information about which variable q points to. Since $k'_1 \preceq k_1$, we have $P'_1(q) = \top$ or $P'_1(q) = P_1(q)$, which yields to $P'_2(q) = \top$ or $P'_2(q) = P_2(q)$.

$$2. \text{ Case } \begin{cases} P'_2(q) = \top \\ P_2(q) = \top \end{cases}.$$

This is the case when the transition (v, c) changes the information about the variable q to $\top$. In this case, we trivially have $P'_2(q) = \top$ or $P'_2(q) = P_2(q)$.

$$3. \text{ Case } \exists q' \in \text{Vars}^{ptr} \begin{cases} P'_2(q) = P'_1(q') \\ P_2(q) = P_1(q') \end{cases}.$$

This is the case when the transition (v, c) changes the information about which variable q points to to the destination of another pointer q' . Since $k'_1 \preceq k_1$, we have $P'_1(q') = \top$ or $P'_1(q') = P_1(q')$, which yields to $P'_2(q) = \top$ or $P'_2(q) = P_2(q)$.

$$4. \text{ Case } \exists v \in \text{Vars}^{mem} \begin{cases} P'_2(q) = v \\ P_2(q) = v \end{cases}.$$

This is the case when the transition (v, c) changes the information about which variables q points to to a variable v . In this case, we trivially have $P'_2(q) = \top$ or $P'_2(q) = P_2(q) = v$.

Let $C_r, C_a \subseteq \mathcal{C}$ the set of constraints respectively subtracted and added to C_1 and C'_1 by the transition (v, c) (depending on the definition of the corresponding semantic rule). We have $\begin{cases} C_2 = (C_1 \setminus C_r) \cup C_a \\ C'_2 = (C'_1 \setminus C_r) \cup C_a \end{cases}$. We have therefore $C'_1 \subseteq C'_2 \Rightarrow C_1 \subseteq C_2$.

We have proven that $C'_2 \subseteq C_2$ and $\forall q \in \text{Vars}^{ptr} (P'_2(q) = \top \vee P'_2(q) = P_2(q))$, which means $k'_2 \preceq k_2$. ■

We extend the previous lemma to paths of finite length.

lemma 5 (Preservation of the Order on Abstract Configurations Along a Path):

$$\forall k_1, k'_1, k_2, k'_2 \in K \forall p \in \mathbf{P} \left(\left\{ \begin{array}{l} k_1 \xrightarrow{*}_p k_2 \\ k'_1 \xrightarrow{*}_p k'_2 \end{array} \right. \wedge k'_1 \preceq k_1 \right) \Rightarrow k'_2 \preceq k_2$$

Proof: By induction on the length of p (a path of length $n + 1$ is a path of length n followed by a last transition (v, c)), noting that since $\preceq$ is an order relation, it is transitive. ■

The lemma 2, p. 21 gives a similar result about loops.

Proposition 5 (Analyzing the normal form in each class is sufficient): Given $S \subseteq \mathbf{P}_{\text{flow}, V_{lsm}}^{\text{flow}, V_{lsm}}$ one loop-equivalence class of $\mathbf{P}_{\text{flow}, V_{lsm}}^{\text{flow}, V_{lsm}}$ and $[s] \in S$ the normal form of S , we have $[s] \in \text{Impossib}(\mathbf{P}) \Rightarrow S \subseteq \text{Impossib}(\mathbf{P})$.

Proof: Let us consider a path $p \in S$. We have by definition $p^\dagger = [s]$. Suppose that $[s] \in \text{Impossib}(\mathbf{P})$ but $p \notin \text{Impossib}(\mathbf{P})$. We can therefore suppose it exists $k_1, k_2 \in K$ $k_1 \xrightarrow{*}_p k_2$ such that $\models k_2$. p is a succession of subpaths, some being cycles. By definition, $[s]$ is the same path, without the cycles. By induction on the length of the sequence of subpaths in p and per lemma 5 for subpaths which are not loops, and per lemma 2 for subpaths which are loops, and by transitivity of $\preceq$, there exists $k'_2 \in K$ such that $k_1 \xrightarrow{*}_{[s]} k'_2$ and $k'_2 \preceq k_2$. By proposition 4, it follows that $\models [s]$, a contradiction with the hypotheses. ■