

Stack Size Analysis for Interrupt-driven Programs*

Krishnendu Chatterjee¹, Di Ma², Rupak Majumdar¹,
Tian Zhao³, Thomas A. Henzinger¹, and Jens Palsberg²

¹ University of California, Berkeley
{c.krish, rupak, tah}@cs.berkeley.edu

² Purdue University

{madi, palsberg}@cs.purdue.edu

³ University of Wisconsin, Milwaukee
tzhao@cs.uwm.edu

Abstract. We study the problem of determining stack boundedness and the exact maximum stack size for three classes of interrupt-driven programs. Interrupt-driven programs are used in many real-time applications that require responsive interrupt handling. In order to ensure responsiveness, programmers often enable interrupt processing in the body of lower-priority interrupt handlers. In such programs a programming error can allow interrupt handlers to be interrupted in cyclic fashion to lead to an unbounded stack, causing the system to crash. For a restricted class of interrupt-driven programs, we show that there is a polynomial-time procedure to check stack boundedness, while determining the exact maximum stack size is PSPACE-complete. For a larger class of programs, the two problems are both PSPACE-complete, and for the largest class of programs we consider, the two problems are PSPACE-hard and can be solved in exponential time.

1 Introduction

Most embedded software runs on resource-constrained processors, often for economic reasons. Once the processor, RAM, etc. have been chosen for an embedded system, the programmers has to fit everything into the available space. For example, on a Z86 processor, the stack exists in the 256 bytes of register space, and it is crucial that the program does not overflow the stack, corrupting other data. Estimating the stack size used by a program is therefore of paramount interest to the correct operation of these systems. A tight upper bound is necessary to check if the program fits into the available memory, and to prevent precious system resources (e.g., registers) from being allocated unnecessarily.

* Palsberg, Ma, and Zhao were supported by the NSF ITR award 0112628. Henzinger, Chatterjee, and Majumdar were supported by the AFOSR grant F49620-00-1-0327, the DARPA grants F33615-C-98-3614 and F33615-00-C-1693, the MARCO grant 98-DT-660, and the NSF grants CCR-0208875 and CCR-0085949.

Stack size analysis is particularly challenging for *interrupt-driven software*. Interrupt-driven software is often used in embedded real-time applications that require fast response to external events. Such programs usually have a fixed number of external interrupt sources, and for each interrupt source, a handler that services the interrupt. When an external interrupt occurs, control is transferred automatically to the corresponding handler if interrupt processing is enabled. To maintain fast response, interrupts should be enabled most of the time, in particular, higher-priority interrupts are enabled in lower-priority handlers. Interrupt handling uses stack space: when a handler is called, a return address is placed on the stack, and if the handler itself gets interrupted, then another return address is placed on the stack, and so on. A programming error occurs when the interrupt handlers can interrupt each other indefinitely, leading to an unbounded stack. Moreover, since stack boundedness violations may occur only for particular interrupt sequences, these errors are difficult to replicate and debug, and standard testing is often inadequate. Therefore, algorithms that statically check for stack boundedness and automatically provide precise bounds on the maximum stack size will be important development tools for interrupt-driven systems.

In this paper, we provide algorithms for the following two problems (defined formally in Section 2.3) for a large class of interrupt-driven programs:

- **Stack boundedness problem.** Given an interrupt-driven program, the stack boundedness problem asks if the stack size is bounded by a finite constant. More precisely, given a program p , the stack boundedness problem returns “yes” if there exists a finite integer K such that on all executions of the program p , the stack size never grows beyond K , and “no” if no such K exists.
- **Exact maximum stack size problem.** Given an interrupt-driven program, the exact maximum stack size problem asks for the maximum possible stack size. More precisely, given a program p , the exact maximum stack size problem returns an integer K such that for all executions of the program p , the stack size never grows beyond K , and such that there is a possible schedule of interrupts and an execution of the program p such that the stack size becomes K ; the problem returns ∞ if there is an execution where the stack can grow unbounded.

We model interrupt-driven programs in the untyped *interrupt calculus* of Palsberg and Ma [3]. The interrupt calculus contains essential constructs for programming interrupt-driven systems. For example, we have found that the calculus can express the core aspects of seven commercial micro-controllers from Greenhill Manufacturing Ltd. A program in the calculus consists of a main part and some interrupt handlers. In the spirit of such processors as the Intel MCS-51 family (8051, etc.), Motorola Dragonball (68000 family), and Zilog Z86, the interrupt calculus supports an interrupt mask register (*imr*). An *imr* value consists of a master bit and one bit for each interrupt source. For example, the Motorola Dragonball processor can handle 22 interrupt sources. An interrupt handler is enabled, if *both* the master bit and the bit for that interrupt handler is set. When an interrupt handler is called, a return address is stored on the stack, and the

master bit is automatically turned off. At the time of return, the master bit is turned back on (however, the handler can turn the master bit on at any point). A program execution has access to:

- the interrupt mask register, which can be updated during computation,
- a stack for storing return addresses, and
- a memory of integer variables; output is done via memory-mapped I/O.

Each element on the stack is a return address. When we measure the size of the stack, we simply count the number of elements on the stack. Our analysis is approximate: when doing the analysis, we ignore the memory of integer variables and the program statements that manipulate this memory. In particular, we assume that both branches of a conditional depending on the memory state can be taken. Of course, all the problems analyzed in this paper become undecidable if integer variables are considered in the analysis.

We consider three versions of Palsberg and Ma’s interrupt calculus, here presented in increasing order of generality:

- **Monotonic programs.** These are interrupt calculus programs that satisfy the following monotonicity restriction: when a handler is called with an imr value imr_b , then it returns with an imr value imr_r such that $imr_r \leq imr_b$, where \leq is the logical bitwise implication ordering. In other words, every interrupt that is enabled upon return of a handler must have been enabled when the handler was called (but could have possibly been disabled during the execution of the handler).
- **Monotonic enriched programs.** This calculus enriches Palsberg and Ma’s calculus with conditionals on the interrupt mask register. The monotonicity restriction from above is retained.
- **Enriched programs.** These are programs in the enriched calculus, without the monotonicity restriction.

We summarize our results in Table 1. We have determined the complexity of stack boundedness and exact maximum stack size both for monotonic programs and for monotonic programs enriched with tests. For general programs enriched with tests, we have a PSPACE lower bound and an EXPTIME upper bound for both problems; tightening this gap remains an open problem. While the complexities are high, our algorithms are *polynomial* (linear or cubic) in the size of the program, and exponential only in the number of interrupts. In other words, our algorithms are polynomial if the number of interrupts is fixed. Since most real systems have a fixed small number of interrupts (for example Motorola Dragonball processor handles 22 interrupt sources), and the size of programs is the limiting factor, we believe the algorithms should be tractable in practice. Experiments are needed to settle this.

We reduce the stack boundedness and exact stack size problems to state space exploration problems over certain graphs constructed from the interrupt-driven program. We then use the structure of the graph to provide algorithms for the two problems. Our first insight is that for monotonic programs, the maximum stack

Calculus	Problem	Complexity	Reference
Monotonic	Stack boundedness	NLOGSPACE-complete	Theorem 1
	Exact maximum stack size	PSPACE-complete	Theorems 2,4
Monotonic enriched	Stack boundedness	PSPACE-complete	Theorems 3,4
	Exact maximum stack size	PSPACE-complete	Theorems 2,4
Enriched	Stack boundedness	PSPACE-hard, EXPTIME	Theorems 3,6
	Exact maximum stack size	PSPACE-hard, EXPTIME	Theorems 2,6

Table 1. Complexity results

bounds are attained without any intermediate handler return. The polynomial-time algorithm for monotonic programs is reduced to searching for cycles in a graph; the polynomial-space algorithm for determining the exact maximum stack size of monotonic enriched programs is based on finding the longest path in a (possibly exponential) acyclic graph. Finally, we can reduce the stack boundedness problem and exact maximum stack size problem for enriched programs to finding context-free cycles and context-free longest paths in graphs. Our EXPTIME algorithm for enriched programs is based on a novel technique to find the longest context-free path in a DAG. Our lower bounds are obtained by reductions from reachability in a DAG (which is NLOGSPACE-complete), satisfiability of quantified boolean formulas (which is PSPACE-complete), and reachability for polynomial-space Turing Machines (which is PSPACE-complete). We also provide algorithms that determine, given an interrupt-driven program, whether it is monotonic. In the nonenriched case, monotonicity can be checked in polynomial time (NLOGSPACE); in the enriched case, in co-NP. In Section 2, we recall the interrupt calculus of Palsberg and Ma [3]. In Section 3, we consider monotonic programs, in Section 4, we consider monotonic enriched programs, and in Section 5, we consider enriched programs without the monotonicity restriction.

Related work. Palsberg and Ma [3] present a type system and a type checking algorithm for the interrupt calculus that guarantees stack boundedness and certifies that the stack size is within a given bound. Each type contains information about the stack size and serves as documentation of the program. However, this requires extensive annotations from the programmer (especially since the types can be exponential in the size of the program), and the type information is absent in legacy programs. Our work can be seen as related to *type inference* for the interrupt calculus. In particular, we check stack properties of programs without annotations. Thus, our algorithms work for existing, untyped programs. Moreover, from our algorithms, we can infer the types of [3].

Brylow, Damgaard, and Palsberg [1] do stack size analysis by running a context-free reachability algorithm for model checking. They use, essentially, the same abstraction that our EXPTIME algorithm uses for enriched programs. Our paper gives more algorithmic details and clarifies that the complexity is exponential in the number of handlers.

Hughes, Pareto, and Sabry [2, 5] use *sized types* to reason about liveness, termination, and space boundedness of reactive systems. However, they require types with explicit space information, and do not address interrupt handling.

Wan, Taha, and Hudak [7] present event-driven Functional Reactive Programming (FRP), which is designed such that the time and space behavior of a program are necessarily bounded. However, the event-driven FRP programs are written in continuation-style, and therefore do not need a stack. Hence stack boundedness is not among the resource issues considered by Wan et al.

2 The Interrupt Calculus

2.1 Syntax

We recall the (abstract) syntax of the interrupt calculus of [3]. We use x to range over a set of program variables, we use imr to range over bit strings, and we use c to range over integer constants.

$$\begin{array}{ll}
\text{(program)} & p ::= (m, \bar{h}) \\
\text{(main)} & m ::= \text{loop } s \mid s ; m \\
\text{(handler)} & h ::= \text{iret} \mid s ; h \\
\text{(statement)} & s ::= x = e \mid \text{imr} = \text{imr} \wedge \text{imr} \mid \text{imr} = \text{imr} \vee \text{imr} \mid \\
& \quad \text{if0 } (x) s_1 \text{ else } s_2 \mid s_1 ; s_2 \mid \text{skip} \\
\text{(expression)} & e ::= c \mid x \mid x + c \mid x_1 + x_2
\end{array}$$

The pair $p = (m, \bar{h})$ is an *interrupt program* with main program m and interrupt handlers \bar{h} . The over-bar notation \bar{h} denotes a sequence $h_1 \dots h_n$ of handlers. We use the notation $\bar{h}(i) = h_i$. We use a to range over m and h .

2.2 Semantics

We use R to denote a *store*, that is, a partial function mapping program variables to integers. We use σ to denote a *stack* generated by the grammar $\sigma ::= \text{nil} \mid a :: \sigma$. We define the size of a stack as $|\text{nil}| = 0$ and $|a :: \sigma| = 1 + |\sigma|$.

We represent the imr as a bit sequence $imr = b_0 b_1 \dots b_n$, where $b_i \in \{0, 1\}$. The 0th bit b_0 is the master bit, and for $i > 0$, the i th bit b_i is the bit for interrupts from source i , which are handled by handler i . Notice that the master bit is the most significant bit, the bit for handler 1 is the second-most significant bit, and so on. This layout is different from some processors, but it simplifies the notation used later. For example, the imr value $101b$ means that the master bit is set, the bit for handler 1 is not set, and the bit for handler 2 is set. We use the notation $imr(i)$ for bit b_i . The predicate *enabled* is defined as

$$\text{enabled}(imr, i) = (imr(0) = 1) \wedge (imr(i) = 1), \quad i \in 1..n.$$

We use 0 to denote the imr value where all bits are 0. We use t_i to denote the imr value where all bits are 0's except that the i th bit is set to 1. We will use \wedge

to denote bitwise logical conjunction, \vee to denote bitwise logical disjunction, \leq to denote bitwise logical implication, and \neg to denote bitwise logical negation. Notice that $enabled(t_0 \vee t_i, j)$ is true if $i = j$, and false otherwise. The imr values, ordered by \leq , form a lattice with bottom element 0.

A *program state* is a tuple $\langle \bar{h}, R, imr, \sigma, a \rangle$ consisting of interrupt handlers \bar{h} , a store R , an interrupt mask register imr , a stack σ of return addresses, and a program counter a . We refer to a as *the current statement*; it models the instruction pointer of a CPU. We use P to range over program states. If $P = \langle \bar{h}, R, imr, \sigma, a \rangle$, then we use the notation $P.stk = \sigma$. For $p = (m, \bar{h})$, the initial program state for executing p is $P_p = \langle \bar{h}, \lambda x.0, 0, nil, m \rangle$, where the function $\lambda x.0$ is defined on the variables that are used in the program p .

A small-step operational semantics for the language is given by the reflexive, transitive closure of the relation \rightarrow on program states:

$$\langle \bar{h}, R, imr, \sigma, a \rangle \rightarrow \langle \bar{h}, R, imr \wedge \neg t_0, a :: \sigma, \bar{h}(i) \rangle \quad (1)$$

if $enabled(imr, i)$

$$\langle \bar{h}, R, imr, \sigma, iret \rangle \rightarrow \langle \bar{h}, R, imr \vee t_0, \sigma', a \rangle \quad \text{if } \sigma = a :: \sigma' \quad (2)$$

$$\langle \bar{h}, R, imr, \sigma, \text{loop } s \rangle \rightarrow \langle \bar{h}, R, imr, \sigma, s; \text{loop } s \rangle \quad (3)$$

$$\langle \bar{h}, R, imr, \sigma, x = e; a \rangle \rightarrow \langle \bar{h}, R\{x \mapsto eval_R(e)\}, imr, \sigma, a \rangle \quad (4)$$

$$\langle \bar{h}, R, imr, \sigma, imr = imr \wedge imr'; a \rangle \rightarrow \langle \bar{h}, R, imr \wedge imr', \sigma, a \rangle \quad (5)$$

$$\langle \bar{h}, R, imr, \sigma, imr = imr \vee imr'; a \rangle \rightarrow \langle \bar{h}, R, imr \vee imr', \sigma, a \rangle \quad (6)$$

$$\langle \bar{h}, R, imr, \sigma, (\text{if0 } (x) s_1 \text{ else } s_2); a \rangle \rightarrow \langle \bar{h}, R, imr, \sigma, s_1; a \rangle \quad \text{if } R(x) = 0 \quad (7)$$

$$\langle \bar{h}, R, imr, \sigma, (\text{if0 } (x) s_1 \text{ else } s_2); a \rangle \rightarrow \langle \bar{h}, R, imr, \sigma, s_2; a \rangle \quad \text{if } R(x) \neq 0 \quad (8)$$

$$\langle \bar{h}, R, imr, \sigma, \text{skip}; a \rangle \rightarrow \langle \bar{h}, R, imr, \sigma, a \rangle \quad (9)$$

where the function $eval_R(e)$ is defined as:

$$\begin{aligned} eval_R(c) &= c & eval_R(x + c) &= R(x) + c \\ eval_R(x) &= R(x) & eval_R(x_1 + x_2) &= R(x_1) + R(x_2). \end{aligned}$$

Rule (1) models that if an interrupt is enabled, then it may occur. The rule says that if $enabled(imr, i)$, then it is a possible transition to push the current statement on the stack, make $\bar{h}(i)$ the current statement, and turn off the master bit in the imr. Notice that we make no assumptions about the interrupt arrivals; any enabled interrupt can occur at any time, and conversely, no interrupt has to occur. Rule (2) models interrupt return. The rule says that to return from an interrupt, remove the top element of the stack, make the removed top element the current statement, and turn on the master bit. Rule (3) is an unfolding rule for loops, and Rules (4)–(9) are standard rules for statements. Let \rightarrow^* denote the reflexive transitive closure of \rightarrow .

A *program execution* is a sequence $P_p \rightarrow P_1 \rightarrow P_2 \rightarrow \dots \rightarrow P_k$ of program states. Consider a program execution γ of the form $P_p \rightarrow^* P_i \rightarrow P_{i+1} \rightarrow^* P_j \rightarrow P_{j+1}$ with $P_i = \langle \bar{h}, R, imr_b, \sigma, a \rangle$ and $P_j = \langle \bar{h}, R', imr', \sigma', a' \rangle$. The handler $\bar{h}(i)$ is called in γ with imr_b from state P_i and returns with imr_r from state P_j if

$$\begin{aligned} P_i \rightarrow P_{i+1} &= \langle \bar{h}, R, imr_b \wedge \neg t_0, a :: \sigma, \bar{h}(i) \rangle \quad \text{and } enabled(imr_b, i), \\ P_j \rightarrow P_{j+1} &= \langle \bar{h}, R', imr_r, \sigma, a \rangle \quad \text{and } \sigma' = a :: \sigma, \end{aligned}$$

```

imr = imr or 111b
loop { imr = imr or 111b }
handler 1 {
  imr = imr and 101b
  imr = imr or 100b
  iret
}

handler 2 {
  imr = imr and 110b
  imr = imr or 010b
  imr = imr or 100b
  imr = imr and 101b
  iret
}

```

Fig. 1. A program in the interrupt calculus

and $P_k.stk \neq \sigma$ for all $i < k \leq j$. We say that there is no handler call in γ between P_i and P_j if for all $i \leq k < j$, the transition $P_k \rightarrow P_{k+1}$ is not a transition of the form (1). Similarly, given an execution $P_p \rightarrow^* P_i \rightarrow^* P_j$, there is no handler return between P_i and P_j if for all $i \leq k < j$, the transition $P_k \rightarrow P_{k+1}$ is not a transition of the form (2).

2.3 Stack Size Analysis

We consider the following problems of stack size analysis.

- **Stack boundedness problem** Given an interrupt program p , the stack boundedness problem returns “yes” if there exists a finite integer K such that for all program states P' , if $P_p \rightarrow^* P'$, then $|P'.stk| \leq K$; and returns “no” if there is no such K .
- **Exact maximum stack size problem** For a program state P we define $maxStackSize(P)$ as the least $K \geq 0$ such that for all P' , if $P \rightarrow^* P'$, then $|P'.stk| \leq K$; and “infinite” in case no such K exists. The exact maximum stack size problem is given an interrupt program p and returns $maxStackSize(P_p)$.

Figure 1 shows an example of a program in the interrupt calculus. The bit sequences such as 111b are imr values. Notice that each of the two handlers can be called from different program points with different imr values. The bodies of the two handlers manipulate the imr, and both are at some point during the execution open to the possibility of being interrupted by the other handler. However, the maximum stack size is 3. This stack size happens if handler 1 is called first, then handler 2, and then handler 1 again, at which time there are three return addresses on the stack.

We shall analyze interrupt programs under the usual program analysis assumption that all paths in the program are executable. More precisely, our analysis assumes that each data assignment statement $x = e$ in the program has been replaced by skip, and each conditional `if0 (x) s1 else s2` has been replaced by `if0 (*) s1 else s2`, where $*$ denotes nondeterministic choice. While this is an overapproximation of the actual set of executable paths, we avoid trivial undecidability results for deciding if a program path is actually executable. Since the

data manipulation in interrupt programs is usually independent of the manipulation of the imr, this is a valid assumption.

3 Monotonic Interrupt Programs

We first define monotonic interrupt programs and then analyze the stack boundedness and exact maximum stack size problems for such programs. A handler h_i of program p is *monotonic* if for every execution γ of p , if h_i is called in γ with an imr value imr_b and returns with an imr value imr_r , then $imr_r \leq imr_b$. The program p is *monotonic* if all handlers $h_1 \dots h_n$ of p are monotonic. The handler h_i of p is *monotonic in isolation* if for every execution γ of p , if h_i is called in γ with an imr value imr_b from a state P_i and returns with an imr value imr_r from a state P_j such that there is no handler call between P_i and P_j , then $imr_r \leq imr_b$.

We first show that a program $p = (m, \bar{h})$ is monotonic iff every handler $h_i \in \bar{h}$ is monotonic in isolation. Moreover, a handler h_i is monotonic in isolation iff, whenever h_i is called with imr value $t_0 \vee t_i$ from state P_i and returns with imr_r from state P_j , with no handler calls between P_i and P_j , then $imr_r \leq t_0 \vee t_i$. These observations can be used to efficiently check if an interrupt program is monotonic: for each handler, we check that the return value imr_r of the imr when called with $t_0 \vee t_i$ satisfies $imr_r \leq t_0 \vee t_i$.

Proposition 1. *It can be checked in linear time (NLOGSPACE) if an interrupt program is monotonic.*

3.1 Stack Boundedness

We now analyze the complexity of stack boundedness of monotonic programs. Our main insight is that the maximum stack size is achieved without any intermediate handler returns. First observe that if handler h is enabled when the imr is imr_1 , then it is enabled for all imr $imr_2 \geq imr_1$. We argue the case where the maximum stack size is finite, the same argument can be formalized in case the maximum stack size is infinite. Fix an execution sequence that achieves the maximum stack size. Let h be the last handler that returned in this sequence (if there is no such h then we are done). Let the sequence of statements executed be $s_0, s_1, \dots, s_{i-1}, s_i, \dots, s_j, s_{j+1}, \dots$ where s_i was the starting statement of h and s_j the iret statement of h . Suppose h was called with imr_b and returned with imr_r such that $imr_r \leq imr_b$. Consider the execution sequence of statements $s_0, s_1, \dots, s_{i-1}, s_{j+1}, \dots$ with the execution of handler h being omitted. In the first execution sequence the imr value while executing statement s_{j+1} is imr_r and in the second sequence the imr value is imr_b . Since $imr_r \leq imr_b$ then repeating the same sequence of statements and same sequence of calls to handlers with h omitted gives the same stack size. Following a similar argument, we can show that all handlers that return intermediately can be omitted without changing the maximum stack size attained.

Lemma 1. *For a monotonic program p , let P_{\max} be a program state such that $P_p \rightarrow^* P_{\max}$ and for any state P' , if $P_p \rightarrow^* P'$ then $|P_{\max}.stk| \geq |P'.stk|$. Then there is a program state P'' such that $P_p \rightarrow^* P''$, $|P''.stk| = |P_{\max}.stk|$, and there is no handler return between P_p and P'' .*

We now give a polynomial-time algorithm for the stack boundedness problem for monotonic programs. The algorithm reduces the stack boundedness question to the presence of cycles in the *enabled graph* of a program. Let $h_1 \dots h_n$ be the n handlers of the program. Given the code of the handlers we build the enabled graph $G = \langle V, E \rangle$ as follows:

- There is a node for each handler, i.e., $V = \{h_1, h_2, \dots, h_n\}$.
- Let the instructions of h_i be $C_i = i_1, i_2, \dots, i_m$. There is an edge between (h_i, h_j) if any of the following condition holds
 1. There is l, k such that $l \leq k$, the instruction at i_l is $\text{imr} = \text{imr} \vee \text{imr}$ with $t_0 \leq \text{imr}$, the instruction at i_k is $\text{imr} = \text{imr} \vee \text{imr}$ with $t_j \leq \text{imr}$ and for all statement i_m between i_l and i_k if i_m is $\text{imr} = \text{imr} \wedge \text{imr}$ then $t_0 \leq \text{imr}$.
 2. There is l, k such that $l \leq k$, the instruction at i_l is $\text{imr} = \text{imr} \vee \text{imr}$ with $t_j \leq \text{imr}$, the instruction at i_k is $\text{imr} = \text{imr} \vee \text{imr}$ with $t_0 \leq \text{imr}$ and for all statement i_m between i_l and i_k if i_m is $\text{imr} = \text{imr} \wedge \text{imr}$ then $t_j \leq \text{imr}$.
 3. There is l such that the instruction at i_l is $\text{imr} = \text{imr} \vee \text{imr}$ with $t_0 \leq \text{imr}$ and for all statement i_m between i_1 and i_l if i_m is $\text{imr} = \text{imr} \wedge \text{imr}$ then $t_i \leq \text{imr}$.

Since we do not model the program variables, we can analyze the code of h_i and detect all outgoing edges (h_i, h_j) in time linear in the length of h_i . We only need to check that there is an \vee statement with an imr constant with j th bit 1 and then the master bit is turned on with no intermediate disabling of the j th bit or vice versa. Hence the enabled graph for program p can be constructed in time $n^2 \times |p|$ (where $|p|$ denotes the length of p).

Let G_p be the enabled graph for a monotonic interrupt program p . If G_p has a cycle, then the stack is unbounded, that is, for all positive integer K there is a program state P' such that $P_p \rightarrow^* P'$ and $|P'.stk| > K$. Since cycles in the enabled graph can be found in NLOGSPACE, the stack boundedness problem for monotonic programs is in NLOGSPACE; note that the enabled graph of a program can be generated on the fly in logarithmic space. Hardness for NLOGSPACE follows from the hardness of DAG reachability.

Theorem 1. *Stack boundedness for monotonic interrupt programs can be checked in time linear in the size of the program and quadratic in the number of handlers. The complexity of stack boundedness for monotonic interrupt programs is NLOGSPACE-complete.*

In case the stack is bounded, we can get a simple upper bound on the stack size as follows. Let G_p be the enabled graph for a monotonic interrupt program p . If G_p is a DAG, and the node h_i of G_p has order k in topological sorting order,

then we can prove by induction that the corresponding handler h_i of p can occur at most $2^{(k-1)}$ times in the stack. Hence for any program state P' such that $P_p \rightarrow^* P'$, we have $|P'.stk| \leq 2^n - 1$. In fact, this bound is tight: there is a program with n handlers that achieves a maximum stack size of $2^n - 1$. We show that starting with an imr value of all 1's one can achieve the maximum stack length of $2^n - 1$ and the stack remain bounded. We give an inductive strategy to achieve this. With one handler which does not turn itself on we can have a stack length 1 starting with imr value 11. By induction hypothesis, using $n-1$ handlers starting with imr value all 1's we can achieve a stack length of $2^{n-1} - 1$. Now we add the n th handler and modify the previous $n-1$ handlers such that they do not change the bit for the n th handler. The n -th handler turns on every bit except itself, and then turns on the master bit. The following sequence achieves a stack size of $2^n - 1$. First, the first $n-1$ handlers achieve a stack size of $2^{n-1} - 1$ using the inductive strategy. After this, the n th handler is called. It enables the $n-1$ handlers but disables itself. Hence the sequence of stack of $2^{n-1} - 1$ can be repeated twice and the n the handler can occur once in the stack in between. The total length of stack is thus $1 + (2^{n-1} - 1) + (2^{n-1} - 1) = 2^n - 1$. Since none of the other handlers can turn the n th handler on, the stack size is in fact bounded. However, the exact maximum stack size problem can be solved only in PSPACE. We defer the algorithm to the next section, where we solve the problem for a more general class of programs.

3.2 Maximum Stack Size

We now prove that the exact maximum stack size problem is PSPACE-hard. We define a subclass of monotonic interrupt calculus which we call simple interrupt calculus and show the exact maximum stack size problem is already PSPACE-hard for this class.

For imr' , imr'' where $imr'(0) = 0$ and $imr''(0) = 0$, define $\mathcal{H}(imr'; imr'')$ to be the interrupt handler

$$\begin{aligned} imr &= imr \wedge \neg imr'; \\ imr &= imr \vee (t_0 \vee imr''); \\ imr &= imr \wedge \neg(t_0 \vee imr''); \\ &iret. \end{aligned}$$

A program $p = (m, \bar{h})$ is *simple* if the main program h has the form

$$imr = imr \vee (imr_S \vee t_0); \quad \text{loop skip}$$

where $imr_S(0) = 0$, and every interrupt handler in \bar{h} has the form $\mathcal{H}(imr'; imr'')$. Intuitively, a handler of a simple interrupt program first disables some handlers, then enables other handlers and enables interrupt handling. This opens the door to the handler being interrupted by other handlers. After that, it disables interrupt handling, and makes sure that the handlers that are enabled on exit are a subset of those that were enabled on entry to the handler.

For a handler $h(i)$ of the form $\mathcal{H}(imr'; imr'')$, we define function $f_i(imr) = imr \wedge (\neg imr') \vee imr''$. Given a simple interrupt program p , we define a directed graph $G(p) = (V, E)$ where $V = \{imr \mid imr(0) = 0\}$ and $E = \{(imr, f_i(imr)) \mid t_i \leq imr\}$.

The edge $(imr, f_i(imr))$ in $G(p)$ represents the call to the interrupt handler $h(i)$ when imr value is imr . We define imr_S as the start node of $G(p)$ and we define $\mathcal{M}(imr)$ as the longest path in $G(p)$ from node imr . The notation $\mathcal{M}(imr)$ is ambiguous because it leaves the graph unspecified; however, in all cases below, the graph in question can be inferred from the context.

Lemma 2. *For a simple program p , we have $maxStackSize(P_p) = |\mathcal{M}(imr_S)|$.*

We now show PSPACE-hardness for simple interrupt calculus. Our proof is based on a polynomial-time reduction from the *quantified boolean satisfiability* (QSAT) problem [4]. The proof is technical, and given in the full paper.

We illustrate the reduction by a small example. Suppose we are given a QSAT instance $S = \exists x_2 \forall x_1 \phi$ with $\phi = (l_{11} \vee l_{12}) \wedge (l_{21} \vee l_{22}) = (x_2 \vee \neg x_1) \wedge (x_2 \vee x_1)$. We construct a simple interrupt program $p = (m, \bar{h})$ with an imr register, where $\bar{h} = \{h(x_i), h(\bar{x}_i), h(w_i), h(\bar{w}_i), h(l_{ij}) \mid i, j = 1, 2\}$ are 12 handlers. The imr contains 13 bits: a master bit, and each remaining bit 1-1 maps to each handler in \bar{h} . Let $\mathcal{D} = \{x_i, \bar{x}_i, w_i, \bar{w}_i, l_{ij} \mid i, j = 1, 2\}$. We use t_x , where $x \in \mathcal{D}$, to denote the imr value where all bits are 0's except the bit corresponding to handler $h(x)$ is set to 1. The initial imr value imr_S is set to $imr_S = t_{x_2} \vee t_{\bar{x}_2}$.

We now construct \bar{h} . Let $E(h(x))$, $x \in \mathcal{D}$, be the set of handlers that $h(x)$ enables. This *enable* relation between the handlers of our example is illustrated in Figure 2, where there is an edge from $h(x_i)$ to $h(x_j)$ iff $h(x_i)$ enables $h(x_j)$. Let $D(h(x))$, $x \in \mathcal{D}$, be the set of handlers that $h(x)$ disables. Let $L = \{h(l_{ij}) \mid i, j = 1, 2\}$. The $D(h(x)), x \in \mathcal{D}$, are defined as follows:

$$D(h(x_2)) = D(h(\bar{x}_2)) = \{h(x_2), h(\bar{x}_2)\}, \quad (10)$$

$$D(h(x_1)) = \{h(x_1)\}, \quad D(h(\bar{x}_1)) = \{h(\bar{x}_1)\}, \quad (11)$$

$$D(h(w_2)) = D(h(\bar{w}_2)) = \{h(x_1), h(\bar{x}_1)\} \cup \{h(w_i), h(\bar{w}_i) \mid i = 1, 2\} \cup L, \quad (12)$$

$$D(h(w_1)) = D(h(\bar{w}_1)) = \{h(w_1), h(\bar{w}_1)\} \cup L, \quad (13)$$

$$D(h(l_{ij})) = \{h(l_{i1}), h(l_{i2})\} \cup \{h(w_k) \mid \text{if } l_{ij} = \neg x_k\} \cup \{h(\bar{w}_k) \mid \text{if } l_{ij} = x_k\} \quad (14)$$

If $h(x) = \mathcal{H}(imr'; imr'')$, then $imr' = \bigvee_{h(y) \in E(h(x))} t_y$ and $imr'' = \bigvee_{h(z) \in D(h(x))} t_z$, where $x, y, z \in \mathcal{D}$.

We claim that the QSAT instance S is satisfiable iff $|\mathcal{M}(imr_S)| = 10$, where $imr_S = t_{x_2} \vee t_{\bar{x}_2}$. We sketch the proof as follows.

Let $imr_L = \bigvee_{h(l) \in L} t_l$, where $l \in \mathcal{D}$. From (14) and Figure 2, it can be shown that $|\mathcal{M}(imr_L)| = 2$. From Figure 2, we have $E(h(x_1)) = \{h(w_1)\} \cup L$; and together with (13), and (14), it can be shown that

$$|\mathcal{M}(t_{x_1})| = 1 + |\mathcal{M}(t_{w_1} \vee imr_L)| \leq 2 + |\mathcal{M}(imr_L)| = 4$$

and the equality holds iff $\exists j_1, j_2 \in 1, 2$, such that $l_{1j_1}, l_{2j_2} \neq \neg x_1$, because otherwise handler $h(w_1)$ would be surely disabled. Similarly, it can be shown that

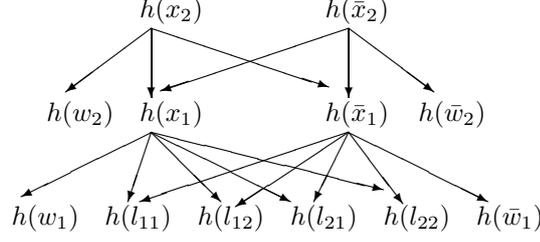


Fig. 2. Enable relation of interrupt handlers

$|\mathcal{M}(t_{\bar{x}_1})| \leq 4$, and that

$$|\mathcal{M}(t_{x_1} \vee t_{\bar{x}_1})| \leq |\mathcal{M}(t_{x_1})| + |\mathcal{M}(t_{\bar{x}_1})| \leq 8,$$

where the equality holds iff $\exists j_1, j_2$, such that $l_{1j_1}, l_{2j_2} \neq \neg x_1$ and $\exists j'_1, j'_2$, such that $l_{1j'_1}, l_{2j'_2} \neq x_1$. From Figure 2, we have $E(h(x_2)) = \{h(w_2), h(x_1), h(\bar{x}_1)\}$. Thus,

$$|\mathcal{M}(t_{x_2})| = 1 + |\mathcal{M}(t_{w_2} \vee t_{x_1} \vee t_{\bar{x}_1})| \leq 2 + |\mathcal{M}(t_{x_1} \vee t_{\bar{x}_1})| = 10,$$

and it can be shown from (12) and (14), that the equality holds iff $\exists j_1, j_2$ such that $l_{ij_1}, l_{ij_2} \neq \neg x_2, \neg x_1$ and $\exists j'_1, j'_2$ such that $l_{ij'_1}, l_{ij'_2} \neq \neg x_2, x_1$, which implies that both $x_2 = \text{true}, x_1 = \text{true}$ and $x_2 = \text{true}, x_1 = \text{false}$ are satisfiable truth assignments to ϕ . Similarly, it can be shown that $|\mathcal{M}(t_{\bar{x}_2})| = 10$ iff both $x_2 = \text{false}, x_1 = \text{true}$ and $x_2 = \text{false}, x_1 = \text{false}$ are satisfiable truth assignments to ϕ .

From (10), we have $|\mathcal{M}(t_{x_2} \vee t_{\bar{x}_2})| = \max(|\mathcal{M}(t_{x_2})|, |\mathcal{M}(t_{\bar{x}_2})|)$. Therefore, $|\mathcal{M}(imr_S)| = 10$ iff there exists x_2 such that for all x_1 , ϕ is satisfiable, or equivalently iff S is satisfiable. For our example, S is satisfiable since $\exists x_2 = \text{true}$ such that $\forall x_1$, ϕ is satisfiable. Correspondingly, $|\mathcal{M}(imr_S)| = |\mathcal{M}(x_2)| = 10$.

Theorem 2. *The exact maximum stack size problem for monotonic interrupt programs is PSPACE-hard.*

4 Monotonic Enriched Interrupt Programs

We now introduce an enriched version of the interrupt calculus, where we allow conditionals on the interrupt mask register. The conditional can test if some bit of the imr is on, and then take the bitwise or of the imr with a constant bit sequence; or it can test if some bit of the imr is off, and then take the bitwise and of the imr with a constant. The syntax for *enriched interrupt programs* is given by the syntax from Section 2 together with the following clauses:

(statement) $s ::= \dots \mid \text{if}(\text{bit } i \text{ on}) \text{ imr} = \text{imr} \vee \text{imr}' \mid \text{if}(\text{bit } i \text{ off}) \text{ imr} = \text{imr} \wedge \text{imr}'$

The small-step operational semantics is given below:

$$\begin{aligned} \langle \bar{h}, R, imr, \sigma, \text{if}(\text{bit } i \text{ on}) \text{ imr} = \text{imr} \vee \text{imr}'; a \rangle &\rightarrow \langle \bar{h}, R, imr \vee \text{imr}', \sigma, a \rangle && \text{if } imr(i) = 1 \\ \langle \bar{h}, R, imr, \sigma, \text{if}(\text{bit } i \text{ on}) \text{ imr} = \text{imr} \vee \text{imr}'; a \rangle &\rightarrow \langle \bar{h}, R, imr, \sigma, a \rangle && \text{if } imr(i) = 0 \\ \langle \bar{h}, R, imr, \sigma, \text{if}(\text{bit } i \text{ off}) \text{ imr} = \text{imr} \wedge \text{imr}'; a \rangle &\rightarrow \langle \bar{h}, R, imr \wedge \text{imr}', \sigma, a \rangle && \text{if } imr(i) = 0 \\ \langle \bar{h}, R, imr, \sigma, \text{if}(\text{bit } i \text{ off}) \text{ imr} = \text{imr} \wedge \text{imr}'; a \rangle &\rightarrow \langle \bar{h}, R, imr, \sigma, a \rangle && \text{if } imr(i) = 1 \end{aligned}$$

Unlike the conditional statement `if0 (x) s1 else s2` on data that has been over-approximated, our analysis will be path sensitive in the `imr`-conditional.

Proposition 2. *Monotonicity of enriched interrupt programs can be checked in time exponential in the number of handlers (in co-NP).*

For monotonic enriched interrupt programs, both the stack boundedness problem and the exact maximum stack size problem are PSPACE-complete. To show this, we first show that the stack boundedness problem is PSPACE-hard by a generic reduction from polynomial-space Turing machines. We fix a PSPACE-complete Turing machine M . Given input x , we construct in polynomial time a program p such that M accepts x iff p has an unbounded stack. We have two handlers for each tape cell (one representing zero, and the other representing one), and a handler for each triple (i, q, b) of head position i , control state q , and bit b . The handlers encode the working of the Turing machine in a standard way. The main program sets the bits corresponding to the initial state of the Turing machine, with x written on the tape. Finally, we have an extra handler that enables itself (and so can cause an unbounded stack) which is set only when the machine reaches an accepting state.

Theorem 3. *The stack boundedness problem for monotonic enriched interrupt programs is PSPACE-hard.*

We now give a PSPACE algorithm to check the exact maximum stack size. Since we restrict our programs to be monotonic it follows from Lemma 1 that the maximum length of the stack can be achieved with no handler returning in between. Given a program p with m statements and n handlers, we label the statements as pc_1, \dots, pc_m . Let PC denote the set of all statements, i.e., $PC = \{pc_1, \dots, pc_m\}$. Consider the graph G_p where there is a node v for every statement with all possible `imr` values (i.e., $v = \langle pc, imr \rangle$ for some value among PC and some `imr` value). Let $v = \langle pc, imr \rangle$ and $v' = \langle pc', imr' \rangle$ be two nodes in the graph. There is an edge between v, v' in G if any of the following two conditions hold:

- on executing the statement at pc with `imr` value imr the control goes to pc' and the value of `imr` is imr' . The weight of this edge is 0.
- pc' is a starting address of a handler h_i and $enabled(imr, i)$ and $imr' = imr \wedge \neg t_0$. The weight of this edge is 1.

We also have a special node in the graph called *target* and add edges to *target* of weight 0 from all those nodes which correspond to a $pc \in PC$ which is a *iret* statement. This graph is exponential in the size of the input as there are $O(|PC| \times 2^n)$ nodes in the graph. The starting node of the graph is the node with pc_1 and $imr = 0$. If there is a node in the graph which is the starting address of a handler h and which is reachable from the start node and also self-reachable then the stack length would be infinite. This is because the sequence of calls from the starting statement to the handler h is first executed and then the cycle of handler calls is repeated infinitely many times. As the handler h is

in stack when it is called again the stack would grow infinite. Since there is a sequence of interrupts which achieves the maximum stack length without any handler returning in between (follows from Lemma 1) if there is no cycle in G_p we need to find the longest path in the DAG G_p .

Theorem 4. *The exact maximum stack size for monotonic enriched interrupt programs can be found in time linear in the size of the program and exponential in the number of handlers. The complexity of exact maximum stack size for monotonic enriched interrupt programs is PSPACE.*

In polynomial space one can generate in lexicographic order all the nodes that have a pc value of the starting statement of a handler. If such a node is reachable from the start node, and also self-reachable, then the stack size is infinite. Since the graph is exponential, this can be checked in PSPACE. If no node has such a cycle, we find the longest path from the start node to the target. Again, since longest path in a DAG is in NLOGSPACE, this can be achieved in PSPACE. It follows that both the stack boundedness and exact maximum stack size problems for monotonic enriched interrupt programs are PSPACE-complete.

5 Nonmonotonic Enriched Interrupt Programs

In this section we consider interrupt programs with tests, but do not restrict handlers to be monotonic. We give an EXPTIME algorithm to check stack boundedness and find the exact maximum stack size for this class of programs. The algorithm involves computing longest context-free paths in context-free DAGs, a technique that may be of independent interest.

5.1 Longest Paths in Acyclic Context-free Graphs

We define a context-free graph as in [6]. Let Σ be a collection of k opening and closing parentheses, i.e., $\Sigma = \{(^1, (^2, \dots, (^k,)^1),)^2, \dots,)^k\}$. A *context-free graph* is a tuple $G = (V, E, \Sigma)$ where V is a set of nodes and $E \subseteq (V \times V \times (\Sigma \cup \{\tau\}))$ is a set of labeled edges (and τ is a special symbol not in Σ). We associate with each edge a *weight function* $wt : E \rightarrow \{0, +1, -1\}$ defined as follows:

- $wt(e) = 0$ if e is of the form (v, v', τ) ,
- $wt(e) = -1$ if e is of the form $(v, v',)^i$ for some i ,
- $wt(e) = 1$ if e is of the form $(v, v', (^i)$ for some i .

Let $\Sigma = \{(^1, (^2, \dots, (^k,)^1),)^2, \dots,)^k\}$ be an alphabet of matched parentheses. Let \mathcal{L} be the language generated by the context-free grammar

$$\begin{aligned} M &\rightarrow M(^i S && \text{for } 1 \leq i \leq k \\ S &\rightarrow \epsilon \mid (^i S)^i S && \text{for } 1 \leq i \leq k \end{aligned}$$

from the starting symbol M . Thus \mathcal{L} defines words of matched parentheses with possibly some opening parentheses mismatched. A *context-free path* π in

a context-free graph G is a sequence of vertices v_1, v_2, \dots, v_k such that for all $i = 1 \dots k - 1$, there is an edge between v_i and v_{i+1} , i.e., there is a letter $\sigma \in \Sigma \cup \{\tau\}$ such that $(v_i, v_{i+1}, \sigma) \in E$ and the projection of the labels along the edges of the path on to Σ is a word in \mathcal{L} . Given a context-free path π with edges e_1, e_2, \dots, e_k the *cost* of the path $Cost(\pi)$ is defined as $\sum_i wt(e_i)$. Note that $Cost(\pi) \geq 0$ for any context-free path π . A context-free graph G is a *context-free DAG* iff there is no cycle C of G such that $\sum_{e \in C} wt(e) > 0$. Given a context-free DAG $G = (V, E, \Sigma)$ we define an ordering $order : V \rightarrow \mathbb{N}$ of the vertices satisfying the following condition: if there is a path π in G from vertex v_i to v_j and $Cost(\pi) > 0$ then $order(v_j) < order(v_i)$. This ordering is well defined for context-free DAGs. Let G be a context-free DAG G , and let $V = \{v_1, v_2, \dots, v_n\}$ be the ordering of the vertex set consistent with $order$ (i.e., $order(v_i) = i$). We give a polynomial-time procedure to find the longest context-free path from any node v_i to v_1 in G .

Algorithm 1 Function LongestContextFreePath

Input: A context-free DAG G , a vertex v_1 of G

Output: For each vertex v of G , return the length of the longest context-free path from v to v_1 , and 0 if there is no context-free path from v to v_1

1. For each vertex $v_j \in V$: $val[v_j] = 0$
 2. Construct the transitive closure matrix T such that

$$T[i, j] = 1 \text{ iff there is a context-free path from } i \text{ to } j$$
 3. For $j = 1$ to n :
 - 3.1 For each immediate successor v_i of v_j such that the edge e_{v_j, v_i} from v_j to v_i satisfies $wt(e_{v_j, v_i}) \geq 0$:

$$val[v_j] = \max\{val[v_j], val[v_i] + wt(e_{v_j, v_i})\}$$
 - 3.2 For each vertex $v_i \in V$:
 - 3.2.1 if $(T[i, j])$ (v_j is context-free reachable from v_i)

$$val[v_i] = \max\{val[v_i], val[v_j]\}$$
-

The correctness proof of our algorithm uses a function Num from paths to \mathbb{N} . Given a path π we define $Num(\pi)$ as $\max\{order(v) \mid v \text{ occurs in } \pi\}$. Given a node v let $L_v = \{L_1, L_2, \dots, L_k\}$ be the set of longest paths from v to v_1 . Then we define $Num_{v_1}(v) = \min\{Num(L_i) \mid L_i \in L_v\}$. The correctness of the algorithm follows from the following set of observations. First, if there is a longest path L from a node v to v_1 such that L starts with an opening parenthesis (ⁱ that is not matched along the path L then $order(v) = Num_{v_1}(v)$. A node v in the DAG G satisfies the following conditions.

- If $Num_{v_1}(v) = order(v) = j$ then within the execution of Statement 3.1 of the j -th iteration of Loop 3 of function LongestContextFreePath, $val[v]$ is equal to the cost of a longest path from v to v_1 .
- If $order(v) < Num_{v_1}(v) = j$ then by the j -th iteration of Loop 3 of function LongestContextFreePath $val[v]$ is equal to the cost of a longest path from v to v_1 .

Finally, notice that at the end of function `LongestContextFreePath(G, v_1)`, for each vertex v , the value of $val[v]$ is equal to the longest context-free path to v_1 , and equal to zero if there is no context-free path to v_1 . In the Function `LongestContextFreePath` the statement 3.2.1 gets executed at most n^2 times since the loop 3 gets executed n times at most and the nested loop 3.2 also gets executed n times at most. The context-free transitive closure can be constructed in $O(n^3)$ time [8]. Hence the complexity of our algorithm is polynomial and it runs in time $O(n^2 + n^3) = O(n^3)$.

Theorem 5. *The longest context-free path of a context-free DAG can be found in time cubic in the size of the graph.*

To complete our description of the algorithm, we must check if a given context-free graph is a context-free DAG, and generate the topological ordering *order* for a context-free DAG. We give a polynomial-time procedure to check given a context-free graph whether it is a DAG. Given a context-free graph $G = (V, E, \Sigma)$ let its vertex set be $V = \{1, 2, \dots, n\}$. For every node $k \in V$ the graph G can be unrolled as a DAG for depth $|V|$ and it can be checked if there is a path π from k to k such that $Cost(\pi) > 0$. We give the construction of the DAG below.

Given the graph G and a node k we create a context-free DAG $G_k = (V_k, E_k, \Sigma)$ as follows:

1. $V_k = \{k_0\} \cup \{(i, j) \mid 1 \leq i \leq n-2, 1 \leq j \leq n\} \cup \{k_{n-1}\}$
2. $E_k = \{ \langle k_0, (1, j), * \rangle \mid \langle k, j, * \rangle \in E \} \cup \{ \langle (i, j), (i+1, j'), * \rangle \mid \langle j, j', * \rangle \in E \}$
 $\cup \{ \langle (n-2, j), k_{n-1}, * \rangle \mid \langle j, k, * \rangle \in E \} \cup \{ \langle k_0, (1, k) \rangle \} \cup \{ \langle (i, k), (i+1, k) \rangle \}$

where $*$ can represent a opening parenthesis, closing parenthesis or can be τ . Notice that the edges in the last line ensure that if there is a cycle of positive cost from k to itself with length $t < n$ then it is possible to go from k_0 to $(n-t-1, k)$ and then to reach k_{n-1} by a path of positive cost.

We can find the longest context-free path from k_0 to k_n in G_n (by the function `LongestContextFreePath`). If the length is positive, then there is a positive cycle in G from k to k . If for all nodes the length of the longest path in G_n is 0, then G is a context-free DAG and the longest context-free path can be computed in G . Given a context-free DAG G we can find the *order*(v) of the vertices in polynomial time. If a vertex v can reach v' and v' can reach v put them in the same group of vertices. Both the path from v to v' and v' to v must be cost 0 else there would be a cycle of positive cost. Hence the ordering of vertices within a group can be arbitrary. We can topologically order the graph induced by the groups and then assign an order to the vertices where vertices in the same group are ordered arbitrarily.

5.2 Stack Size Analysis

We present an algorithm to check for stack boundedness and exact maximum stack size. The idea is to perform context-free longest path analysis on the state

space of the program. Given a program p with m statements and n handlers, we label the statements as pc_1, pc_2, \dots, pc_m . Let $PC = \{pc_1, \dots, pc_m\}$ as before. We construct a context-free graph [6] $G_p = \langle V, E, \Sigma \rangle$, called the *state graph of p* , where $\Sigma = \{(1, (2, \dots, (m,)^1,)^2, \dots)^m\}$ as follows:

- $V = PC \times IMR$, where IMR is the set of all 2^n possible imr values.
- $E \subseteq (V \times V \times \Sigma) \cup (V \times V)$
 1. Handler call: $(v, v', (i) \in E$ iff $v = (pc_i, imr_1)$ and $v' = (pc_j, imr_2)$ and pc_j is the starting address of some handler h_j such that $enabled(imr_1, j)$ and $imr_2 = imr_1 \wedge \neg t_0$.
 2. Handler return: $(v, v', (i) \in E$ iff $v = (pc_i, imr_1)$ and $v' = (pc_j, imr_2)$ and pc_j is the ired statement of some handler and $imr_1 = imr_2 \vee t_0$.
 3. Statement execution: $(v, v') \in E$ iff $v = (pc_i, imr_1)$ and $v' = (pc_j, imr_2)$ and executing the statement at pc_i with imr value imr_1 the control goes to pc_j and the imr value is imr_2 .

The vertex $(pc_1, 0)$ is the starting vertex of G_p . Let G'_p be the induced subgraph of G_p which contains only nodes which are context-free reachable from the start node. If G'_p is not a context-free DAG then we report that stack is unbounded, else we create a new DAG G''_p adding a new vertex *target* and adding edges to *target* from all nodes of G'_p of weight 0 and find the value of the longest context-free path from the start vertex to *target* in the DAG G''_p .

Algorithm 2 Function StackSizeGeneral

Input: Enriched interrupt program p

Output: $maxStackSize(P_p)$

1. Build the state graph $G_p = \langle V, E, \Sigma \rangle$ from the program p
 2. Let $V' = \{v' \mid \text{there is a context-free path from the starting vertex to } v'\}$
 3. Let G'_p be the subgraph of G_p induced by the vertex set V'
 4. If G'_p is not a context-free DAG then return “infinite”
 5. Else create $G''_p = \langle V'', E'', \Sigma \rangle$ as follows :
 - 5.1 $V'' = V' \cup \{target\}$
 - 5.2 $E'' = E' \cup \{(v, target) \mid v \in V'\}$
 6. Return the value of the longest context-free path from the starting vertex to *target*
-

From the construction of the state graph, it follows that there is a context-free path from a vertex $v = (pc, imr)$ to $v' = (pc', imr')$ in the state graph G_p iff there exists stores R, R' and stacks σ, σ' such that $\langle \bar{h}, R, imr, \sigma, pc \rangle \rightarrow^* \langle \bar{h}, R', imr', \sigma', pc' \rangle$. Moreover, if G'_p is the reachable state graph then there exists K such that for all P' such that $P_p \rightarrow^* P'$ we have $|P'.stk| \leq K$ iff G'_p is a context-free DAG. To see this, first notice that if G'_p is not a context-free DAG then there is a cycle of positive cost. Traversing this cycle infinitely many times makes the stack grow unbounded. On the other hand, if the stack is unbounded then there is a program address that is visited infinitely many times with the

same imr value and the stack grows between the successive visits. Hence there is a cycle of positive cost in G'_p . These observations, together with Theorem 5 show that function `StackSizeGeneral` correctly computes the exact maximum stack size of an interrupt program p .

Theorem 6. *The exact maximum stack size of nonmonotonic enriched interrupt programs can be found in time cubic in the size of the program and exponential in the number of handlers.*

The number of vertices in G_p is $m \times 2^n$, for m program statements and n interrupt handlers. It follows from Theorem 5 and the earlier discussion that the steps 1, 2, 3, 4, 5, and 6 of `StackSizeGeneral` can be computed in time polynomial in G_p . Since G_p is exponential in the size of the input program p , we have an EXPTIME procedure for determining the exact maximum stack size of nonmonotonic enriched interrupt programs. Notice also that our syntax ensures that all statements that modify the imr are monotonic: if $imr_1 \leq imr_2$, and for $i = 1, 2$, we have $P(imr_i) \rightarrow P'(imr'_i)$ for any two program states $P(imr)$ and $P'(imr)$ parameterized by imr , then $imr'_1 \leq imr'_2$. In particular, we only allow bitwise or's on the imr if we test if a bit is set, and only allow bitwise and's on the imr if we test if a bit is unset. Indeed, we can extend the syntax of the enriched calculus to allow any imr operations, and the above algorithm still solves the exact maximum stack size problem, with no change in complexity.

We leave open whether the exact maximum stack size problem for nonmonotonic interrupts programs, in the nonenriched and enriched cases, is EXPTIME-hard or PSPACE-complete (PSPACE-hardness follows from Theorem 3).

References

1. D. Brylow, N. Damgaard, and J. Palsberg. Static checking of interrupt-driven software. In *ICSE: International Conference on Software Engineering*, pp. 47–56. ACM/IEEE, 2001.
2. J. Hughes, L. Pareto, and A. Sabry. Proving the correctness of reactive systems using sized types. In *POPL: Principles of Programming Languages*, pp. 410–423. ACM, 1996.
3. J. Palsberg and D. Ma. A typed interrupt calculus. In *FTRTFT: Formal Techniques in Real-Time and Fault-tolerant Systems*, LNCS 2469, pp. 291–310. Springer, 2002.
4. C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
5. L. Pareto. *Types for Crash Prevention*. PhD thesis, Chalmers University of Technology, 2000.
6. T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL: Principles of Programming Languages*, pp. 49–61. ACM, 1995.
7. Z. Wan, W. Taha, and P. Hudak. Event-driven FRP. In *PADL: Practical Aspects of Declarative Languages*, LNCS 2257, pp. 155–172. Springer, 2002.
8. M. Yannakakis. Graph-theoretic methods in database theory. In *PODS: Principles of Database Systems*, pp. 203–242. ACM, 1990.