

Quantitative Information Flow Analysis

Hauptseminar WS'2010/11

Sergey Grebenshchikov

June 30, 2011

Contents

1 Preliminaries	2
1.1 Secret information and information leaks	2
1.2 Programs and computations	3
2 Exact static analysis	6
2.1 Qualitative information flow	6
2.2 Refinement for leak discovery	6
2.3 Quantitative information flow	8
3 Approximation and randomization	9
3.1 Approximation of information leakage	10
3.2 Randomized approximation	11
4 Approximative dynamic analysis	14
4.1 Assembly programs and computations	14
4.2 Dynamic tainting and bit-tracking	14
4.3 Flow graphs	15

Introduction

The goal of *quantitative information flow* (QIF) analysis is to quantify the amount of sensitive information revealed (*leaked*) by a program. This problem is relevant for any program that handles sensitive or secret data - anonymization protocols, electronic voting, payment or banking systems, appointment organizers providing busy/free data and networked multiplayer games, among others. In a nutshell, we seek an answer to the question:

“How much information can an attacker gain about sensitive data from a program’s observable behavior?”

In the following, we will present several approaches to this question.

1 Preliminaries

1.1 Secret information and information leaks

In order to describe the secret information a program leaks, we need to define what we mean by “secret information”. One approach commonly taken (e.g in [4], [7]) is to consider program inputs consisting of a *high* (secret) and a *low* (public) part, and assume the low input and the entire source code (i.e. transition system) and output of a program to be public. In this presentation we assume the *entire* input as secret and the entire output and source code as public.

Information-theoretic entropy

In order to quantify the information leaked by a program, we need to introduce a quantitative measure of information. *Information-theoretic entropy* (also referred to as *Shannon entropy*) [6] is such a measure.

Let (Ω_X, \Pr) be a discrete probability space and Ω a measurable space. A *random variable* X is a map $\Omega_X \rightarrow \Omega$. If the probability distribution of X is $p_X : \Omega \rightarrow [0, 1]$, we write $X \sim p_X$. For $x \in \Omega$ we use the suggestive notation $\Pr[X = x]$ for $p_X(x)$. We denote the fact that a random variable Y has the same probability distribution as a random variable X by writing $Y \simeq X$. A random variable can be obtained by applying a function to another random variable. We write $Y = f(X)$ for $Y = f \circ X$. We define the *binary entropy* $H[X]$ of X as follows:

$$H[X] := - \sum_{x \in \Omega} \Pr[X = x] \log_2 \Pr[X = x] \quad (1)$$

If X is the identity function $\text{id}_{\Omega_X} : \Omega_X \rightarrow \Omega_X$ and is uniformly distributed on Ω_X , i.e. $\forall x \in \Omega_X : \Pr[X = x] = \frac{1}{|\Omega_X|}$, we can derive a simple expression for $H[X]$:

$$\begin{aligned} H[X] &= - \sum_{x \in \Omega_X} \Pr[X = x] \log_2 \Pr[X = x] \\ &= \frac{1}{|\Omega_X|} \sum_{x \in \Omega_X} \log_2 |\Omega_X| \\ &= \log_2 |\Omega_X| \end{aligned}$$

Entropy can be interpreted as describing *uncertainty* of the value of a random variable. In the special case of binary entropy, $H[X]$ is the average number of bits needed to determine the value of X . The above result for a uniform distribution matches our intuition: to identify an element out of a set of n equally probable elements, we need $\log_2 n$ bits.

Given two random variables $X : \Omega_X \rightarrow \Omega$ and $Y : \Omega_Y \rightarrow \Omega'$, we define the *conditional entropy* $H[X|Y = y]$ of X given the value y of Y as follows:

$$H[X|Y = y] := - \sum_{x \in \Omega} \Pr[X = x|Y = y] \log_2 \Pr[X = x|Y = y] \quad (2)$$

The average conditional entropy of X given Y is the average of $H[X|Y = y]$ over $y \in \Omega'$:

$$H[X|Y] := \sum_{y \in \Omega'} \Pr[Y = y] H[X|Y = y] \quad (3)$$

Applying the definition of conditional probability, we can derive the following equation for $H[U|V]$:

$$H[U|V] = H[U, V] H[V]$$

The average conditional entropy $H[X|Y]$ describes the uncertainty of the value of X after observing the value of Y .

Example 1 (Conditional entropy). Let $\Omega_X := \{1, 2, 3, 4\}$. Consider the two random variables X and Y :

$$\begin{aligned} X : \Omega_X &\rightarrow \Omega_X & Y : \Omega_X &\rightarrow \{0, 1\} \\ X = \text{id}_{\Omega_X} & & Y &= \begin{cases} 0 & X \leq 2 \\ 1 & \text{otherwise.} \end{cases} \end{aligned}$$

Assume a uniform probability $\Pr[x] = \frac{1}{|\Omega_X|} = \frac{1}{4}$ for all elementary events $x \in \Omega_X$. The entropy of X is $H[X] = \log_2 |\Omega_X| = 2$. Intuitively, observing the result of Y gives us a binary distinction about the value of X , so for $X|Y$ we expect an entropy reduced by one bit. For the elementary and conditional probabilities we obtain:

$$\begin{aligned} \Pr[Y = 0] &= \Pr[Y = 1] = \frac{1}{2} \\ \Pr[X = x|Y = 0] &= \begin{cases} \frac{1}{2} & x \leq 2 \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Plugging the probabilities into the formulas for conditional entropy, we get:

$$\begin{aligned} H[X|Y = 0] &= - \sum_{x=1}^4 \Pr[X = x, Y = 0] \log_2 \Pr[X = x, Y = 0] \\ &= - \sum_{x=1}^2 \frac{1}{4} \log_2 \frac{1}{4} = 1 \\ H[X|Y = 1] &= - \sum_{x=3}^4 \Pr[X = x, Y = 1] \log_2 \Pr[X = x, Y = 1] \\ &= - \sum_{x=3}^4 \frac{1}{4} \log_2 \frac{1}{4} = 1 \end{aligned}$$

For the average conditional entropy we obtain, as expected:

$$\begin{aligned} H[X|Y] &= \sum_{y \in \{0,1\}} \Pr[Y = y] H[X|Y = y] \\ &= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1 \end{aligned}$$

The remaining uncertainty of X is 1 bit.

1.2 Programs and computations

We consider programs \mathcal{P} of the form $\mathcal{P} = (\mathbb{S}, \mathbb{I}, \mathbb{F}, \mathcal{T})$, where

- \mathbb{S} is a set of program states
- $\mathbb{I} \subseteq \mathbb{S}$ is a set of *initial* states
- $\mathbb{F} \subseteq \mathbb{S}$ is a set of *final* states
- \mathcal{T} is a set of *transitions*, such that for each $\tau \in \mathcal{T}$ a *transition relation* $\rho_\tau \subseteq \mathbb{S} \times \mathbb{S}$ is given.

We define the *program transition relation* ρ of \mathcal{P} as the union of its transition relations:

$$\rho := \bigcup_{\tau \in \mathcal{T}} \rho_\tau$$

Using ρ , we can define the *input-output relation* ρ_{io} of \mathcal{P} as the transitive closure of ρ , restricted to initial and final states:

$$\rho_{\text{io}} := \rho^* \cap (I \times F)$$

We require that ρ_{io} is a total function $I \rightarrow F$, i.e.:

$$\forall s \in I : \exists_1 s' \in F : (s, s') \in \rho_{\text{io}}$$

A final state $s' \in F$ is *reachable* from an initial state $s \in I$ iff $(s, s') \in \rho_{\text{io}}$. We write F_{reach} for the set of reachable final states:

$$F_{\text{reach}} := \{s' \in F \mid \exists s \in I : s' \text{ reachable from } s\} \subseteq F$$

For a final state $s' \in F$, the *preimage* $\mathcal{P}^{-1}(s')$ is the set of all initial states s from which s' is reachable:

$$\mathcal{P}^{-1}(s') := \{s \in I \mid (s, s') \in \rho_{\text{io}}\}$$

An execution step $s \xrightarrow{\rho} s'$ of \mathcal{P} can be described using the *post operator*:

$$\begin{aligned} \text{Post} &: 2^{S \times S} \times S \rightarrow 2^S \\ \text{Post}_\rho(s \in S) &:= \{s' \in S \mid (s, s') \in \rho\} \end{aligned}$$

The definition of Post naturally extends to sets of states:

$$\begin{aligned} \text{Post} &: 2^{S \times S} \times (2^S \cup S) \rightarrow 2^S \\ \text{Post}_\rho(X \subseteq S) &:= \{s' \in S \mid \exists s \in S : (s, s') \in \rho\} \end{aligned}$$

The set of reachable final states F_{reach} can be computed by iterating Post :

$$\begin{aligned} F_{\text{reach}} &= F \cap \left(\bigcup_{k=0}^{\infty} \text{Post}_\rho^k(I) \right) \\ &= F \cap \text{lfp}(\text{Post}_\rho, I) \end{aligned}$$

Computing the above fixpoint explicitly can be prohibitively expensive. This is due to the large number of states of almost any practical program. We can, however, employ approximative methods to remedy this problem.

Over-approximation of F_{reach}

Abstract interpretation overcomes the limitations of explicit reachability computation by approximating the Post operator. This approximation is achieved by *abstracting* sets of states. For a set X of states, an *abstraction* of X is a superset of X . We compute abstractions using an *abstraction function* $\alpha : 2^S \rightarrow 2^S$, a monotone (w.r.t set inclusion) function that maps sets of states to abstractions. Instead of iterating Post_ρ , we iterate the composition Post_ρ^\sharp of Post_ρ with α :

$$\begin{aligned} \text{Post}_\rho^\sharp &:= \alpha \circ \text{Post}_\rho \\ F_{\text{reach}}^\sharp &:= F \cap \text{lfp}(\text{Post}_\rho^\sharp, \alpha(I)) \end{aligned}$$

By the monotonicity of α , the resulting set F_{reach}^\sharp is an over-approximation of F_{reach} , i.e.:

$$F_{\text{reach}} \subseteq F_{\text{reach}}^\sharp$$

The abstraction function α can be chosen such that the fixpoint can be computed in a finite number of steps.

Under-approximation of F_{reach}

Abstract interpretation does not yet provide practical approaches to under-approximation. Instead, we symbolically execute a program for a selected set of initial states $\{s_1, \dots, s_n\} \subseteq I$. We assume a symbolic execution of \mathcal{P} on a initial state s_i yields a tuple $(\pi_i, s'_i) \in \mathcal{T}^+ \times F$, where π_i is a path $s_i \rightarrow s'_i$ and s'_i is the final state reached. We can derive an under-approximation F_{reach}^b of F_{reach} by collecting the final states encountered along the paths:

$$F_{\text{reach}}^b := \bigcup_{i=1}^n \{s' \mid \exists s \in S : (s, s') \in (\rho_{\pi_i} \cap (I \times F))\}$$

We know each final state in F_{reach}^b to be reachable in \mathcal{P} . We obtain the inclusion:

$$F_{\text{reach}}^b \subseteq F_{\text{reach}}$$

We will use the methods from this section for the approximative QIF algorithms in section 3.

Example 2 (Programs and computations). Consider the program given in Fig. 2.

The set of states is given by the set of valuations of program variables and the

```

0: byte f(byte x) { // 0 <= x <= 255;
1:     byte ret=0;
2:     if(x <= 127)
3:         ret=x;
4:     else
5:         ret=0;
6:     return ret;
}
```


Figure 1: Program for Example 2

program counter:

$$S = \{\langle \text{ret} = u, x = v, \text{pc} = w \rangle \mid u, v \in \{0, \dots, 255\}, w \in \{0, \dots, 6\}\}$$

The set I of initial states is given by

$$I = \{s \in S \mid s.\text{pc} = 0 \wedge s.\text{ret} = 0\}$$



$$\begin{aligned}
\rho_{\tau_0} &\equiv \text{pc} = 0 \wedge 0 \leq x \leq 255 \wedge x' = x \wedge \text{ret}' = 0 \wedge \text{pc}' = 1 \\
\rho_{\tau_1} &\equiv \text{pc} = 1 \wedge x \leq 127 \wedge x' = x \wedge \text{ret}' = x \wedge \text{pc}' = 3 \\
\rho_{\tau_2} &\equiv \text{pc} = 1 \wedge 128 \leq x \wedge x' = x \wedge \text{ret}' = 0 \wedge \text{pc}' = 5 \\
\rho_{\tau_3} &\equiv \text{pc} = 3 \wedge x' = x \wedge \text{ret}' = \text{ret} \wedge \text{pc}' = 6 \\
\rho_{\tau_4} &\equiv \text{pc} = 5 \wedge x' = x \wedge \text{ret}' = \text{ret} \wedge \text{pc}' = 6 \\
\rho_{l_0} &= (\text{pc} = 0 \wedge 0 \leq x \leq 127 \wedge x' = x \wedge \text{ret}' = x \wedge \text{pc}' = 6) \vee \\
&\quad (\text{pc} = 0 \wedge 128 \leq x \leq 255 \wedge x' = x \wedge \text{ret}' = 0 \wedge \text{pc}' = 6)
\end{aligned}$$

Figure 2: Transition system for Example 2

The set F of final states is given by

$$F = \{s \in S \mid s.\text{pc} = 6\}$$

The set of \mathcal{T} of transitions is given by $\mathcal{T} = \{\tau_1, \tau_2, \tau_3, \tau_4\}$ and is shown in Fig. 2 with the respective transition relations. The program returns values less than or equal to 127 unchanged and returns 0 otherwise, thus the set F_{reach} of reachable final states is given by:

$$\begin{aligned} F_{\text{reach}} = & \{s \in F \mid 0 \leq s.x \leq 127 \wedge s.\text{ret} = x \wedge s.\text{pc} = 6\} \\ & \cup \{s \in F \mid 128 \leq s.x \leq 255 \wedge s.\text{ret} = 0 \wedge s.\text{pc} = 6\} \end{aligned}$$

2 Exact static analysis

In this section, we present the conceptual framework and algorithm for the exact static QIF analysis method presented in [1]. This presentation differs from [1] in that we assume no high/low structure on the program states.

2.1 Qualitative information flow

We use *set partitions* to characterize partial knowledge about the underlying set. We call the elements of a partition *blocks*. A partition P of a set X models that each $x \in X$ is known up to its enclosing subset of X in P . We say that a partition P is *more precise* than P' (in symbols: $P \sqsubseteq P'$) if each partition block of P is contained in a partition block of P' :

$$P \sqsubseteq P' :\iff \forall A \in P : \exists B \in P' : A \subseteq B$$

Let $\mathcal{P} = (S, I, F, \mathcal{T})$ be a program. We require that \mathcal{P} implements a total function, i.e. from each initial state we can reach exactly one final state and, as a consequence, $I = \bigcup_{s' \in F_{\text{reach}}} \mathcal{P}^{-1}(s')$. The preimage operator thus induces a partition Π on the set of initial states I :

$$\Pi := \{\mathcal{P}^{-1}(s') \mid s' \in F_{\text{reach}}\}$$

We can qualitatively characterize the information a program reveals about its input in terms of the partition Π . For example, $\Pi = \{I\}$ models a program that does not reveal any information about its input. On the other hand, $\Pi = \{\{s\} \mid s \in I\}$ models a program that reveals its input completely.

2.2 Refinement for leak discovery

In order to construct the preimage partition Π , we compute an equivalence relation $\approx \subseteq I \times I$ such that Π is equal to the set of equivalence classes of \approx :

$$\Pi = I / \approx = \{[s]_{\approx} \mid s \in I\}$$

Given a program \mathcal{P} and an equivalence relation $R \subseteq I \times I$, there is an *information leak with respect to R* if there is a pair of paths π_1, π_2 corresponding to computations that lead from R -equivalent initial states $s_1 \sim_R s_2$ to unequal final states $s'_1 \neq s'_2$:

$$\begin{aligned} \text{Leak}_{\mathcal{P}}(R, \pi_1, \pi_2) \equiv & \\ & \exists s_1, s_2 \in I : \exists s'_1, s'_2 \in F_{\text{reach}} : \\ & (s_1, s'_1) \in \rho_{\pi_1} \wedge (s_2, s'_2) \in \rho_{\pi_2} \wedge s_1 \sim_R s_2 \wedge s'_1 \neq s'_2 \end{aligned}$$

A refinement algorithm

We construct \approx by successively refining a candidate relation R . R is initialized to the coarsest equivalence relation $R_0 = I \times I$ – i.e. $I/R_0 = \{I\}$ – and then successively refined until no more leaks can be discovered. This final relation is the largest relation over I such that no leaks in the above sense exist. We refine a relation R with leak paths π_1, π_2 by distinguishing between initial states that lead to different output states along π_1, π_2 . This is realized by splitting the equivalence classes of the initial states by removing the corresponding tuples from the equivalence relation. The refinement conjunct $\text{Refine}_{\mathcal{P}}$ implementing this is shown in (4). The leak discovery and refinement procedure is formalized in Algorithm 1.

$$\begin{aligned} \text{Refine}_{\mathcal{P}}(\pi_1, \pi_2) \equiv & \quad (4) \\ & \{(s_1, s_2) \in I \times I \mid \forall s'_1, s'_2 \in F_{\text{reach}} : ((s_1, s'_1) \in \rho_{\pi_1} \wedge (s_2, s'_2) \in \rho_{\pi_2}) \rightarrow s'_1 = s'_2\} \end{aligned}$$

Algorithm 1 DISCO – Information leak discovery

```

procedure DISCO( $\mathcal{P}$ ):
   $R \leftarrow I \times I$ 
  while exists  $\pi_1, \pi_2 \in \mathcal{T}^+ : \text{Leak}_{\mathcal{P}}(R, \pi_1, \pi_2)$  do
     $R \leftarrow R \cap \text{Refine}_{\mathcal{P}}(\pi_1, \pi_2)$ 
   $\approx \leftarrow R \cup \text{id}_I$ 
  return  $\approx$ 

```

Theorem 1. *The set of equivalence classes of the relation $R = \text{DISCO}(\mathcal{P})$ computed by DISCO is the preimage partition Π :*

$$\Pi = I/R$$

Proof. “ \subseteq ” Let $X = \mathcal{P}^{-1}(s') \in \Pi$ for some $s' \in F_{\text{reach}}$. Let $s_1, s_2 \in X$. Since we assume \mathcal{P} implements a total function (i.e. $\{\mathcal{P}^{-1}(s') \mid s' \in F_{\text{reach}}\}$ is a partition), no other final state $t' \in F_{\text{reach}}, t' \neq s'$ is reachable from s_1 or s_2 . It follows that $\forall \pi_1, \pi_2 \in \mathcal{T}^+ : (s_1, s'_1) \in \rho_{\pi_1} \wedge (s_2, s'_2) \in \rho_{\pi_2} \rightarrow s'_1 = s'_2 = s'$. Since R is the largest relation such that this property holds for all pairs of equivalent states, we obtain $s_1 \sim_R s_2$. Since s_1, s_2 were chosen arbitrarily from X , there exists an $Y \in I/R$ such that $X \subseteq Y$. Assume there exists a state $t \in Y \setminus X$. Then $\forall \pi_1, \pi_2 \in \mathcal{T}^+ : (s_1, s'_1) \in \rho_{\pi_1}, (t, t') \in \rho_{\pi_2} \rightarrow s'_1 = t'$. By the above argument, however, $t' = s'_1 = s'$ and thus $t \in \mathcal{P}^{-1}(s') = X$. Thus $X = Y \in I/R$.

“ \supseteq ” Let $X \in I/R, s_1, s_2 \in X$, i.e. $s_1 \sim_R s_2$. By construction of R , we have

$$\forall \pi_1, \pi_2 \in \mathcal{T}^+ : \forall s'_1, s'_2 \in F_{\text{reach}} : (s_1, s'_1) \in \rho_{\pi_1} \wedge (s_2, s'_2) \in \rho_{\pi_2} \rightarrow s'_1 = s'_2$$

Since \mathcal{P} implements a total function, there exists one state $s' \in F_{\text{reach}}$, such that $\forall \pi_1, \pi_2 \in \mathcal{T}^+ : (s_1, s'_1) \in \rho_{\pi_1} \wedge (s_2, s'_2) \in \rho_{\pi_2} \rightarrow s'_1 = s'_2 = s'$. It follows that $\{s_1, s_2\} \subseteq \mathcal{P}^{-1}(s')$ and thus $\exists Y = \mathcal{P}^{-1}(s') \in \Pi : X \subseteq Y$. As Π is a partition of I and X and Y have a nonempty intersection, $X = Y$. Thus $X = Y = \mathcal{P}^{-1}(s') \in \Pi$. \square

We have shown how to compute the preimage partition Π as the set of equivalence classes of an equivalence relation $\approx \subseteq I \times I$. We will now utilize Π to express quantitative information flow properties of the underlying program.

2.3 Quantitative information flow

In order to quantify the information leaked by a program, we model a program as a set of random variables and use Shannon entropy as an information measure. We derive a connection between the preimage partition Π and quantitative information flow in terms of Shannon entropy.

Programs as random variables

Given a program $\mathcal{P} = (\mathbb{S}, \mathbb{I}, \mathbb{F}, \mathcal{T})$, we consider the probability space (\mathbb{I}, \Pr) on the set of program inputs \mathbb{I} . In the following, we assume \Pr constant on \mathbb{I} , i.e. for all initial states $s \in \mathbb{I} : \Pr[s] = \frac{1}{|\mathbb{I}|}$. We model the behaviour of a program using the random variables U and V . The variable $U := \text{id}_{\mathbb{I}} : \mathbb{I} \rightarrow \mathbb{I}$ models the choice of program input, while $V := \rho_{\text{io}}(U) : \mathbb{I} \rightarrow \mathbb{F}$ models a computation of \mathcal{P} . For $s \in \mathbb{I}$ and $s' \in \mathbb{F}$, we can derive the following probabilities:

$$\begin{aligned} \Pr[V = s'] &= \Pr[\rho_{\text{io}}(U) = s'] \\ &= \frac{|\mathcal{P}^{-1}(s')|}{|\mathbb{I}|} \\ \Pr[U = s, V = s'] &= \begin{cases} \frac{1}{|\mathbb{I}|} & s \in \mathcal{P}^{-1}(s') \\ 0 & \text{otherwise.} \end{cases} \\ \Pr[U = s | V = s'] &= \frac{\Pr[U = s, V = s']}{\Pr[V = s']} \\ &= \begin{cases} \frac{1}{|\mathcal{P}^{-1}(s')|} & s \in \mathcal{P}^{-1}(s') \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Applying the concept of entropy outlined above, $H[U]$ describes the uncertainty of the program input. We assume that all input events have the same probability, thus $H[U] = \log_2 |\mathbb{I}|$. The conditional entropy $H[U|V = s']$ describes the uncertainty about the program input, *given a specific program output*. Using the probabilities derived above, we can express $H[U|V = s']$ in terms of the size of the preimage set $\mathcal{P}^{-1}(s')$:

$$\begin{aligned} H[U|V = s'] &= - \sum_{s \in \mathbb{I}} \underbrace{\Pr[U = s | V = s']}_{=0 \text{ for } s \notin \mathcal{P}^{-1}(s')} \log_2 \Pr[U = s | V = s'] \\ &= \sum_{s \in \mathcal{P}^{-1}(s')} \frac{1}{|\mathcal{P}^{-1}(s')|} \log_2 |\mathcal{P}^{-1}(s')| \\ &= |\mathcal{P}^{-1}(s')| \cdot \frac{1}{|\mathcal{P}^{-1}(s')|} \log_2 |\mathcal{P}^{-1}(s')| \\ &= \log_2 |\mathcal{P}^{-1}(s')| \end{aligned}$$

Since we are interested in how much information the program reveals *in general*, we consider the average conditional input entropy $H[U|V]$ over all final states. Using the expression for $H[U|V = s']$ derived above and the block sizes of the preimage

partition Π of I we obtain:

$$\begin{aligned}
\mathbb{H}[U|V] &= \sum_{s' \in F} \underbrace{\Pr[V = s']}_{=0 \text{ for } s' \notin F_{\text{reach}}} \mathbb{H}[U|V = s'] \\
&= \sum_{s' \in F_{\text{reach}}} \frac{|\mathcal{P}^{-1}(s')|}{|I|} \mathbb{H}[U|V = s'] \\
&= \sum_{s' \in F_{\text{reach}}} \frac{|\mathcal{P}^{-1}(s')|}{|I|} \log_2 |\mathcal{P}^{-1}(s')| \\
&= \frac{1}{|I|} \sum_{B \in \Pi} |B| \log_2 |B|
\end{aligned}$$

The last equation provides the connection between the quantitative (entropy) and qualitative (preimage partition) viewpoints. Since V is fully determined by U , the joint entropy $\mathbb{H}[U, V]$ of U and V equals the entropy $\mathbb{H}[U]$ of U , thus by the equation derived above, we also obtain

$$\mathbb{H}[U|V] = \mathbb{H}[U, V] - \mathbb{H}[V] = \mathbb{H}[U] - \mathbb{H}[V]$$

The information *leakage* \mathcal{L} of a program is the *reduction in uncertainty* of the value of the input U gained by observing an output V :

$$\mathcal{L} := \mathbb{H}[U] - \mathbb{H}[U|V] = \mathbb{H}[V]$$

Using the expression for $\mathbb{H}[U|V]$ derived above, we can express \mathcal{L} in terms of the block sizes of the preimage partition:

$$\mathcal{L} = \log_2 |I| - \frac{1}{|I|} \sum_{B \in \Pi} |B| \log_2 |B|$$

Example 3 (Information leakage). Again, consider the program from Example 2 given in Fig. 2. Since the program maps values $x \leq 127$ to themselves and values $x \geq 128$ to 0, the preimage partition Π of \mathcal{P} is given by

$$\Pi = \{\{s\} \mid s \in I \wedge 0 \leq s.x \leq 127\} \cup \{\{s \mid s \in I \wedge 128 \leq s.x \leq 255\}\}$$

Using this and $|I| = 256$, we can compute the information leakage \mathcal{L} of \mathcal{P} :

$$\begin{aligned}
\mathbb{H}[U|V] &= \frac{1}{|I|} \cdot \sum_{B \in \Pi} |B| \cdot \log_2 |B| \\
&= \frac{1}{256} \cdot (128 \cdot (1 \cdot \log_2 1) + (127 \cdot \log_2 127)) \\
&= \frac{1}{256} \cdot (127 \cdot \log_2 127) \\
&\approx 3.47 \text{ bit} \\
\mathcal{L} &= \log_2 |I| - \mathbb{H}[U|V] \approx 8 - 3.47 = 4.53 \text{ bit}
\end{aligned}$$

3 Approximation and randomization

In this section we utilize the formulation of information flow as reduction in Shannon entropy derived above and present an approximative static QIF analysis method from [3].

3.1 Approximation of information leakage

Computing the output entropy $H[V]$ exactly requires computing the block sizes of the preimage partition Π . In practice, this is often infeasible, even for finite-state programs, due to the large number of program states. We will show that we can approximate the block sizes of Π . Using this approximation, we can derive bounds on $H[V]$.

Approximation of the block sizes of the preimage partition Π

In order to approximate preimage partition block sizes of a program \mathcal{P} , we also need to approximate the input-output relation ρ_{io} . In order to do this, we construct an *augmented program* $\overline{\mathcal{P}} = (\overline{S}, \overline{I}, \overline{F}, \overline{\mathcal{T}})$, such that reachable final states in \mathcal{P} keep track of the corresponding initial states.

- $\overline{S} := S \times S$
- $\overline{I} := I \times I$
- $\overline{F} := S \times F$
- $\overline{\mathcal{T}} := \{\overline{\tau} \mid \tau \in \mathcal{T}\}$, where $\rho_{\overline{\tau}} := \{((s'', s), (s'', s')) \mid (s, s') \in \rho_{\tau} \wedge s'' \in S\}$

The program relation $\overline{\rho}$ and reachability can be defined analogously:

$$\overline{\rho} := \bigcup_{\overline{\tau} \in \overline{\mathcal{T}}} \rho_{\overline{\tau}} \quad \overline{F}_{\text{reach}} := \{(s, s') \in \overline{F} \mid s \in I, s' \in F, ((s, s), (s, s')) \in \overline{\rho}^*\} \\ = \rho_{\text{io}}$$

We can now apply the approximation methods described above to the augmented program. Using abstract interpretation, we compute the over-approximation $\overline{F}_{\text{reach}}^{\#}$ of $\overline{F}_{\text{reach}}$:

$$\overline{F}_{\text{reach}}^{\#} := \text{lfp}(\alpha \circ \text{Post}_{\overline{\rho}}, \alpha(\overline{I}))$$

In order to compute the under-approximation $\overline{F}_{\text{reach}}^b$, we symbolically execute a given set of program paths $\{\pi_1, \dots, \pi_n\}$ of $\overline{\mathcal{P}}$. We obtain the inclusions:

$$\overline{F}_{\text{reach}}^b \subseteq \overline{F}_{\text{reach}} \subseteq \overline{F}_{\text{reach}}^{\#}$$

This can be used to approximate the block sizes of the preimage partition Π of \mathcal{P} . Let $\overline{\mathcal{P}}^{b-1}(s')$, $\overline{\mathcal{P}}^{\#-1}(s')$ denote the sets of projections of each tuple – in $\overline{F}_{\text{reach}}^b$ and $\overline{F}_{\text{reach}}^{\#}$, respectively – to the first component:

$$\overline{\mathcal{P}}^{b-1}(s') := \{s \mid (s, s') \in \overline{F}_{\text{reach}}^b\} \subseteq I \\ \overline{\mathcal{P}}^{\#-1}(s') := \{s \mid (s, s') \in \overline{F}_{\text{reach}}^{\#}\} \subseteq I$$

We thus have:

$$\overline{\mathcal{P}}^{b-1}(s') \subseteq \mathcal{P}^{-1}(s') \subseteq \overline{\mathcal{P}}^{\#-1}(s')$$

and followingly:

$$|\overline{\mathcal{P}}^{b-1}(s')| \leq |\mathcal{P}^{-1}(s')| \leq |\overline{\mathcal{P}}^{\#-1}(s')|$$

Approximation of input entropy

Using the approximations of \mathcal{P}^{-1} obtained above, we can derive upper and lower bounds for the conditional input entropy $\mathsf{H}[U|V]$.

Lemma 1. *For $s' \in F_{\text{reach}}$ we define:*

$$p_V^{\flat}(s') := \frac{1}{|I|} \max\left\{\left|\overline{\mathcal{P}}^{\flat-1}(s')\right|, 1\right\}$$

$$p_V^{\sharp}(s') := \frac{1}{|I|} \left|\overline{\mathcal{P}}^{\sharp-1}(s')\right|$$

Then p_V^{\flat} and p_V^{\sharp} under- and over-approximate p_V , i.e. for all $s' \in F_{\text{reach}}$:

$$p_V^{\flat}(s') \leq \Pr[V = s'] \leq p_V^{\sharp}(s')$$

Proof. Let $s' \in F_{\text{reach}}$. Above we have shown that $\Pr[V = s'] = |\mathcal{P}^{-1}(s')| |I|^{-1}$ and that $\left|\overline{\mathcal{P}}^{\flat-1}(s')\right|$ and $\left|\overline{\mathcal{P}}^{\sharp-1}(s')\right|$ approximate $|\mathcal{P}^{-1}(s')|$. Since $s' \in F_{\text{reach}}$, at least one initial state leads to s' , thus $|\mathcal{P}^{-1}(s')| \geq 1$. The statement follows. \square

Using the approximations obtained above, we can give lower and upper bounds on $\mathsf{H}[V]$. We show how to derive the lower bound:

$$\begin{aligned} \mathsf{H}[V] &= \sum_{s' \in F_{\text{reach}}} \Pr[V = s'] (-\log_2 \Pr[V = s']) \\ &\geq \sum_{s' \in F_{\text{reach}}^{\flat}} \Pr[V = s'] (-\log_2 \Pr[V = s']) && \text{by } F_{\text{reach}}^{\flat} \subseteq F_{\text{reach}} \\ &\geq \sum_{s' \in F_{\text{reach}}^{\flat}} \Pr[V = s'] (-\log_2 p_V^{\sharp}(s')) && \text{Monotonicity of } -\log_2 \text{ on }]0, 1] \\ &\geq \sum_{s' \in F_{\text{reach}}^{\flat}} p_V^{\flat}(s') (-\log_2 p_V^{\sharp}(s')) \end{aligned}$$

The upper bound follows similarly and we obtain:

$$\begin{aligned} - \sum_{s' \in F_{\text{reach}}^{\flat}} p_V^{\flat}(s') \log_2 p_V^{\sharp}(s') \\ \leq \mathsf{H}[V] \leq \\ - \sum_{s' \in F_{\text{reach}}^{\sharp}} p_V^{\sharp}(s') \log_2 p_V^{\flat}(s') \end{aligned}$$

Above we have shown:

$$\mathsf{H}[U|V] = \mathsf{H}[U] - \mathsf{H}[V]$$

Assuming U is uniformly distributed on I , we have $\mathsf{H}[U] = \log_2 |I|$. Using the bounds on $\mathsf{H}[V]$ we can bound the value of $\mathsf{H}[U|V]$:

$$\begin{aligned} \log_2 |I| + \sum_{s' \in F_{\text{reach}}^{\sharp}} p_V^{\sharp}(s') \log_2 p_V^{\flat}(s') \\ \leq \mathsf{H}[U|V] \leq \\ \log_2 |I| + \sum_{s' \in F_{\text{reach}}^{\flat}} p_V^{\flat}(s') \log_2 p_V^{\sharp}(s') \end{aligned}$$

3.2 Randomized approximation

While the approximation methods above are effective in more cases than exact computation, they require the *enumeration* of the over-approximation $F_{\text{reach}}^{\sharp}$ of the set F_{reach} of reachable final states. This is frequently impractical due to the size of that set. We will show that we can avoid this enumeration.

Estimating entropy

Our goal is to construct a sampling-based algorithm that can compute upper and lower bounds for the conditional input entropy $H[U|V]$. Since $H[U|V] = \log_2 |I| - H[V]$, it is sufficient to approximate the entropy $H[V]$ of V . Thus, as a first step, we show that we can construct an estimator for the entropy of a random variable.

Lemma 2. *Let $|\Omega| = m$, $X : \Omega_X \rightarrow \Omega$, $X \sim p_X$ a random variable, $0 < \delta \in \mathbb{R}$, random variables $X_1 \simeq \dots \simeq X_n \simeq X$ independent instances (i.e. samples) of X . Then*

$$\Pr \left[\left| -\frac{1}{n} \sum_{i=1}^n \log_2 (\Pr[X = X_i]) - H[X] \right| < \delta \right] > 1 - \frac{\log_2 m^2}{n\delta^2}$$

Proof. We define random variables $Y, Y_i : \Omega_X \rightarrow \mathbb{R}$, $Y \simeq Y_i$ for $i = 1 \dots n$ by

$$Y_i = -\log_2 (p_X(X_i))$$

Each Y_i is an unbiased estimator for the entropy $H[X]$ of X :

$$\mathbb{E}[Y_i] = \sum_{y \in \mathbb{R}} \Pr[Y_i = y] \cdot y = - \sum_{x \in \Omega} \Pr[X = x] \cdot \log_2 (p_X(x)) = H[X]$$

Consider a random variable Z that is the average of the Y_i :

$$Z := \frac{1}{n} \sum_{i=1}^n Y_i$$

By linearity of the expectation value, $\mathbb{E}[Z] = H[X]$. For the variance $\text{Var}[Z]$ of Z we obtain:

$$\text{Var}[Z] = \frac{1}{n} \text{Var}[Y]$$

In [2] it is shown that $\text{Var}[Y] \leq (\log_2 m)^2$, thus

$$\text{Var}[Z] \leq \frac{(\log_2(m))^2}{n}$$

Let $0 < \delta \in \mathbb{R}$. We can apply Chebyshev's inequality:

$$\Pr[|Z - \mathbb{E}[Z]| \geq \delta] \leq \frac{\text{Var}[Z]}{\delta^2} \leq \frac{(\log_2 m)^2}{n\delta^2}$$

Using $\mathbb{E}[Z] = H[X]$ and unfolding the definitions of Z and Y_i , we obtain:

$$\Pr \left[\left| -\frac{1}{n} \sum_{i=1}^n \log_2 (p_X(X_i)) - H[X] \right| < \delta \right] > 1 - \frac{(\log_2 m)^2}{n\delta^2}$$

□

We have shown that, given the probability distribution of a random variable X , we can use a sampling estimator Z to bound the entropy $H[X]$ of X . The confidence of $p = 1 - \frac{(\log_2 m)^2}{n\delta^2}$ for $H[X] \in [Z - \delta, Z + \delta]$ allows us to judge the quality of the estimation. Conversely, for a desired confidence $p \in [0, 1[$ and precision $\delta \in]0, \infty[$ we can give the minimum number $n = \left\lceil \frac{(\log_2 m)^2}{(1-p)\delta^2} \right\rceil$ of required samples.

A randomized algorithm

In order to bound the output entropy $H[V]$ for output samples $X_i \simeq V$ using the estimator Z derived above, we would need to know the exact probability distribution of V . This still requires the exact computation of preimage sizes. Our goal is to avoid this. Instead, we can apply the approximations p_V^\sharp and p_V^\flat of the probability distribution p_V of V from the previous section. In terms of the estimators Y_i and Z we derived above, we construct the following over- and under-approximations:

$$\begin{aligned} Y_i^\flat &:= -\log_2(p_V^\flat(X_i)) && \leq Y_i \leq && Y_i^\sharp &:= -\log_2(p_V^\sharp(X_i)) \\ Z^\flat &:= \frac{1}{n} \sum_{i=1}^n Y_i^\flat && \leq Z \leq && Z^\sharp &:= \frac{1}{n} \sum_{i=1}^n Y_i^\sharp \end{aligned}$$

We obtain the following bounds on $H[V]$:

$$Z^\flat - \delta \leq Z - \delta < H[V] < Z + \delta \leq Z^\sharp + \delta$$

Bounding $H[U|V]$ using the bounds on $H[V]$ we obtain:

$$\begin{aligned} H^\flat &:= \log_2 |I| - Z^\sharp \leq H[U|V] + \delta \\ H[U|V] - \delta &\leq \log_2 |I| - Z^\flat =: H^\sharp \end{aligned}$$

By Lemma 2 these bounds are valid with a probability $p > 1 - \frac{(\log_2 |I|)^2}{n\delta^2}$. We can simplify the approximations:

$$\begin{aligned} H^\flat &= \log_2 |I| - Z^\sharp = \log_2 |I| - \frac{1}{n} \sum_{i=1}^n -\log_2(p_V^\flat(X_i)) \\ &= \log_2 |I| + \frac{1}{n} \sum_{i=1}^n \log_2 \left(\frac{|\overline{\mathcal{P}}^{\flat-1}(X_i)|}{|I|} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \log_2 \left(\min \left\{ 1, |\overline{\mathcal{P}}^{\flat-1}(X_i)| \right\} \right) \\ H^\sharp &= \log_2 |I| - Z^\flat = \log_2 |I| - \frac{1}{n} \sum_{i=1}^n -\log_2(p_V^\sharp(X_i)) \\ &= \log_2 |I| + \frac{1}{n} \sum_{i=1}^n \log_2 \left(\frac{|\overline{\mathcal{P}}^{\sharp-1}(X_i)|}{|I|} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \log_2 \left(|\overline{\mathcal{P}}^{\sharp-1}(X_i)| \right) \end{aligned}$$

The estimation procedure derived above is formalized in Algorithm 2.

Algorithm 2 RANT – Randomized quantification

procedure RANT(\mathcal{P}, p, δ):
 $(H^\flat, H^\sharp, n) \leftarrow (0, 0, \left\lceil \frac{(\log_2 |I|)^2}{(1-p)\delta^2} \right\rceil)$
for $i = 1 \dots n$ **do**
 $s_i \leftarrow$ pick randomly from I
 $(\pi_i, s') \leftarrow$ execute \mathcal{P} symbolically on s_i
 $H^\flat \leftarrow H^\flat + \log_2 |\rho_{\pi_i}^{-1}(s')|$
 $H^\sharp \leftarrow H^\sharp + \log_2 |\overline{\mathcal{P}}^{\sharp-1}(s')|$
return $(\frac{H^\flat}{n}, \frac{H^\sharp}{n})$

Theorem 2. For $(H^b, H^\sharp) = \text{RANT}(\mathcal{P}, p, \delta)$ the following holds:

$$\Pr [H^b - \delta \leq H[U|V] \leq H^\sharp + \delta] \geq p$$

Proof. RANT computes the over-approximation H^\sharp and under-approximation H^b of the entropy estimator $H = \log_2 |I| - Z$ for $n = \left\lceil \frac{(\log_2 |I|)^2}{(1-p)\delta^2} \right\rceil$ samples. The statement follows from Lemma 2 and the construction of H^\sharp and H^b . \square

4 Approximative dynamic analysis

The methods presented in the previous sections are *static* analysis methods – they bound the information leakage of a program for all possible executions. In this section, we present a *dynamic* analysis method described in [5]. The goal is to provide a sound upper bound on the information leakage of a *single execution*. Additionally, the method provides for a way of combining the results of several single-execution analyses to obtain more general bounds.

4.1 Assembly programs and computations

We consider assembly-like programs defined by the grammar shown in Fig. 3. The element **reg** represents a register identifier, **addr: reg** represents the usage of the value of a register as a memory address. A register identifier on the left-hand side of an assignment represents the register as a variable, a register identifier on the right-hand side represents its value. The **input** instruction writes a (secret) input value to a register, **output** makes the contents of a register public.

$$\begin{aligned}
 r \in \text{reg} & ::= r_n : n \in \mathbb{N} \\
 c \in \text{cst} & ::= n \in \mathbb{N} \\
 \circ \in \text{binop} & ::= + \mid \text{and} \mid \text{xor} \\
 i \in \text{instr} & ::= \text{reg} := \text{cst} \\
 & \mid \text{reg} := \text{reg} \\
 & \mid \text{reg} := \text{reg} \circ \text{reg} : \circ \in \text{binop} \\
 & \mid \text{reg} := \text{load} (\text{addr} : \text{reg}) \\
 & \mid \text{store} (\text{addr} : \text{reg}) \text{ reg} \\
 & \mid \text{output} \text{ reg} \\
 & \mid \text{input} \text{ reg} \\
 & \mid \text{if} \text{ reg} \geq 0 \text{ goto} \text{ cst} \\
 p \in \text{program} & ::= \epsilon \mid \text{instr program}
 \end{aligned}$$

Figure 3: Syntax for a minimal imperative language

4.2 Dynamic tainting and bit-tracking

Dynamic tainting analysis computes the secrecy status for each piece of data occurring during program execution. A piece of data is considered *tainted* if it contains secret information. The amount of information leaked by the program is bounded by the number of secret bits present in the program’s output. Tainting can be computed on various levels of granularity – variables, bytes and individual bits could

be marked as secret. Once a piece of data is marked as secret, any operation on it must be instrumented to compute the secrecy status of the result. In the case of bit-level tainting, this is referred to as bit-tracking. Each register and memory location is associated with a *secrecy bit mask* which specifies which of its bits are marked as secret. Examples of bit-tracking formulas for several bit-level operations are given in [5].

4.3 Flow graphs

In order to model dynamic QIF analysis as a maximum-flow problem, we construct a directed acyclic graph, which we refer to as the *flow graph*. The flow graph has potentially secret pieces of information as its vertices (*nodes*). The edges connecting the nodes have capacities that model how much secret information can flow from one piece of data to another.

The edge capacities are additionally bounded by a bit-tracking analysis. Each vertex is associated with the secrecy bit mask of the piece of data it represents. An edge originating from a vertex v can carry at most as many bits of secret information as are set in the secrecy bit mask of v .

The introduction of secret information through **input** instructions is modeled as an edge from a special *source node* to a new node representing the input data. The revelation of potentially secret information by **output** instructions is modeled as an edge to a special *sink node*. The *information leakage* of a program is bounded by the maximum flow from the source node to the sink node.

The graph is constructed dynamically by keeping track of the source of each piece of data during an execution. In the algorithm `INSTRFLOW` presented below, this is implemented by mapping each register and memory location to a vertex of the flow graph and updating this mapping upon any modification of data. An example of a flow graph for a specific program execution is given in Fig. 4.

Combining flow graphs

Flow graphs corresponding to different executions of a program can be merged to obtain bounds that are valid for several executions. This is done by identifying edges corresponding to a particular program location in all graphs and merging the source nodes and destination nodes to obtain two nodes and one edge for each set of original edges and nodes. The capacity of the merged edge is the sum of the capacities of the original edges. Finally, the source and sink nodes of all graphs are merged to obtain new sink and source nodes.

Implicit and explicit flows

An *explicit flow* of secret information is one that occurs by copying or mapping over secret data (for example in the assignment $r_i := r_j \circ r_k$, where either of the operands is secret, secret information explicitly flows from the operands to r_i). An *implicit flow* occurs when secret information is revealed indirectly. In the language defined above, this applies to case distinctions over the contents of registers containing secret data (`if $r_i \geq 0$ goto c`).

We over-approximate the secret information flow for the conditional **goto** as an additional 1 bit of flow to any of the following **output** statements. This captures the intuition that the result of the comparison might affect the value of any publicly visible register. We account for implicit flows by adding a chain of “leak nodes”, to which implicit flows from comparisons on registers are directed.

On each **output** operation, a new node representing the output information is created and an edge from the current leak node to the output node is added. This

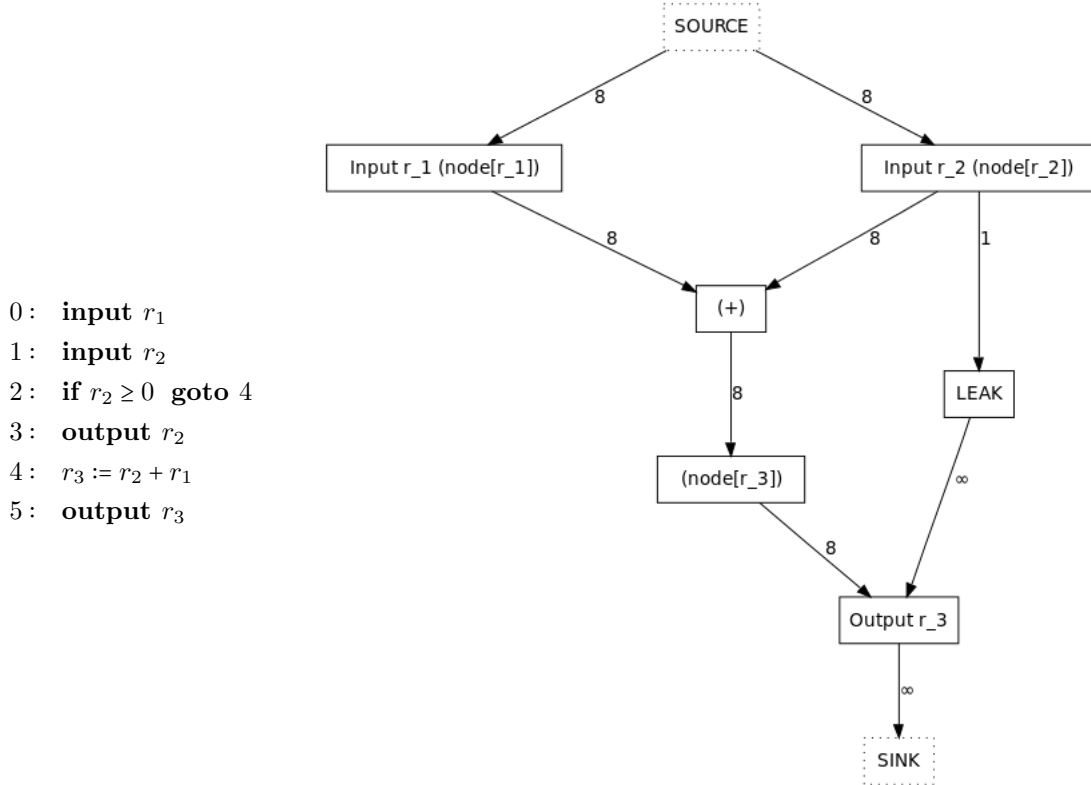


Figure 4: Example program (using 8-bit registers) and the flow graph corresponding to an execution with inputs $(-1,0)$. The upper bound on information leakage, computed as maximum flow, is 9 bits. The labels “node[r.i]” refer to the node map explained in the presentation of the construction algorithm.

models the possibility that a previous comparison on secret data might have an effect on the output value. A new leak node is created and an unbounded edge is added between the old and new leak node to allow implicit flows to propagate to later outputs.

Instrumentation for flow graph construction

The algorithm INSTRFLOW for incremental construction of the flow graph is presented in Algorithm 3. The program is instrumented such that INSTRFLOW is applied to every instruction encountered. For each register and memory location, we keep track of the source of the information it contains. The map `node` associates a register $r_i \in \text{reg}$ or memory location $(\text{addr} : m) \in \text{addr}$ with its node (v, bits) . The node contains the vertex $v \in V$ in the flow graph and the secrecy bitmask $\text{bits} \in (\mathbb{N}_0 \bmod 2^{\text{regbits}})$ for its contents.

When a value is to a register, the register’s `node` mapping is set to the node of the rhs value. If a register or memory location is modified, a new node is created. The node n corresponding to the original information is connected to the node n' representing the modified data and the `node` map is updated accordingly. The capacity c of the new edge $n \rightarrow^c n'$ is at most the number $\#(n.\text{bits})$ of set bits in the secret bit mask of its source node. Implicitly, a new edge $n \rightarrow^c n'$ represents a new edge $n.v \rightarrow^c n'.v$. Binary operations are treated analogously. Source, sink and leak nodes are treated as special cases: edges between leak nodes as well as edges to the

sink node are unbounded, edges from leak nodes to output nodes have a capacity of 1.

Enclosure regions

Additionally to the algorithm `INSTRFLOW` presented above, the method in [5] implements *enclosure regions*, specified by source code annotations. An enclosure region denotes the restriction of implicit flows within a part of the source code to explicitly specified “output locations”. Implicit flows occurring within an enclosure region can only affect the values of the region’s defined output locations. When the enclosure annotations are correct, this method can improve precision while still retaining the soundness of the upper bound on information leakage.

Conclusion

We have presented exact and approximative static methods as well as an approximative dynamic method to quantify information flow properties of programs. While the exact static analysis method [1] yields precise results, it generally cannot be applied in practice due to prohibitive computational cost.

The approximative dynamic analysis method [5] only produces an upper bound on the information leaked by a program for a finite set of program paths, but is computationally cheap and can be applied to pre-compiled binaries.

Randomized approximative static analysis (as presented in [3]) generates upper and lower bounds on information leakage over all paths. Additionally, the method allows a controlled cost-versus-precision trade-off. By providing statistical quality guarantees, it allows the user to judge the tightness and validity of the generated bounds to enable meaningful decisions with respect to an information flow policy.

References

- [1] Michael Backes, Boris Köpf, and Andrey Rybalchenko. Automatic Discovery and Quantification of Information Leaks. In *Proc. IEEE Symp. on Security and Privacy (S&P '09)*, pages 141–153. IEEE, 2009.
- [2] Tugkan Batu, Sanjoy Dasgupta, Ravi Kumar, and Ronitt Rubinfeld. The complexity of approximating entropy. In *Proc. ACM Symp. on Theory of Computing (STOC '02)*, pages 678–687. ACM, 2002.
- [3] Boris Köpf and Andrey Rybalchenko. Approximation and randomization for quantitative information-flow analysis. In *Proceedings of the 23rd IEEE Computer Security Foundations Symposium, CSF 2010, Edinburgh, United Kingdom, July 17-19, 2010*, pages 3–14, 2010.
- [4] Pasquale Malacaria. Assessing security threats of looping constructs. In *Proc. Symp. on Principles of Programming Languages (POPL '07)*, pages 225–235. ACM, 2007.
- [5] Stephen McCamant and Michael D. Ernst. Quantitative information flow as network flow capacity. In *Proc. ACM Conf. on Programming Language Design and Implementation (PLDI '08)*, pages 193–205. ACM, 2008.
- [6] Claude E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423 and 623–656, July and October 1948.

- [7] Geoffrey Smith. On the foundations of quantitative information flow. In *Proc. Intl. Conf. of Foundations of Software Science and Computation Structures (FoSSaCS '09)*, LNCS 5504, pages 288–302. Springer, 2009.

Algorithm 3 INSTRFLOW($i \in \text{instr}$)

```
vars
  const regbits  $\in \mathbb{N}$  : length of a register in bits
  global  $V \ni \{\text{sink}, \text{source}, \text{public}\}$ ,  $E \subseteq V^2 \times \mathbb{N}_0$  : nodes, edges with capacities
  global leak  $\in V$ 
  global node  $\subseteq (\text{reg} \cup \text{addr}) \times V \times (\mathbb{N}_0 \bmod 2^{\text{regbits}})$ 
end vars
begin
  match i with
  |  $r_i := c$   $\rightarrow$ 
    node[ $r_i$ ] := public
  |  $r_i := r_j$   $\rightarrow$ 
    node[ $r_i$ ] := node[ $r_j$ ]
  |  $r_i := r_j \circ r_k$   $\rightarrow$ 
    begin
       $m, n$  := fresh nodes
       $n.\text{bits} := m.\text{bits} := \text{opBits}(\circ, r_j, r_k)$ 
       $c_1 := \#(\text{node}[r_j].\text{bits})$ 
       $c_2 := \#(\text{node}[r_k].\text{bits})$ 
       $c_3 := \#(n.\text{bits})$ 
       $E := E \cup \{(\text{node}[r_j] \rightarrow^{c_1} n), (\text{node}[r_k] \rightarrow^{c_2} n), (n \rightarrow^{c_3} m)\}$ 
      node[ $r_i$ ] :=  $m$ 
    end
  | input  $r_i$   $\rightarrow$ 
    begin
       $n$  := fresh node with  $n.\text{bits} = (1^{\text{regbits}})_2$ 
      node[ $r_i$ ] :=  $n$ 
       $E := E \cup \{(\text{source} \rightarrow^{\text{regbits}} n)\}$ 
    end
  | output  $r_i$   $\rightarrow$ 
    begin
       $n, l$  := fresh nodes
       $c := \#(\text{node}[r_i].\text{bits})$ 
       $E := E \cup \{(\text{node}[r_i] \rightarrow^c n), (\text{leak} \rightarrow^\infty n), (\text{leak} \rightarrow^\infty l), (n \rightarrow^\infty \text{sink})\}$ 
      leak :=  $l$ 
    end
  | if  $r_i \geq 0$  goto  $loc$   $\rightarrow$ 
       $E := E \cup \{(\text{node}[r_i] \rightarrow^1 \text{leak})\}$ 
  | store ( $\text{addr} : r_a$ )  $r_v$   $\rightarrow$ 
      node[ $\text{addr} : r_a$ ] := node[ $r_v$ ]
  |  $r_i := \text{load}$  ( $\text{addr} : r_a$ )  $\rightarrow$ 
      node[ $r_i$ ] := node[ $\text{addr} : r_a$ ]
```

end.
