

Security of Multithreaded Programms by Compilation

Paper written by Barthe, Rezk, Russo and Sabelfeld [1]

Pascal Wittmann

TU Darmstadt

Seminar “Formal Specification”

December 1–2, 2011

Outline

- Why formal methods?
- Security problems of multithreaded programs.
- Discussion of a solution.
- Other/related solutions.
- Conclusion / Outlook.

Why formal methods?

- Modeling precisely a part of the world
- Formulate the problem unambiguous
- Leaving unimportant things underspecified
- Improve the understanding of the problem
- Use abstraction to cover a large number of cases

Security problems of multithreaded programs

- There are private (*high*) and public (*low*) variables
- The attacker can observe low-level variables
- Sequential:
 - explicit flows: `lo := hi`
 - implicit flows: `if hi then lo := 1 else lo := 0`
- Concurrent:
 - internal timing leak:


```
if hi {sleep(100)}; lo := 1 || sleep(50); lo := 0
```
 - other example: `hi := 0; lo = hi || hi := private-data`
- External timing leaks are not covered
- Advantages of formal methods
 - Applicable on a wide range of schedulers and bytecode
 - Verification without running the program

Security problems of multithreaded programs

- There are private (*high*) and public (*low*) variables
- The attacker can observe low-level variables
- Sequential:
 - explicit flows: `lo := hi`
 - implicit flows: `if hi then lo := 1 else lo := 0`
- Concurrent:
 - internal timing leak:
`if hi {sleep(100)}; lo := 1 || sleep(50); lo := 0`
 - other example: `hi := 0; lo = hi || hi := private-data`
- External timing leaks are not covered
- Advantages of formal methods
 - Applicable on a wide range of schedulers and bytecode
 - Verification without running the program

Security problems of multithreaded programs

- There are private (*high*) and public (*low*) variables
- The attacker can observe low-level variables
- Sequential:
 - explicit flows: `lo := hi`
 - implicit flows: `if hi then lo := 1 else lo := 0`
- Concurrent:
 - internal timing leak:
 - `if hi {sleep(100)}; lo := 1 || sleep(50); lo := 0`
 - other example: `hi := 0; lo = hi || hi := private-data`
- External timing leaks are not covered
- Advantages of formal methods
 - Applicable on a wide range of schedulers and bytecode
 - Verification without running the program

Security problems of multithreaded programs

- There are private (*high*) and public (*low*) variables
- The attacker can observe low-level variables
- Sequential:
 - explicit flows: `lo := hi`
 - implicit flows: `if hi then lo := 1 else lo := 0`
- Concurrent:
 - internal timing leak:
 - `if hi {sleep(100)}; lo := 1 || sleep(50); lo := 0`
 - other example: `hi := 0; lo = hi || hi := private-data`
- External timing leaks are not covered
- Advantages of formal methods
 - Applicable on a wide range of schedulers and bytecode
 - Verification without running the program

Security problems of multithreaded programs

- There are private (*high*) and public (*low*) variables
- The attacker can observe low-level variables
- Sequential:
 - explicit flows: `lo := hi`
 - implicit flows: `if hi then lo := 1 else lo := 0`
- Concurrent:
 - internal timing leak:
 - `if hi {sleep(100)}; lo := 1 || sleep(50); lo := 0`
 - other example: `hi := 0; lo = hi || hi := private-data`
- External timing leaks are not covered
- Advantages of formal methods
 - Applicable on a wide range of schedulers and bytecode
 - Verification without running the program

Discussion of a solution

- Syntax & Semantic of multithreaded programs
 - Program
 - State & Security environment
 - History & Scheduler
- Type system & it's soundness
- The next function
- Concrete instantiation
 - Transfer rules
 - Defining the next function

Program

We have a set of sequential Instructions *SeqIns* and a primitive start *pc* that spawns a new thread.

Definition (Program P)

- 1 A set of program points \mathcal{P} , with a distinguished entry point 1 and exit point *exit*
- 2 A map from \mathcal{P} to *Ins*, where $Ins = SeqIns \cup \{startpc\}$ and $pc \in \mathcal{P} \setminus \{stop\}$. This map is referred to as $P[i]$.

Further, a relation $\mapsto \subseteq \mathcal{P} \times \mathcal{P}$ that describes possible successor instructions and its reflexive and transitive closure \mapsto^* .

Program

We have a set of sequential Instructions *SeqIns* and a primitive start *pc* that spawns a new thread.

Definition (Program P)

- 1 A set of program points \mathcal{P} , with a distinguished entry point 1 and exit point *exit*
- 2 A map from \mathcal{P} to *Ins*, where $Ins = SeqIns \cup \{startpc\}$ and $pc \in \mathcal{P} \setminus \{stop\}$. This map is referred to as $P[i]$.

Further, a relation $\mapsto \subseteq \mathcal{P} \times \mathcal{P}$ that describes possible successor instructions and its reflexive and transitive closure \mapsto^* .

Program

We have a set of sequential Instructions $SeqIns$ and a primitive start pc that spawns a new thread.

Definition (Program P)

- 1 A set of program points \mathcal{P} , with a distinguished entry point 1 and exit point $exit$
- 2 A map from \mathcal{P} to Ins , where $Ins = SeqIns \cup \{startpc\}$ and $pc \in \mathcal{P} \setminus \{stop\}$. This map is referred to as $P[i]$.

Further, a relation $\mapsto \subseteq \mathcal{P} \times \mathcal{P}$ that describes possible successor instructions and it's reflexive and transitive closure \mapsto^* .

State

We have a set of local states, `LocState` and a global memory `GMemory`. In Addition we have a set of thread identifiers `Thread`.

Definition (State)

- 1 `SeqState` is a product $\text{LocState} \times \text{GMemory}$
- 2 `ConcState` is a product $(\text{Thread} \rightarrow \text{LocState}) \times \text{GMemory}$

Accessors for a state s :

- `s.lst` and `s.gmem` are projections on the first and second component
- `s.act` is the set of active threads
- `s.pc(tid)` retrieves the current program point of the thread `tid`

State

We have a set of local states, `LocState` and a global memory `GMemory`. In Addition we have a set of thread identifiers `Thread`.

Definition (State)

- 1 `SeqState` is a product $\text{LocState} \times \text{GMemory}$
- 2 `ConcState` is a product $(\text{Thread} \rightarrow \text{LocState}) \times \text{GMemory}$

Accessors for a state s :

- `s.lst` and `s.gmem` are projections on the first and second component
- `s.act` is the set of active threads
- `s.pc(tid)` retrieves the current program point of the thread `tid`

State

We have a set of local states, `LocState` and a global memory `GMemory`. In Addition we have a set of thread identifiers `Thread`.

Definition (State)

- 1 `SeqState` is a product $\text{LocState} \times \text{GMemory}$
- 2 `ConcState` is a product $(\text{Thread} \rightarrow \text{LocState}) \times \text{GMemory}$

Accessors for a state s :

- $s.lst$ and $s.gmem$ are projections on the first and second component
- $s.act$ is the set of active threads
- $s.pc(tid)$ retrieves the current program point of the thread tid

Security environment

We assume a set of levels $\text{Level} = \{low, high\}$ where $low < high$ with an attacker on level low .

Definition (Security environment)

- ① A function $se : \mathcal{P} \rightarrow \text{Level}$
- ② A program point $i \in \mathcal{P}$ is:
 - low if $se(i) = low$, written $L(i)$
 - high if $se(i) = high$, written $H(i)$
 - always high if $\forall j \in \mathcal{P}. (i \mapsto^* j) \rightarrow se(j) = high$, written $AH(i)$

Now we classify threads in (where s is a `ConcState`):

$$s.lowT = \{tid \in s.act \mid L(s.pc(tid))\}$$

$$s.highT = \{tid \in s.act \mid H(s.pc(tid))\}$$

$$s.ahighT = \{tid \in s.act \mid AH(s.pc(tid))\}$$

$$s.hidT = \{tid \in s.act \mid H(s.pc(tid)) \wedge \neg AH(s.pc(tid))\}$$

Security environment

We assume a set of levels $\text{Level} = \{low, high\}$ where $low < high$ with an attacker on level low .

Definition (Security environment)

- ① A function $se : \mathcal{P} \rightarrow \text{Level}$
- ② A program point $i \in \mathcal{P}$ is:
 - low if $se(i) = low$, written $L(i)$
 - high if $se(i) = high$, written $H(i)$
 - always high if $\forall j \in \mathcal{P}. (i \mapsto^* j) \rightarrow se(j) = high$, written $AH(i)$

Now we classify threads in (where s is a `ConcState`):

$$s.lowT = \{tid \in s.act \mid L(s.pc(tid))\}$$

$$s.highT = \{tid \in s.act \mid H(s.pc(tid))\}$$

$$s.ahighT = \{tid \in s.act \mid AH(s.pc(tid))\}$$

$$s.hidT = \{tid \in s.act \mid H(s.pc(tid)) \wedge \neg AH(s.pc(tid))\}$$

Security environment

We assume a set of levels $\text{Level} = \{low, high\}$ where $low < high$ with an attacker on level low .

Definition (Security environment)

- ① A function $se : \mathcal{P} \rightarrow \text{Level}$
- ② A program point $i \in \mathcal{P}$ is:
 - low if $se(i) = low$, written $L(i)$
 - high if $se(i) = high$, written $H(i)$
 - always high if $\forall j \in \mathcal{P}. (i \mapsto^* j) \rightarrow se(j) = high$, written $AH(i)$

Now we classify threads in (where s is a `ConcState`):

$$s.lowT = \{tid \in s.act \mid L(s.pc(tid))\}$$

$$s.highT = \{tid \in s.act \mid H(s.pc(tid))\}$$

$$s.ahighT = \{tid \in s.act \mid AH(s.pc(tid))\}$$

$$s.hidT = \{tid \in s.act \mid H(s.pc(tid)) \wedge \neg AH(s.pc(tid))\}$$

History & Scheduler

Definition (History)

A History *History* is a list of pairs (tid, l) , where $tid \in \text{Thread}$ and $l \in \text{Level}$.

Definition (Scheduler)

A scheduler is a function $pickt : \text{ConcState} \times \text{History} \rightarrow \text{Thread}$ that satisfies these conditions:

- 1 Always picks active threads
- 2 if $s.\text{hidT} \neq \emptyset$ then $pick(s, h) \in s.\text{hidT}$
- 3 Only uses low names and the low part of the history to pick a low thread

History & Scheduler

Definition (History)

A History *History* is a list of pairs (tid, l) , where $tid \in \text{Thread}$ and $l \in \text{Level}$.

Definition (Scheduler)

A scheduler is a function $pickt : \text{ConcState} \times \text{History} \rightarrow \text{Thread}$ that satisfies these conditions:

- 1 Always picks active threads
- 2 if $s.\text{hidT} \neq \emptyset$ then $pick(s, h) \in s.\text{hightT}$
- 3 Only uses low names and the low part of the history to pick a low thread

Type system

LType is a poset (\leq is reflexive, antisymmetric, transitive) of local types.

Intuition of the type judgements: $se, i \vdash s \Rightarrow t$ means if executing program point i the type changes from s to t w.r.t a security environment se .

Definition (Typable program)

A program is typable (written $se, \mathcal{S} \vdash P$) if

- 1 for all initial program points holds $\mathcal{S}(i) = t_{init}$ and
- 2 $\forall i, j \in \mathcal{P} : (i \mapsto j) \rightarrow \exists s \in \text{LType} . se, i \vdash \mathcal{S}(i) \Rightarrow s \wedge \mathcal{S}(j) \leq s$

where $\mathcal{S} : \mathcal{P} \rightarrow \text{LType}$ and se a security environment.

Type system

LType is a poset (\leq is reflexive, antisymmetric, transitive) of local types.

Intuition of the type judgements: $se, i \vdash s \Rightarrow t$ means if executing program point i the type changes from s to t w.r.t a security environment se .

Definition (Typable program)

A program is typable (written $se, \mathcal{S} \vdash P$) if

- 1 for all initial program points holds $\mathcal{S}(i) = t_{init}$ and
- 2 $\forall i, j \in \mathcal{P} : (i \mapsto j) \rightarrow \exists s \in \text{LType} . se, i \vdash \mathcal{S}(i) \Rightarrow s \wedge \mathcal{S}(j) \leq s$

where $\mathcal{S} : \mathcal{P} \rightarrow \text{LType}$ and se a security environment.

Soundness of the type system

Definition (Noninterfering program)

\sim_g is a indistinguishability relation on global memories. A program is noninterfering iff for all global memories $\mu_1, \mu'_1, \mu_2, \mu'_2$ the following holds

$$(\mu_1 \sim_g \mu_2 \wedge P, \mu_1 \Downarrow \mu'_1 \wedge P, \mu_2 \Downarrow \mu'_2) \rightarrow \mu'_1 \sim_g \mu'_2$$

Theorem

If the scheduler is secure and $se, S \vdash P$, then P is noninterfering

Due to this theorem it is possible to typecheck the bytecode (which was compiled type-preserving) to proof the non-existence of internal timing leaks.

The proof is not part of this presentation, but I'll show the next function on which the proof relies.

Soundness of the type system

Definition (Noninterfering program)

\sim_g is a indistinguishability relation on global memories. A program is noninterfering iff for all global memories $\mu_1, \mu'_1, \mu_2, \mu'_2$ the following holds

$$(\mu_1 \sim_g \mu_2 \wedge P, \mu_1 \Downarrow \mu'_1 \wedge P, \mu_2 \Downarrow \mu'_2) \rightarrow \mu'_1 \sim_g \mu'_2$$

Theorem

If the scheduler is secure and $se, S \vdash P$, then P is noninterfering

Due to this theorem it is possible to typecheck the bytecode (which was compiled type-preserving) to proof the non-existence of internal timing leaks.

The proof is not part of this presentation, but I'll show the next function on which the proof relies.

Soundness of the type system

Definition (Noninterfering program)

\sim_g is a indistinguishability relation on global memories. A program is noninterfering iff for all global memories $\mu_1, \mu'_1, \mu_2, \mu'_2$ the following holds

$$(\mu_1 \sim_g \mu_2 \wedge P, \mu_1 \Downarrow \mu'_1 \wedge P, \mu_2 \Downarrow \mu'_2) \rightarrow \mu'_1 \sim_g \mu'_2$$

Theorem

If the scheduler is secure and $se, S \vdash P$, then P is noninterfering

Due to this theorem it is possible to typecheck the bytecode (which was compiled type-preserving) to proof the non-existence of internal timing leaks.

The proof is not part of this presentation, but I'll show the next function on which the proof relies.

The next function

If the execution of program point i results in a high thread, the function $\text{next} : \mathcal{P} \rightarrow \mathcal{P}$ calculates the program point in which the thread becomes visible again.

The next function has to fulfill the following properties:

$$\text{Dom}(\text{next}) = \{i \in \mathcal{P} \mid H(i) \wedge \neg AH(i)\} \quad (1)$$

$$i, j \in \text{Dom}(\text{next}) \wedge i \mapsto j \Rightarrow \text{next}(i) = \text{next}(j) \quad (2)$$

$$i \in \text{Dom}(\text{next}) \wedge L(j) \wedge i \mapsto j \Rightarrow \text{next}(i) = j \quad (3)$$

$$j, k \in \text{Dom}(\text{next}) \wedge L(i) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = \text{next}(k) \quad (4)$$

$$i, j \in \text{Dom}(\text{next}) \wedge L(k) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = k \quad (5)$$

The next function

If the execution of program point i results in a high thread, the function $\text{next} : \mathcal{P} \rightarrow \mathcal{P}$ calculates the program point in which the thread becomes visible again.

The next function has to fulfill the following properties:

$$\text{Dom}(\text{next}) = \{i \in \mathcal{P} \mid H(i) \wedge \neg AH(i)\} \quad (1)$$

$$i, j \in \text{Dom}(\text{next}) \wedge i \mapsto j \Rightarrow \text{next}(i) = \text{next}(j) \quad (2)$$

$$i \in \text{Dom}(\text{next}) \wedge L(j) \wedge i \mapsto j \Rightarrow \text{next}(i) = j \quad (3)$$

$$j, k \in \text{Dom}(\text{next}) \wedge L(i) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = \text{next}(k) \quad (4)$$

$$i, j \in \text{Dom}(\text{next}) \wedge L(k) \wedge i \mapsto j \wedge i \mapsto k \wedge j \neq k \Rightarrow \text{next}(j) = k \quad (5)$$

Source and target language

- Simple language with `if`, `;`, `:=`, `while` and `fork`
- Assembly
 - `push n` — push value on the stack
 - `load x` — push value of variable on the stack
 - `store x` — store first element of the stack in `x`
 - `goto j / ifeq j` — un-/conditional jump to `j`
 - `start j` — create a new thread starting in `j`

Source and target language

- Simple language with `if`, `;`, `:=`, `while` and `fork`
- Assembly
 - `push n` — push value on the stack
 - `load x` — push value of variable on the stack
 - `store x` — store first element of the stack in `x`
 - `goto j / ifeq j` — un-/conditional jump to `j`
 - `start j` — create a new thread starting in `j`

Transfer rules

LType = *Stack*(Level)

$$\frac{P[i] = \text{store } x \quad \text{se}(i) \sqcup k \leq \Gamma(x)}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow st}$$

$$\frac{P[i] = \text{ifeq } j \quad \forall j' \in \text{reg}(i), k \leq \text{se}(j')}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow \text{lift}_k(st)}$$

where $\text{reg} : \mathcal{P} \rightarrow \mathfrak{P}(\mathcal{P})$ computes the control dependence region. $\text{lift}_k(st)$ is the point-wise extension of $\lambda k'. k \sqcup k'$. $\Gamma(x)$ expresses the chosen security policy by assigning a security level to each variable.

Similar rules have to be established for the other commands of the target language.

Transfer rules

LType = *Stack*(Level)

$$\frac{P[i] = \text{store } x \quad \text{se}(i) \sqcup k \leq \Gamma(x)}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow st}$$

$$\frac{P[i] = \text{ifeq } j \quad \forall j' \in \text{reg}(i), k \leq \text{se}(j')}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow \text{lift}_k(st)}$$

where $\text{reg} : \mathcal{P} \rightarrow \mathfrak{P}(\mathcal{P})$ computes the control dependence region. $\text{lift}_k(st)$ is the point-wise extension of $\lambda k'.k \sqcup k'$. $\Gamma(x)$ expresses the chosen security policy by assigning a security level to each variable.

Similar rules have to be established for the other commands of the target language.

Transfer rules

LType = *Stack*(Level)

$$\frac{P[i] = \text{store } x \quad \text{se}(i) \sqcup k \leq \Gamma(x)}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow st}$$

$$\frac{P[i] = \text{ifeq } j \quad \forall j' \in \text{reg}(i), k \leq \text{se}(j')}{\text{se}, i \vdash_{\text{seq}} k :: st \Rightarrow \text{lift}_k(st)}$$

where $\text{reg} : \mathcal{P} \rightarrow \mathfrak{P}(\mathcal{P})$ computes the control dependence region. $\text{lift}_k(st)$ is the point-wise extension of $\lambda k'. k \sqcup k'$. $\Gamma(x)$ expresses the chosen security policy by assigning a security level to each variable.

Similar rules have to be established for the other commands of the target language.

Concurrent extension

The transfer rules are extended by the following rules:

$$\frac{P[i] \in \text{SeqIns} \quad se, i \vdash_{seq} s \Rightarrow t}{se, i \vdash s \Rightarrow t}$$

$$\frac{P[i] = \text{start } pc \quad se(i) \leq se(pc)}{se, i \vdash s \Rightarrow s}$$

We label the program points where control flow can branch or side effects can occur.

$$c ::= [x := e]^n \mid c; c \mid [if\ e\ then\ c\ else\ c']^n \mid [while\ e\ do\ c]^n \\ \mid [fork(c)]^n$$

With this labeling we can define control dependence regions for the source language (`sregion`) and derive them for the target language (`tregion`).

Concurrent extension

The transfer rules are extended by the following rules:

$$\frac{P[i] \in \text{SeqIns} \quad se, i \vdash_{seq} s \Rightarrow t}{se, i \vdash s \Rightarrow t}$$

$$\frac{P[i] = \text{start } pc \quad se(i) \leq se(pc)}{se, i \vdash s \Rightarrow s}$$

We label the program points where control flow can branch or side effects can occur.

$$c ::= [x := e]^n \mid c; c \mid [if\ e\ then\ c\ else\ c']^n \mid [while\ e\ do\ c]^n \\ \mid [fork(c)]^n$$

With this labeling we can define control dependence regions for the source language (`sregion`) and derive them for the target language (`tregion`).

sregion & tregion

Definition (sregion)

sregion(n) is defined as the set of labels that are inside a branching command $[c]^n$, except those inside `fork`.

Definition (tregion)

tregion(n) is defined for $[c]^n$ as the set of instructions/labels obtained by compiling $[c']^{n'}$ where $n' \in \text{sregion}(n)$. If c is `while` then $n \in \text{tregion}(n)$.

Excerpt of the compilation function C :

```
C(c) = let (lc, T) = S(c, []);
      in goto (#T + 2) :: T :: lc :: return
S(fork(c), T) = let (lc, T') = S(c, T);
              in (start (#T' + 2), T' :: lc :: return)
```

sregion & tregion

Definition (sregion)

$sregion(n)$ is defined as the set of labels that are inside a branching command $[c]^n$, except those inside `fork`.

Definition (tregion)

$tregion(n)$ is defined for $[c]^n$ as the set of instructions/labels obtained by compiling $[c']^{n'}$ where $n' \in sregion(n)$. If c is `while` then $n \in tregion(n)$.

Excerpt of the compilation function C :

```
C(c) = let (lc, T) = S(c, []);
      in goto (#T + 2) :: T :: lc :: return
S(fork(c), T) = let (lc, T') = S(c, T);
              in (start (#T' + 2), T' :: lc :: return)
```

sregion & tregion

Definition (sregion)

$sregion(n)$ is defined as the set of labels that are inside a branching command $[c]^n$, except those inside fork.

Definition (tregion)

$tregion(n)$ is defined for $[c]^n$ as the set of instructions/labels obtained by compiling $[c']^{n'}$ where $n' \in sregion(n)$. If c is while then $n \in tregion(n)$.

Excerpt of the compilation function C:

```
C(c) = let (lc, T) = S(c, []);
      in goto (#T + 2) :: T :: lc :: return
S(fork(c), T) = let (lc, T') = S(c, T);
               in (start (#T' + 2), T' :: lc :: return)
```

junction points & next function

Definition (junction point)

For every branching point $[c]^n$ in the source program we define

$$jun(n) = \max\{i \mid i \in tregion(n)\} + 1$$

To identify the outermost branching points that involves secrets we extend the type system. A source program is typeable ($\vdash_{\circ} c : E$ where E maps labels to security levels) and judgments of the form $\vdash_{\alpha} [c]_{\alpha'}^n : E$. One example typing rule (\circ public, \bullet secret):

$$\frac{\vdash e : H \quad \vdash_{\bullet} c : E \quad E = \text{lift}_H(E, sregion(n))}{\vdash_{\circ} [\text{while } e \text{ do } c]_{\bullet}^n : E}$$

Definition (next)

For alle branching program points c such that $\vdash_{\circ} [c]_{\bullet}^n$ *next* is defined as $\forall k \in tregion(n) . next(k) = jun(n)$.

junction points & next function

Definition (junction point)

For every branching point $[c]^n$ in the source program we define

$$jun(n) = \max\{i \mid i \in tregion(n)\} + 1$$

To identify the outermost branching points that involves secrets we extend the type system. A source program is typeable ($\vdash_{\circ} c : E$ where E maps labels to security levels) and judgments of the form $\vdash_{\alpha} [c]_{\alpha'}^n : E$. One example typing rule (\circ public, \bullet secret):

$$\frac{\vdash e : H \quad \vdash_{\bullet} c : E \quad E = \text{lift}_H(E, sregion(n))}{\vdash_{\circ} [\text{while } e \text{ do } c]_{\bullet}^n : E}$$

Definition (next)

For alle branching program points c such that $\vdash_{\circ} [c]_{\bullet}^n$ *next* is defined as $\forall k \in tregion(n) . next(k) = jun(n)$.

junction points & next function

Definition (junction point)

For every branching point $[c]^n$ in the source program we define

$$jun(n) = \max\{i \mid i \in tregion(n)\} + 1$$

To identify the outermost branching points that involves secrets we extend the type system. A source program is typeable ($\vdash_{\circ} c : E$ where E maps labels to security levels) and judgments of the form $\vdash_{\alpha} [c]_{\alpha'}^n : E$. One example typing rule (\circ public, \bullet secret):

$$\frac{\vdash e : H \quad \vdash_{\bullet} c : E \quad E = \text{lift}_H(E, sregion(n))}{\vdash_{\circ} [\text{while } e \text{ do } c]_{\bullet}^n : E}$$

Definition (next)

For alle branching program points c such that $\vdash_{\circ} [c]_{\bullet}^n$ next is defined as $\forall k \in tregion(n) . next(k) = jun(n)$.

Other/related solutions

- Protection/hiding based approaches
 - Volpano & Smith [4][5][3] use a `protect(c)` primitive
 - Russo & Sabelfeld [2] use `hide` and `unhide` primitives
- Low-determinism approaches
 - Zdancewic and Myres [6] disallow races on public data
- External-timing based approaches
 - here the attacker is more powerful: he can measure execution time
 - this causes much more restrictiveness (e.g. loops with secret guards are disallowed)

Comparison with Zdancewi and Myres[6]

- Introduces a relative complex language λ_{SEC}^{PAR}
- Also uses a type system to enforce security
- Uses the same notion of noninterference
- Observational determinism is defined as the indistinguishability of memory access traces

$$(m \approx_{\zeta} m' \wedge m \Downarrow T \wedge m' \Downarrow T') \Rightarrow T \approx_{\zeta} T'$$

Thus it rejects Programs like $!o := 1 \parallel !o := 0$

- In contrast to the paper discussed here, λ_{SEC}^{PAR} provides support for synchronization using *join patterns*

Adaption to the JVM

- JVML's sequential type system is compatible with bytecode verification, thus it's compatible with the concurrent type system.
- The scheduler is mostly left unspecified, thus introducing a secure scheduler is possible.
- Issues
 - Method calls have a big-step semantic
 - This approach does not deal with synchronization

Adaption to the JVM

- JVML's sequential type system is compatible with bytecode verification, thus it's compatible with the concurrent type system.
- The scheduler is mostly left unspecified, thus introducing a secure scheduler is possible.
- Issues
 - Method calls have a big-step semantic
 - This approach does not deal with synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Conclusion

- Proof of noninterference for a concurrent low-level language
- Proof of type-preserving compilation in context of concurrency
- Scheduler is driven by the security environment
- Independent of the scheduling algorithm
- No useful secure programs are rejected
- No need to trust the compiler, checking can be done at target level (without running the program)
- Programmer does not need to know about internal timing leaks
- No restrictions on dynamic thread creation
- What needs to be done? Extension for real world languages e.g. adding support for synchronization

Bibliography I

- [1] Gilles Barthe, Tamara Rezk, Alejandro Russo, and Andrei Sabelfeld.
Security of multithreaded programs by compilation.
In In Proc. 12th European Symposium on Research in Computer Security, pages 2–18. Springer-Verlag, 2007.
- [2] Alejandro Russo and Andrei Sabelfeld.
Securing interaction between threads and the scheduler.
In IEEE Computer Security Foundations Symposium, pages 177–189, 2006.
- [3] G. Smith and D. Volpano.
A sound type system for secure flow analysis.
In J. Computer Security 4, pages 167–187, 1996.

Bibliography II

- [4] G. Smith and D. Volpano.
Secure information flow in a multi-threaded imperative language.
In ACM Symp. on Principles of Programming Languages, pages 355–364, 1998.
- [5] G. Smith and D. Volpano.
Probalistic noninterference in a concurrent language.
In J. Computer Security 7, pages 231–253, 1999.
- [6] Steve Zdancewic and Andrew C. Myers.
Observational determinism for concurrent program security.
In In Proc. 16th IEEE Computer Security Foundations Workshop, pages 29–43, 2003.