

S³MS

Security of Software and
Services for Mobile Systems

Security-by-Contract: Toward a Semantics for Digital Signatures on Mobile Code*

N. Dragoni, F. Massacci, K. Naliuka and I. Siahaan
Department of Information and Communication Technologies
University of Trento, ITALY
dragoni@dit.unitn.it

EUROPKI'07 - Mallorca, Balearic Islands, Spain

* Research partly supported by the project EU-IST-STREP-S3MS (<http://www.s3ms.org>)

Talk Outline

- ⊙ Introduction
- ⊙ Security-By-Contract Framework
 - ⊙ Stakeholders
 - ⊙ Contract and Policy
 - ⊙ Application/Service Life-Cycle
- ⊙ Contract-Policy Matching
 - ⊙ Problem
 - ⊙ Algorithm

Lack of Applications for Mobile Devices

- Mobile devices are increasingly popular and powerful
- Yet, the growth in computing power of nomadic devices has not been supported by a comparable growth in available software
- For instance, on high-end mobile phones we cannot even remotely find the amount of *third party software* that was available on our old PC

A Reason: Security Model for Mobile Code

- One of the reasons for this lack of applications is also the security model adopted for mobile phones.
- The current security model approach (for instance, the JAVA MIDP 2.0) is based on **trust relationships**: *mobile code is run if its origin is trusted*.
- This essentially boils down to *mobile code is accepted if it is digitally signed by a trusted party*.
- The level of trust of the “trusted party” determines the privileges of the code by essentially segregating it into an appropriate trust domain.

Trust Relationship Problem (1)

The problem with trust relationship, i.e. digital signatures on mobile code, is twofold:

1. At first we can only **reject or accept the signature**.

This means that interoperability in a domain is either total or not existing: an application from a not-so-trusted source can be denied network access, but it cannot be denied access to a specific protocol, or to a specific domain.

- E.g. if a payment service is available on a platform and an application for paying parking meters is loaded, the user cannot block the application from performing large payments.

2. The second (and major) problem, is that **there is no semantics attached to the signature**. This is a problem for both code producers and consumers.

Trust Relationship Problem (2)

- From the point of view of the **code consumers**, they must essentially **accept the code “as-is”** without the possibility of making informed decisions.
 - One might well trust SuperGame Inc. to provide excellent games and yet might decide to rule out games that keep playing while the battery falls below 20%. At present such choice is not possible.
- From the point of view of the **code producers**, they produce code with **unbounded liability**. They cannot declare *which* security actions the code will **do**. By signing the code they essentially declare that they did it.

Trust Relationship Problem (2)

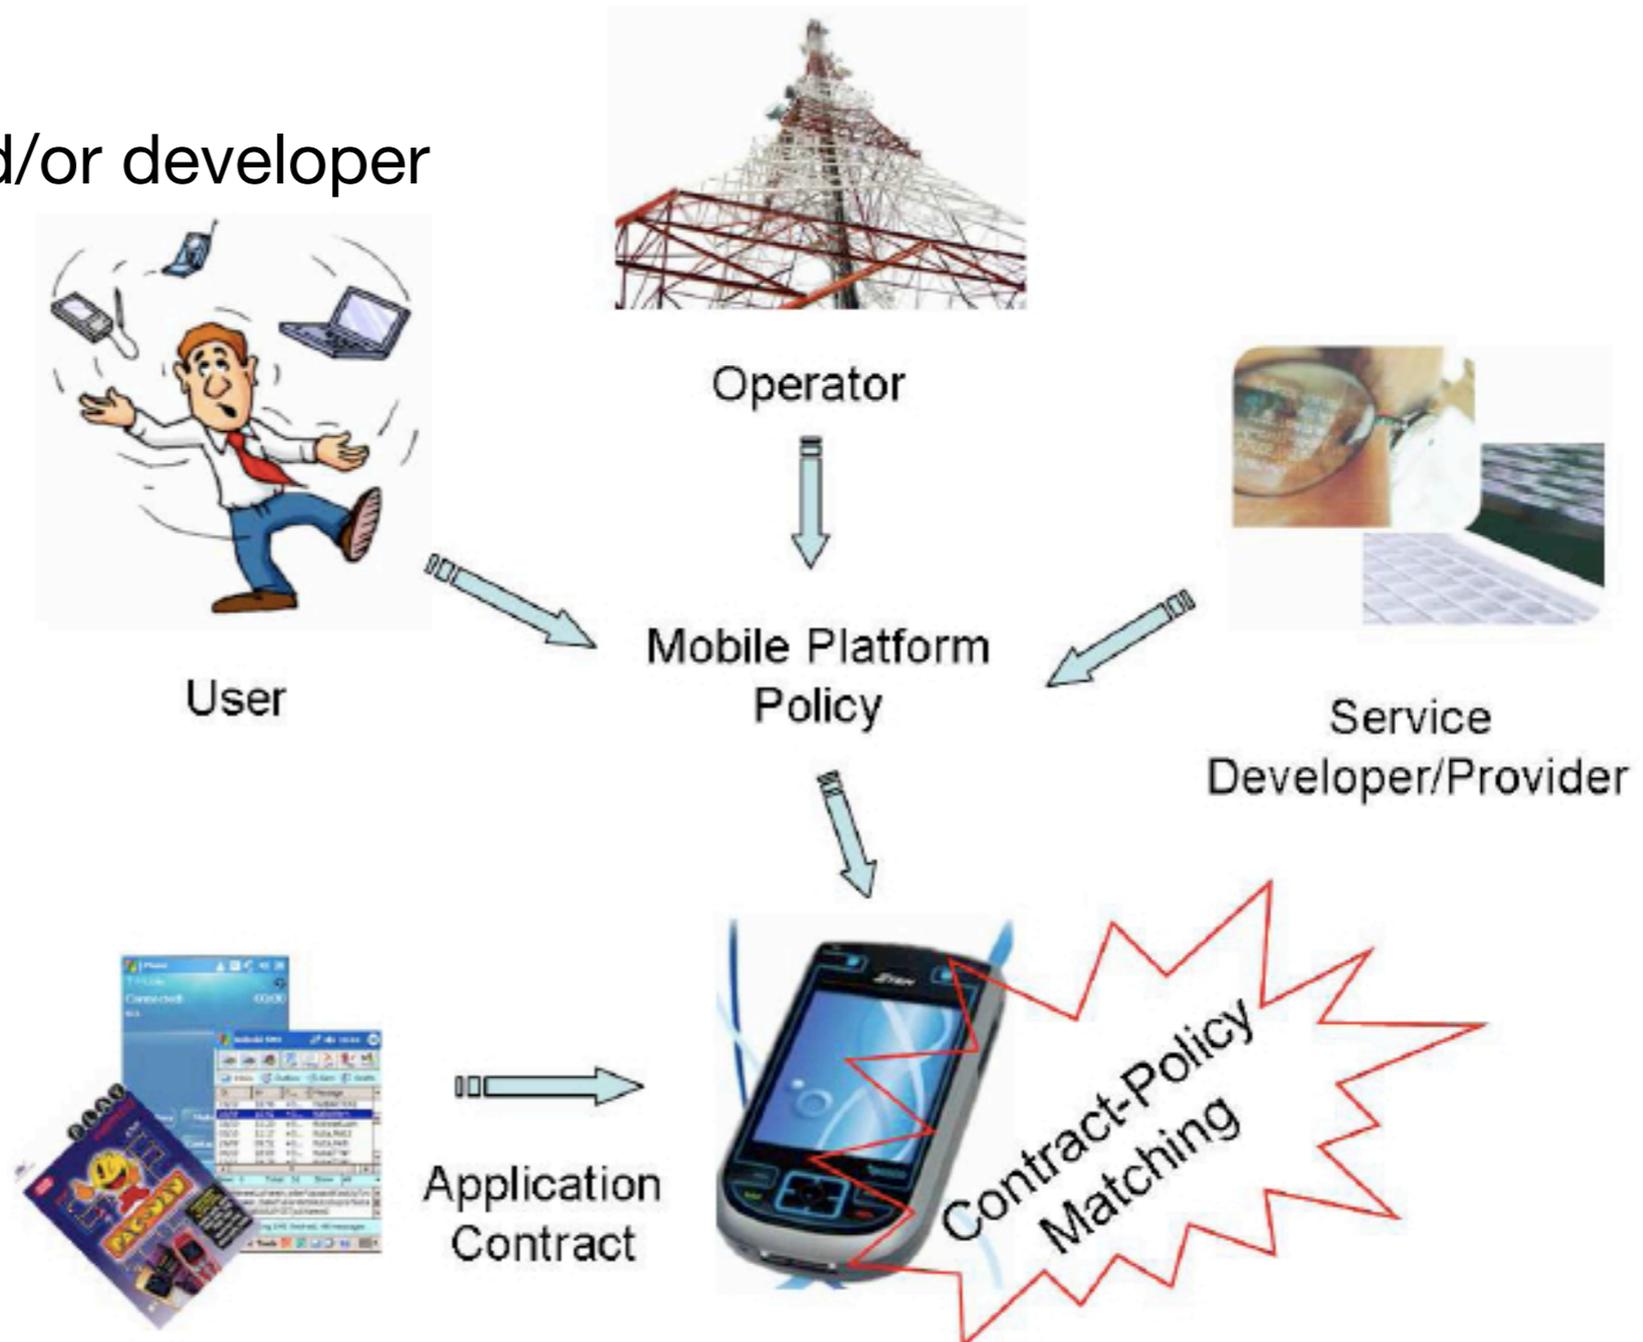
- From the point of view of the **code consumers**, they must essentially **accept the code “as-is”** without the possibility of making informed decisions.
 - One might well trust SuperGame Inc. to provide excellent games and yet might decide to rule out games that keep playing while the battery falls below 20%. At present such choice is not possible.
- From the point of view of the **code producers**, they produce code with **unbounded liability**. They cannot declare *which* security actions the code will **do**. By signing the code they essentially declare that they did it.

The consequence is that **injecting an application in the mobile market is a time consuming operation as developers must essentially convince the operators that their code will not do anything harmful.**

SxC Framework: Stakeholders

The Security-by-Contract framework is essentially shaped by three groups of stakeholders:

1. mobile operator
2. service provider and/or developer
3. mobile user



SxC Framework: Contract

The mobile code developers are responsible to provide a description of the **security behavior** that their code provides.



SxC Framework: Contract

The mobile code developers are responsible to provide a description of the **security behavior** that their code provides.

Contract: a contract is a formal complete and correct specification of the behavior of an application for what concerns relevant security actions (Virtual Machine API Calls, Operating System Calls).

What's in a code's contract?

- ➔ security features of application
- ➔ (security) interactions with its host platform
- ➔ proof-of-compliance that code satisfies contract



SxC Framework: Contract

The mobile code developers are responsible to provide a description of the **security behavior** that their code provides.

Contract: a contract is a formal complete and correct specification of the behavior of an application for what concerns relevant security actions (Virtual Machine API Calls, Operating System Calls).

What's in a code's contract?

- ➔ security features of application
- ➔ (security) interactions with its host platform
- ➔ proof-of-compliance that code satisfies contract



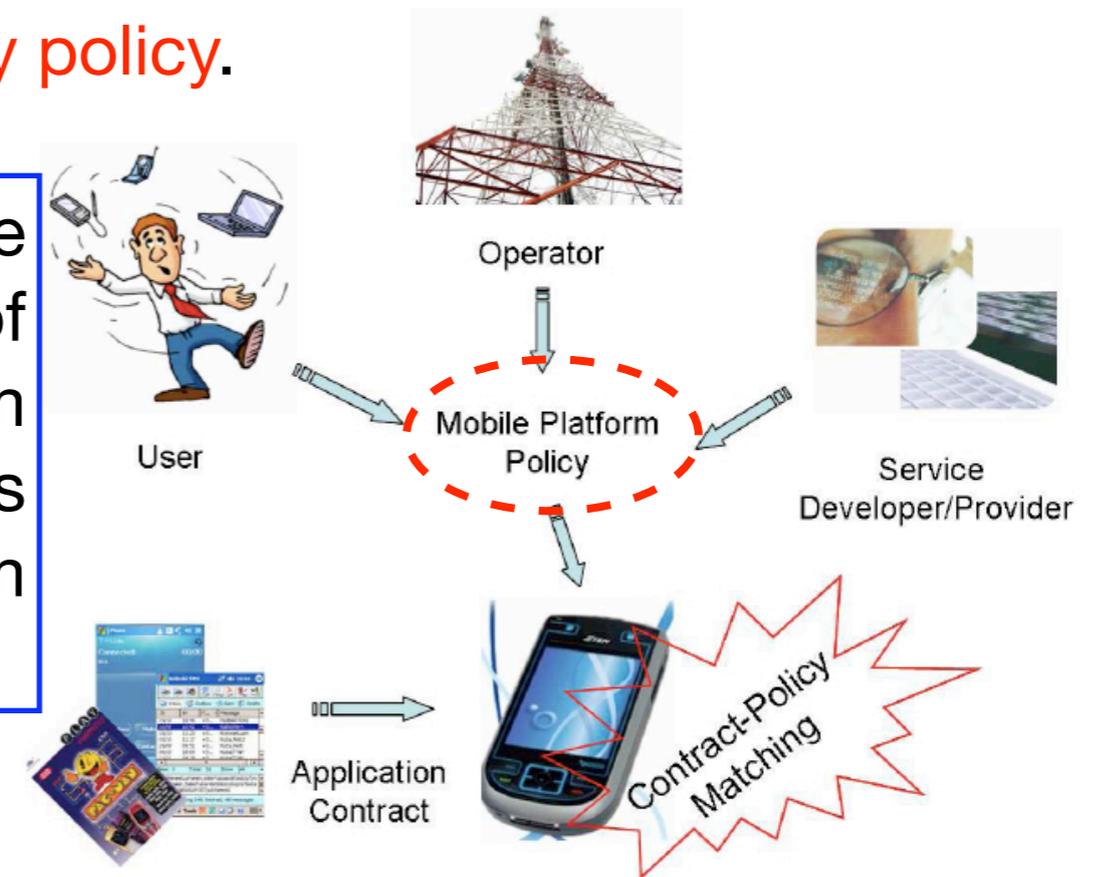
By signing the code the developer certifies that the code complies with the stated claims on its security-relevant behavior.

SxC Framework: Policy

On the other side we can see that users and mobile phone operators are interested that all codes that are deployed on their platform are secure.

In other words they must declare their **security policy**.

Policy: a policy is a formal complete specification of the *acceptable* behavior of applications to be executed on the platform for what concerns relevant security actions (Virtual Machine API Calls, Operating System Calls).



What's in a platform's policy?

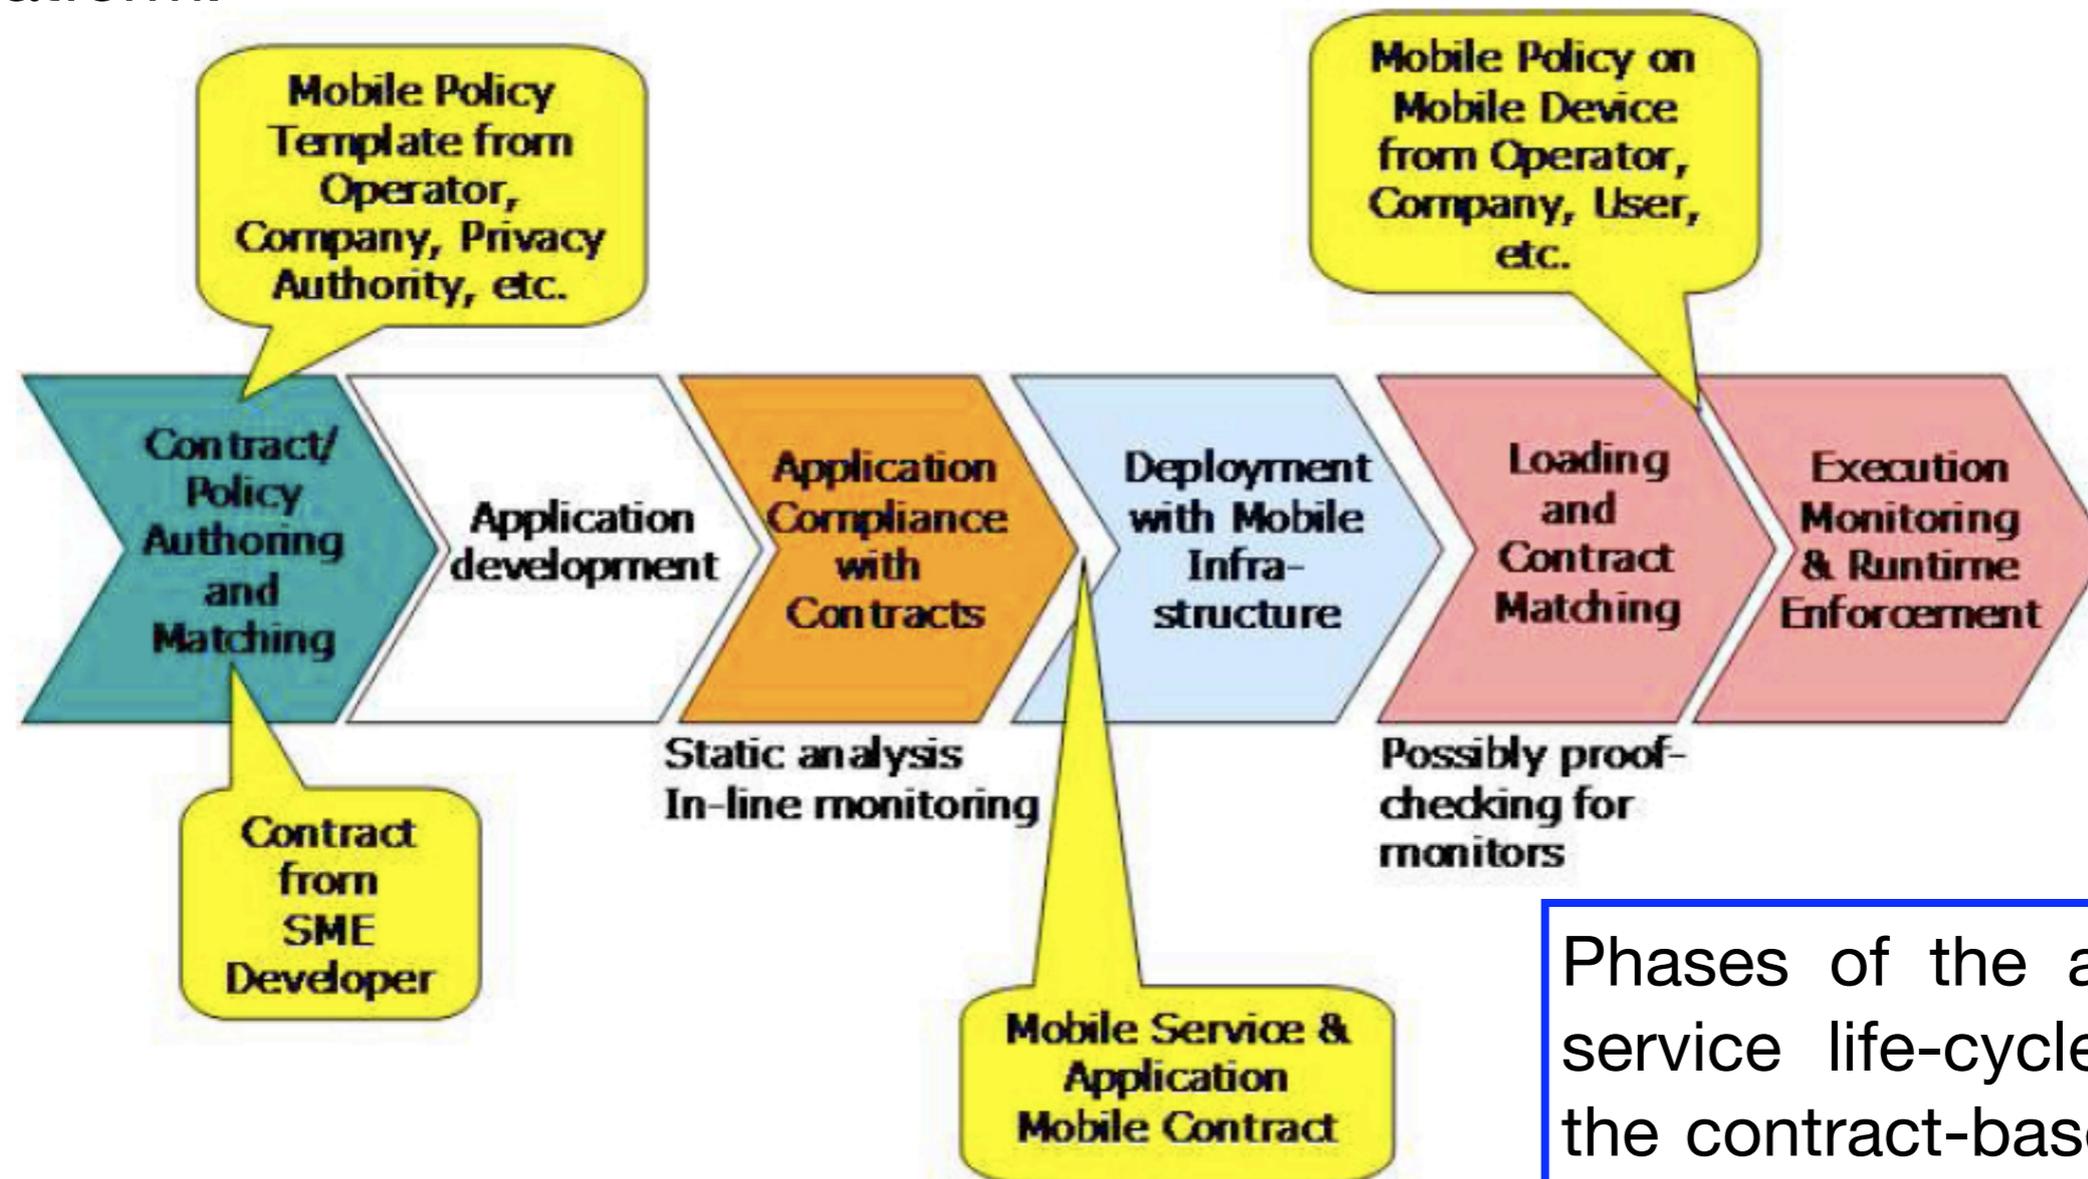
- ➔ platform contractual requirements on application
- ➔ fine-grained resource control (e.g. silently initiate a phone call or send a SMS)
- ➔ memory usage, secure and insecure web connections, user privacy protection

How a Contract/Policy Should Look Like?

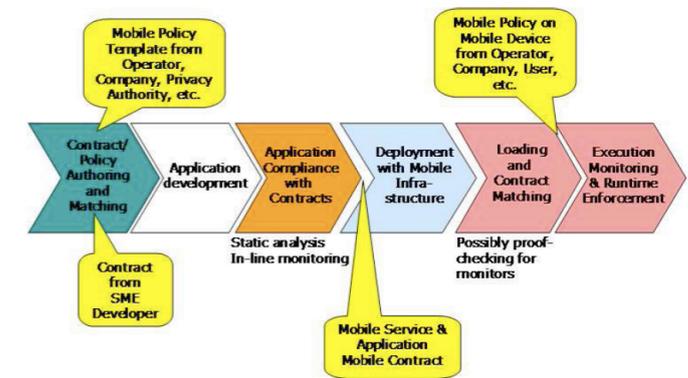
1. The application sends no more than a number messages in each session
2. The application only loads each image from the network once
3. The delay between two periodic invocations of the MIDlet is at least T
4. The application does not initiate calls to international numbers
5. The application only uses files whose name matches a given pattern
6. The application does not send MMS messages
7. The application connects only to its origin domain
8. The application must close all files that it opens
9. The application only receives SMS messages on a specific port
- ...

Application/Service Life Cycle (1)

A contract should be negotiated and enforced during development, at time of delivery and loading, and during execution of the application by the mobile platform.



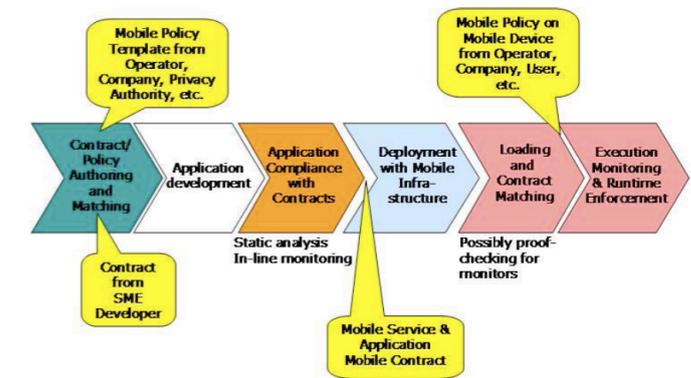
Phases of the application/service life-cycle in which the contract-based security paradigm is present.



Application/Service Life Cycle (2)

In order to guarantee that an application complies with its desired contract (or the policy requested on a particular platform) we should consider the stage where such enforcement can be done.

Development	Deployment		Execution
(I) At design and development time	(II) After design but before shipping the application	(III) When downloading the application	(IV) During the execution of the application

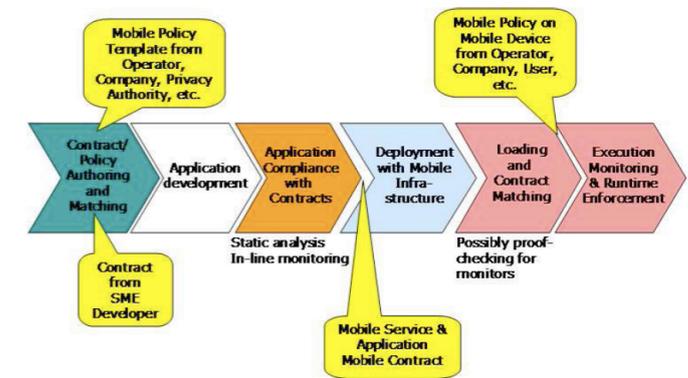


Application/Service Life Cycle (2)

In order to guarantee that an application complies with its desired contract (or the policy requested on a particular platform) we should consider the stage where such enforcement can be done.

Development	Deployment		Execution
(I) At design and development time	(II) After design but before shipping the application	(III) When downloading the application	(IV) During the execution of the application

- ▶ Enforcing at level (I) can be achieved by appropriate design rules and require developer support

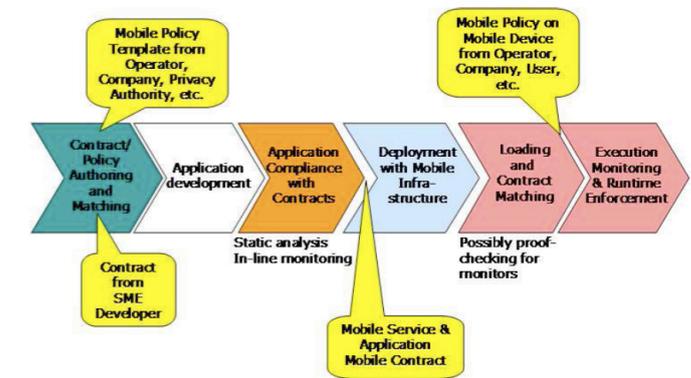


Application/Service Life Cycle (2)

In order to guarantee that an application complies with its desired contract (or the policy requested on a particular platform) we should consider the stage where such enforcement can be done.

Development	Deployment		Execution
(I) At design and development time	(II) After design but before shipping the application	(III) When downloading the application	(IV) During the execution of the application

- ▶ Enforcing at level (I) can be achieved by appropriate design rules and require developer support
- ▶ (II) and (III) can be carried out through **(automatic) verification techniques**. Such verifications can take place
 - ▶ before downloading (**static verification** by developers and operators followed by a contract coming with a *trusted signature*) or
 - ▶ as a combination of pre and post-loading operations (e.g., through **in-line monitors** and **proof carrying code**)



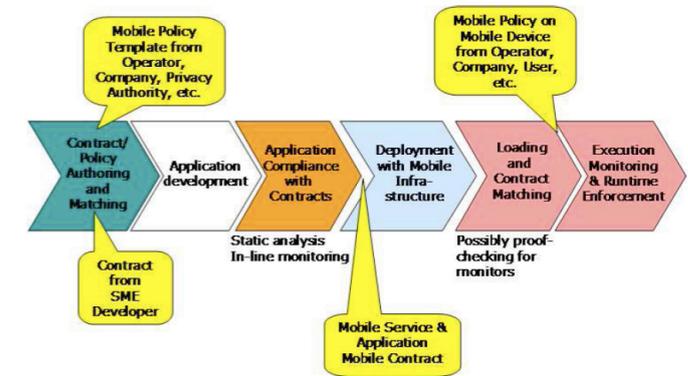
Application/Service Life Cycle (2)

In order to guarantee that an application complies with its desired contract (or the policy requested on a particular platform) we should consider the stage where such enforcement can be done.

Development	Deployment		Execution
(I) At design and development time	(II) After design but before shipping the application	(III) When downloading the application	(IV) During the execution of the application

- ▶ Enforcing at level (I) can be achieved by appropriate design rules and require developer support
- ▶ (II) and (III) can be carried out through **(automatic) verification techniques**. Such verifications can take place
 - ▶ before downloading (**static verification** by developers and operators followed by a contract coming with a *trusted signature*) or
 - ▶ as a combination of pre and post-loading operations (e.g., through **in-line monitors** and **proof carrying code**)
- ▶ (IV) can be implemented by **run-time checking**

Contract/Policy Matching



One of the key problems in the overall security-by-contract life-cycle is the **contract-policy matching issue**.

Given:

- ▶ *a contract that an application carries with itself*
 - ▶ *a policy that a platform specifies*
- is the contract compliant with the policy?*



- Intuitively, matching should succeed if and only if by executing the application on the platform every behaviour of the application that satisfies its contract also satisfies the platform's policy.
- Contract-policy matching represents a common problem in the life-cycle because **it must be done at all levels**: both for development and run-time operation.

Contract Specification

A single contract/policy is specified as a list of *disjoint* rules instead of one giant specification describing all possible security properties. A rule is defined according to the following grammar:

```
<RULE> :=  
    SCOPE [ OBJECT <class> |  
            SESSION |  
            MULTISESSION ]  
    RULEID <identifier>  
    <formal specification>
```

Contract Specification

A single contract/policy is specified as a list of *disjoint* rules instead of one giant specification describing all possible security properties. A rule is defined according to the following grammar:

<RULE> :=

SCOPE [**OBJECT** <class> |
 SESSION |
 MULTISESSION]

RULEID <identifier>
<formal specification>

Scope definition reflects at which scope the specified contract will be applied.

OBJECT: the obligation must be fulfilled by each object of a given type.

SESSION: the obligation must be fulfilled by each run of the application separately.

MULTISESSION: the obligation must be fulfilled by all runs of the application as a whole.

Contract Specification

A single contract/policy is specified as a list of *disjoint* rules instead of one giant specification describing all possible security properties. A rule is defined according to the following grammar:

```
<RULE> :=  
    SCOPE [ OBJECT <class> |  
            SESSION |  
            MULTISESSION ]  
    RULEID <identifier>  
    <formal specification>
```

The tag RULEID identifies the area of the contract (which security-relevant actions the policy concerns, for example “files” or “connections”).

Contract Specification

A single contract/policy is specified as a list of *disjoint* rules instead of one giant specification describing all possible security properties. A rule is defined according to the following grammar:

```
<RULE> :=  
    SCOPE [ OBJECT <class> |  
            SESSION |  
            MULTISESSION ]  
    RULEID <identifier>  
    <formal specification>
```

The tag RULEID identifies the area of the contract (which security-relevant actions the policy concerns, for example “files” or “connections”).

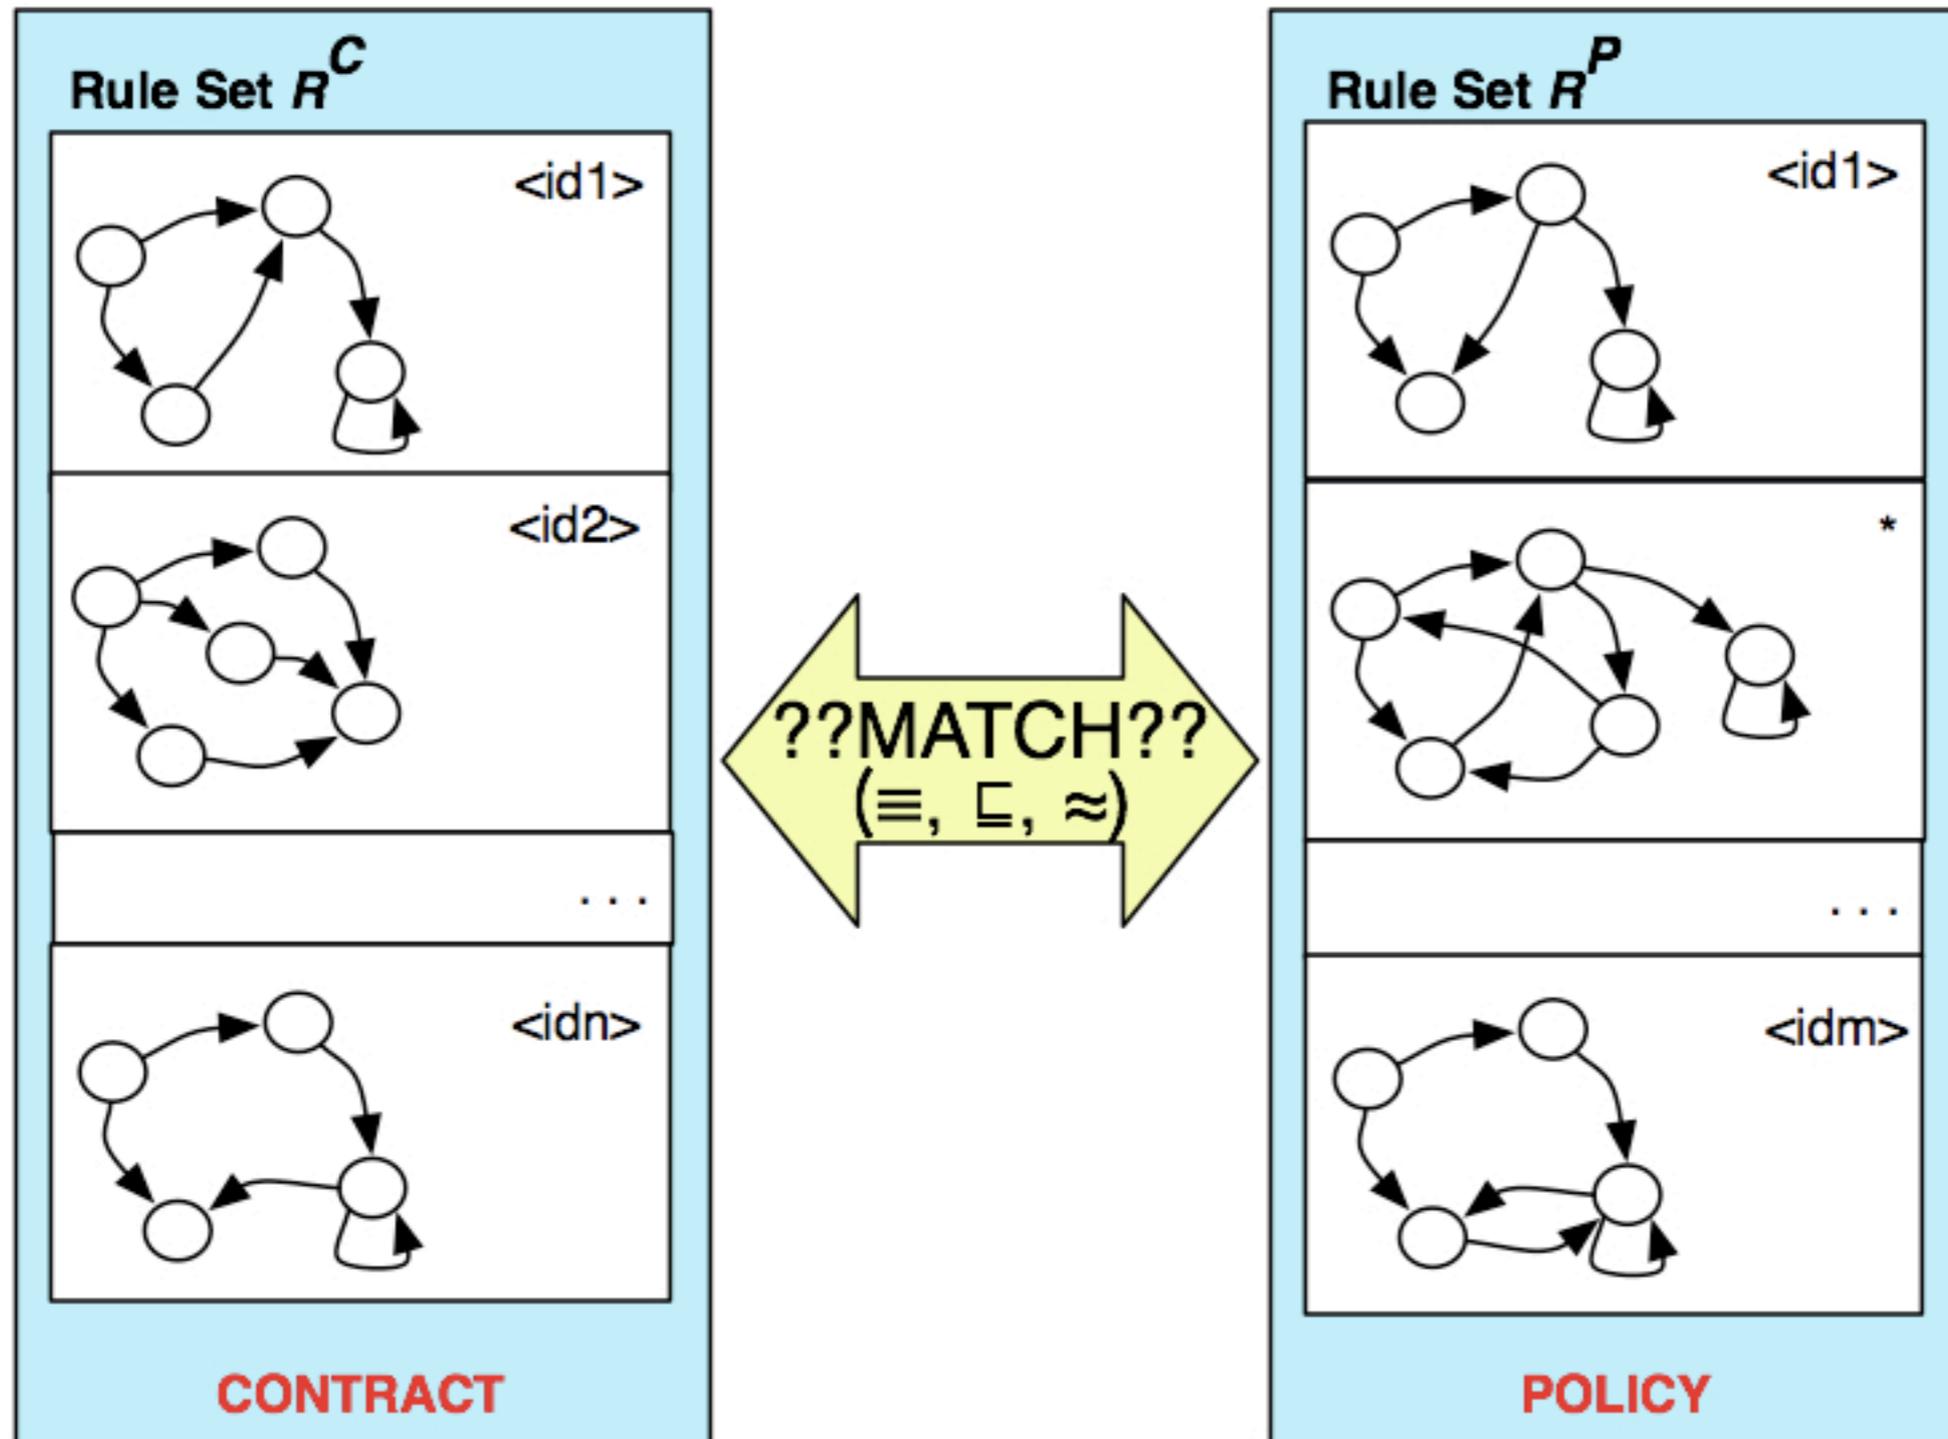
The <formal specification> part of a rule gives a rigorous and not ambiguous definition of the semantics of the rule.

Abstract Constructs

We have identified the following **abstract operators** (C and P indicate a generic contract and policy respectively):

- [Combine Operator \oplus] $\text{Spec} = \oplus_{i=1, \dots, n} \text{Spec}_i$
It combines all the rule formal specifications $\text{Spec}_1, \dots, \text{Spec}_n$ in a new specification Spec .
- [Simulate Operator \approx] $\text{Spec}^C \approx \text{Spec}^P$
It returns 1 if rule formal specification Spec^C simulates rule formal specification Spec^P , 0 otherwise.
- [Contained-By Operator \sqsubseteq] $\text{Spec}^C \sqsubseteq \text{Spec}^P$
It returns 1 if the behavior specified by Spec^C is among the behaviors that are allowed by Spec^P , 0 otherwise.

Contract-Policy Matching Problem



Matching Algorithm (1)

Algorithm 1. MatchContracts Function

Input: rule set \mathcal{R}^C , rule set \mathcal{R}^P

Output: 1 if \mathcal{R}^C matches \mathcal{R}^P , 0 otherwise

```

1:  $\langle \mathcal{R}_{SESSION}^C, \mathcal{R}_{MULTISESSION}^C, \{\mathcal{R}_{class}^C\}_{class \in \zeta^C} \rangle \Leftarrow \text{Partition}(\mathcal{R}^C)$ 
2:  $\langle \mathcal{R}_{SESSION}^P, \mathcal{R}_{MULTISESSION}^P, \{\mathcal{R}_{class}^P\}_{class \in \zeta^P} \rangle \Leftarrow \text{Partition}(\mathcal{R}^P)$ 
3: if MatchRules( $\mathcal{R}_{SESSION}^C, \mathcal{R}_{SESSION}^P$ ) then
4:   if MatchRules( $\mathcal{R}_{MULTISESSION}^C, \mathcal{R}_{MULTISESSION}^P$ ) then
5:     for all  $class \in \zeta^P$  do // for all classes in policy
6:       if MatchRules( $\mathcal{R}_{class}^C, \mathcal{R}_{class}^P$ ) then // if  $class \notin \zeta^C$ , then  $\mathcal{R}_{class}^C = \emptyset$ 
7:         skip
8:       else
9:         return(0)
10:      end if
11:    end for
12:    return(1)
13:  end if
14: end if
15: return(0)

```

Matching Algorithm (1)

Algorithm 1. MatchContracts Function

Input: rule set \mathcal{R}^C , rule set \mathcal{R}^P

Output: 1 if \mathcal{R}^C matches \mathcal{R}^P , 0 otherwise

```

1:  $\langle \mathcal{R}_{SESSION}^C, \mathcal{R}_{MULTISESSION}^C, \{\mathcal{R}_{class}^C\}_{class \in \zeta^C} \rangle \Leftarrow \text{Partition}(\mathcal{R}^C)$ 
2:  $\langle \mathcal{R}_{SESSION}^P, \mathcal{R}_{MULTISESSION}^P, \{\mathcal{R}_{class}^P\}_{class \in \zeta^P} \rangle \Leftarrow \text{Partition}(\mathcal{R}^P)$ 
3: if MatchRules( $\mathcal{R}_{SESSION}^C, \mathcal{R}_{SESSION}^P$ ) then
4:   if MatchRules( $\mathcal{R}_{MULTISESSION}^C, \mathcal{R}_{MULTISESSION}^P$ ) then
5:     for all  $class \in \zeta^P$  do // for all classes in policy
6:       if MatchRules( $\mathcal{R}_{class}^C, \mathcal{R}_{class}^P$ ) then // if class  $\notin \zeta^C$ , then  $\mathcal{R}_{class}^C = \emptyset$ 
7:         skip
8:       else
9:         return(0)
10:      end if
11:     end for
12:     return(1)
13:   end if
14: end if
15: return(0)

```

Algorithm 2. Partition Procedure

Input: rule set \mathcal{R}

Output: $\langle \mathcal{R}_{SESSION}, \mathcal{R}_{MULTISESSION}, \{\mathcal{R}_{class}\}_{class \in \zeta} \rangle$

```

1:  $\mathcal{R}_{SESSION} \Leftarrow \{r \in \mathcal{R} \mid \text{Scope}(r) = \text{SESSION}\}$ 
2:  $\mathcal{R}_{MULTISESSION} \Leftarrow \{r \in \mathcal{R} \mid \text{Scope}(r) = \text{MULTISESSION}\}$ 
3: for all  $class \in \zeta$  do // for all classes in contract/policy
4:    $\mathcal{R}_{class} \Leftarrow \{r \in \mathcal{R} \mid \text{Scope}(r) = \text{OBJECT} \langle class \rangle\}$ 
5: end for

```

Matching Algorithm (2)

Algorithm 3. MatchRules Function

Input: rule set \mathcal{R}^C , rule set \mathcal{R}^P

Output: 1 if \mathcal{R}^C matches \mathcal{R}^P , 0 otherwise

```

1:  $L^C \Leftarrow \{ (ID^C, Spec^C) \mid \langle \text{scope}, ID^C, Spec^C \rangle \in \mathcal{R}^C \}$ 
2:  $L^P \Leftarrow \{ (ID^P, Spec^P) \mid \langle \text{scope}, ID^P, Spec^P \rangle \in \mathcal{R}^P \}$ 
3: for all  $(ID^P, Spec^P) \in L^P$  do
4:   if MatchSpec( $L^C, (ID^P, Spec^P)$ ) then
5:     skip
6:   else // may return  $\emptyset$  for efficiency
7:      $L_{failed}^P \Leftarrow L_{failed}^P \cup (ID^P, Spec^P)$ 
8:   end if
9: end for
10: if  $L_{failed}^P = \emptyset$  then
11:   return(1)
12: else
13:   return( MatchSpec ( (  $*$ ,  $\bigoplus_{(ID^C, Spec^C) \in L^C}$  ), (  $*$ ,  $\bigoplus_{(ID^P, Spec^P) \in L_{failed}^P}$  ) ) ) )
14: end if

```

Matching Algorithm (3)

Algorithm 4. MatchSpec Function

Input: $L^C = \langle (ID_1^C, Spec_1^C), \dots, (ID_n^C, Spec_n^C) \rangle, (ID^P, Spec^P)$

Output: 1 if L^C matches $(ID^P, Spec^P)$, 0 otherwise

```

1: if  $\exists (ID^C, Spec^C) \in L^C \wedge ID^C = ID^P$  then
2:   if  $HASH(Spec^C) = HASH(Spec^P)$  then
3:     return(1)
4:   else if  $Spec^C \approx Spec^P$  then
5:     return(1)
6:   else if  $Spec^C \sqsubseteq Spec^P$  then
7:     return(1)
8:   else // Restriction: if same ID then same specification must match
9:     return(0)
10:  end if
11: else
12:   MatchSpec $\left(\left(*, \bigoplus_{(ID^C, Spec^C) \in L^C}\right), (*, Spec^P)\right)$ 
13: end if

```

Matching as Language Inclusion

Algorithm 5. FSA-Based Implementation of \sqsubseteq

Input: two FSA-based rule specifications Aut^C and Aut^P

Output: 1 if $A_C \sqsubseteq A_P$, 0 otherwise

1: Complement the automaton Aut^P

2: Construct the automaton Aut^I that accepts $\mathcal{L}_{\text{Aut}^C} \cap \overline{\mathcal{L}_{\text{Aut}^P}}$

3: **if** Aut^I is empty **then**

4: return(1)

5: **else**

6: return(0)

7: **end if**

Matching as Language Inclusion

Algorithm 5. FSA-Based Implementation of \sqsubseteq

Input: two FSA-based rule specifications Aut^C and Aut^P

Output: 1 if $A_C \sqsubseteq A_P$, 0 otherwise

- 1: Complement the automaton Aut^P
- 2: Construct the automaton Aut^I that accepts $\mathcal{L}_{\text{Aut}^C} \cap \overline{\mathcal{L}_{\text{Aut}^P}}$
- 3: **if** Aut^I is empty **then**
- 4: return(1)
- 5: **else**
- 6: return(0)
- 7: **end if**

on-the-fly

Conclusion

The contributions of the paper are threefold:

1. First, we have proposed the **security-by-contract framework** providing a description of the overall life-cycle of mobile code in this setting.
2. Then we have described a tentative structure for a contractual language.
3. Finally, we have proposed a number of algorithms for one of the key steps in the life-cycle process: the issue of contract-policy matching.

The main novelty of the proposed framework is that it would provide a semantics for digital signatures on mobile code thus being a step in the transition from trusted code to trustworthy code.

Conclusion

The contributions of the paper are threefold:

1. First, we have proposed the **security-by-contract framework** providing a description of the overall life-cycle of mobile code in this setting.
2. Then we have described a tentative structure for a contractual language.
3. Finally, we have proposed a number of algorithms for one of the key steps in the life-cycle process: the issue of contract-policy matching.

The main novelty of the proposed framework is that it would provide a semantics for digital signatures on mobile code thus being a step in the transition from trusted code to trustworthy code.

THANKS!

Appendix A - Definitions

1. **Exact Matching:** matching should succeed if and only if by executing the application on the platform every trace that satisfies the application's contract also satisfies the platform's policy:

$$\text{Traces} \left(\bigoplus_{i=1, \dots, n} \text{Spec}_i^C \right) \subseteq \text{Traces} \left(\bigoplus_{i=1, \dots, m} \text{Spec}_i^P \right)$$

2. **Sound Sufficient Matching:** matching should fail if by executing the application on the platform there might be an application trace that satisfies the contract and does not satisfy the policy.
3. **Complete Matching:** matching should succeed if by executing the application on the platform every traces satisfying the contract also satisfy the policy.

By applying Def. 2 we might reject “good” applications that are however too difficult or too complex to perform. On the other hand, Def. 3 may allow “bad” applications to run but it will certainly accept all “good” ones (and “bad” applications can later be detected, for instance, by run-time monitoring).

Appendix B - Abstract Constructs

We have identified the following **abstract operators** (C and P indicate a generic contract and policy respectively):

- [Combine Operator \oplus] $\text{Spec} = \oplus_{i=1, \dots, n} \text{Spec}_i$
It combines all the rule formal specifications $\text{Spec}_1, \dots, \text{Spec}_n$ in a new specification Spec .
- [Simulate Operator \approx] $\text{Spec}^C \approx \text{Spec}^P$
It returns 1 if rule formal specification Spec^C simulates rule formal specification Spec^P , 0 otherwise.
- [Contained-By Operator \sqsubseteq] $\text{Spec}^C \sqsubseteq \text{Spec}^P$
It returns 1 if the behavior specified by Spec^C is among the behaviors that are allowed by Spec^P , 0 otherwise.
- [Traces Operator] $\mathcal{S} = \text{Traces}(\text{Spec})$
It returns the set \mathcal{S} of all the possible sequences of actions that can be performed according to the formal specification Spec .