

An operational architecture for privacy-by-design


Prashant Agrawal*

Department of Computer Science and Engineering, IIT Delhi

*Joint work with Anubhuti Singh, Malavika Raghavan, Subodh Sharma and Subhashis Banerjee


Impossibility of absolute privacy

Absolute privacy goal (aka inferential privacy): A should not obtain any information about an individual that B cannot obtain without access to DB



Impossibility of absolute privacy



Absolute privacy goal (aka inferential privacy): A should not obtain any information about an individual that B cannot obtain without access to DB



If the adversary has arbitrary side-information, above absolute privacy goal is impossible to achieve. ([Dwork '05](#))

Impossibility of absolute privacy



Absolute privacy goal (aka inferential privacy): A should not obtain any information about an individual that B cannot obtain without access to DB



If the adversary has arbitrary side-information, above absolute privacy goal is impossible to achieve. (Dwork '05)

Impossibility of absolute privacy

Absolute privacy goal (aka inferential privacy): A should not obtain any information about an individual that B cannot obtain without access to DB




If the adversary has arbitrary side-information, above absolute privacy goal is impossible to achieve. (Dwork '05)

Observe: Privacy of director's salary is compromised even if the director is not in the DB

So now what?

- **“Privacy is dead, get over it”**
- **“Burn everything down”**
- **Differential privacy:**
 - Make sure that the privacy protection is as good (or as bad) as the case when you (specifically) did not even choose to participate in the database
 - Rather narrow view of the privacy harms
 - Mainly a statistical notion, to allow privacy-preserving analytics and machine learning on personal information
- **Puttaswamy judgment:**
 - The proportionality test for balancing utility and privacy

An economic perspective: Who really cares about privacy?



- Us, the individuals
- Not the corporations, not the government
- They care about the utility of our data
- Privacy will always take the back seat, especially if it conflicts with utility
- Individuals by themselves are often powerless, naive and ignorant

The false notion of consent

- Consent is broken, as evidenced by the customary “*I Agree*”
- Consent can be overridden
- Unfamiliarity with legal rights, technology
- Inability to envisage or judge potential harms of digitisation use cases, both to self and society
- Unfamiliarity with privacy management tools

What does the court say?

The proportionality test (Puttaswamy I and II)

- Must be sanctioned by law
- Must be necessary in a democratic society for a legitimate state aim
- Extent of interference must be proportionate to the aim
 - Rational nexus with the objective
 - Least intrusive for the purpose
 - Must not have disproportionate impact (balancing)
- There must be procedural guarantees against abuse from such interference

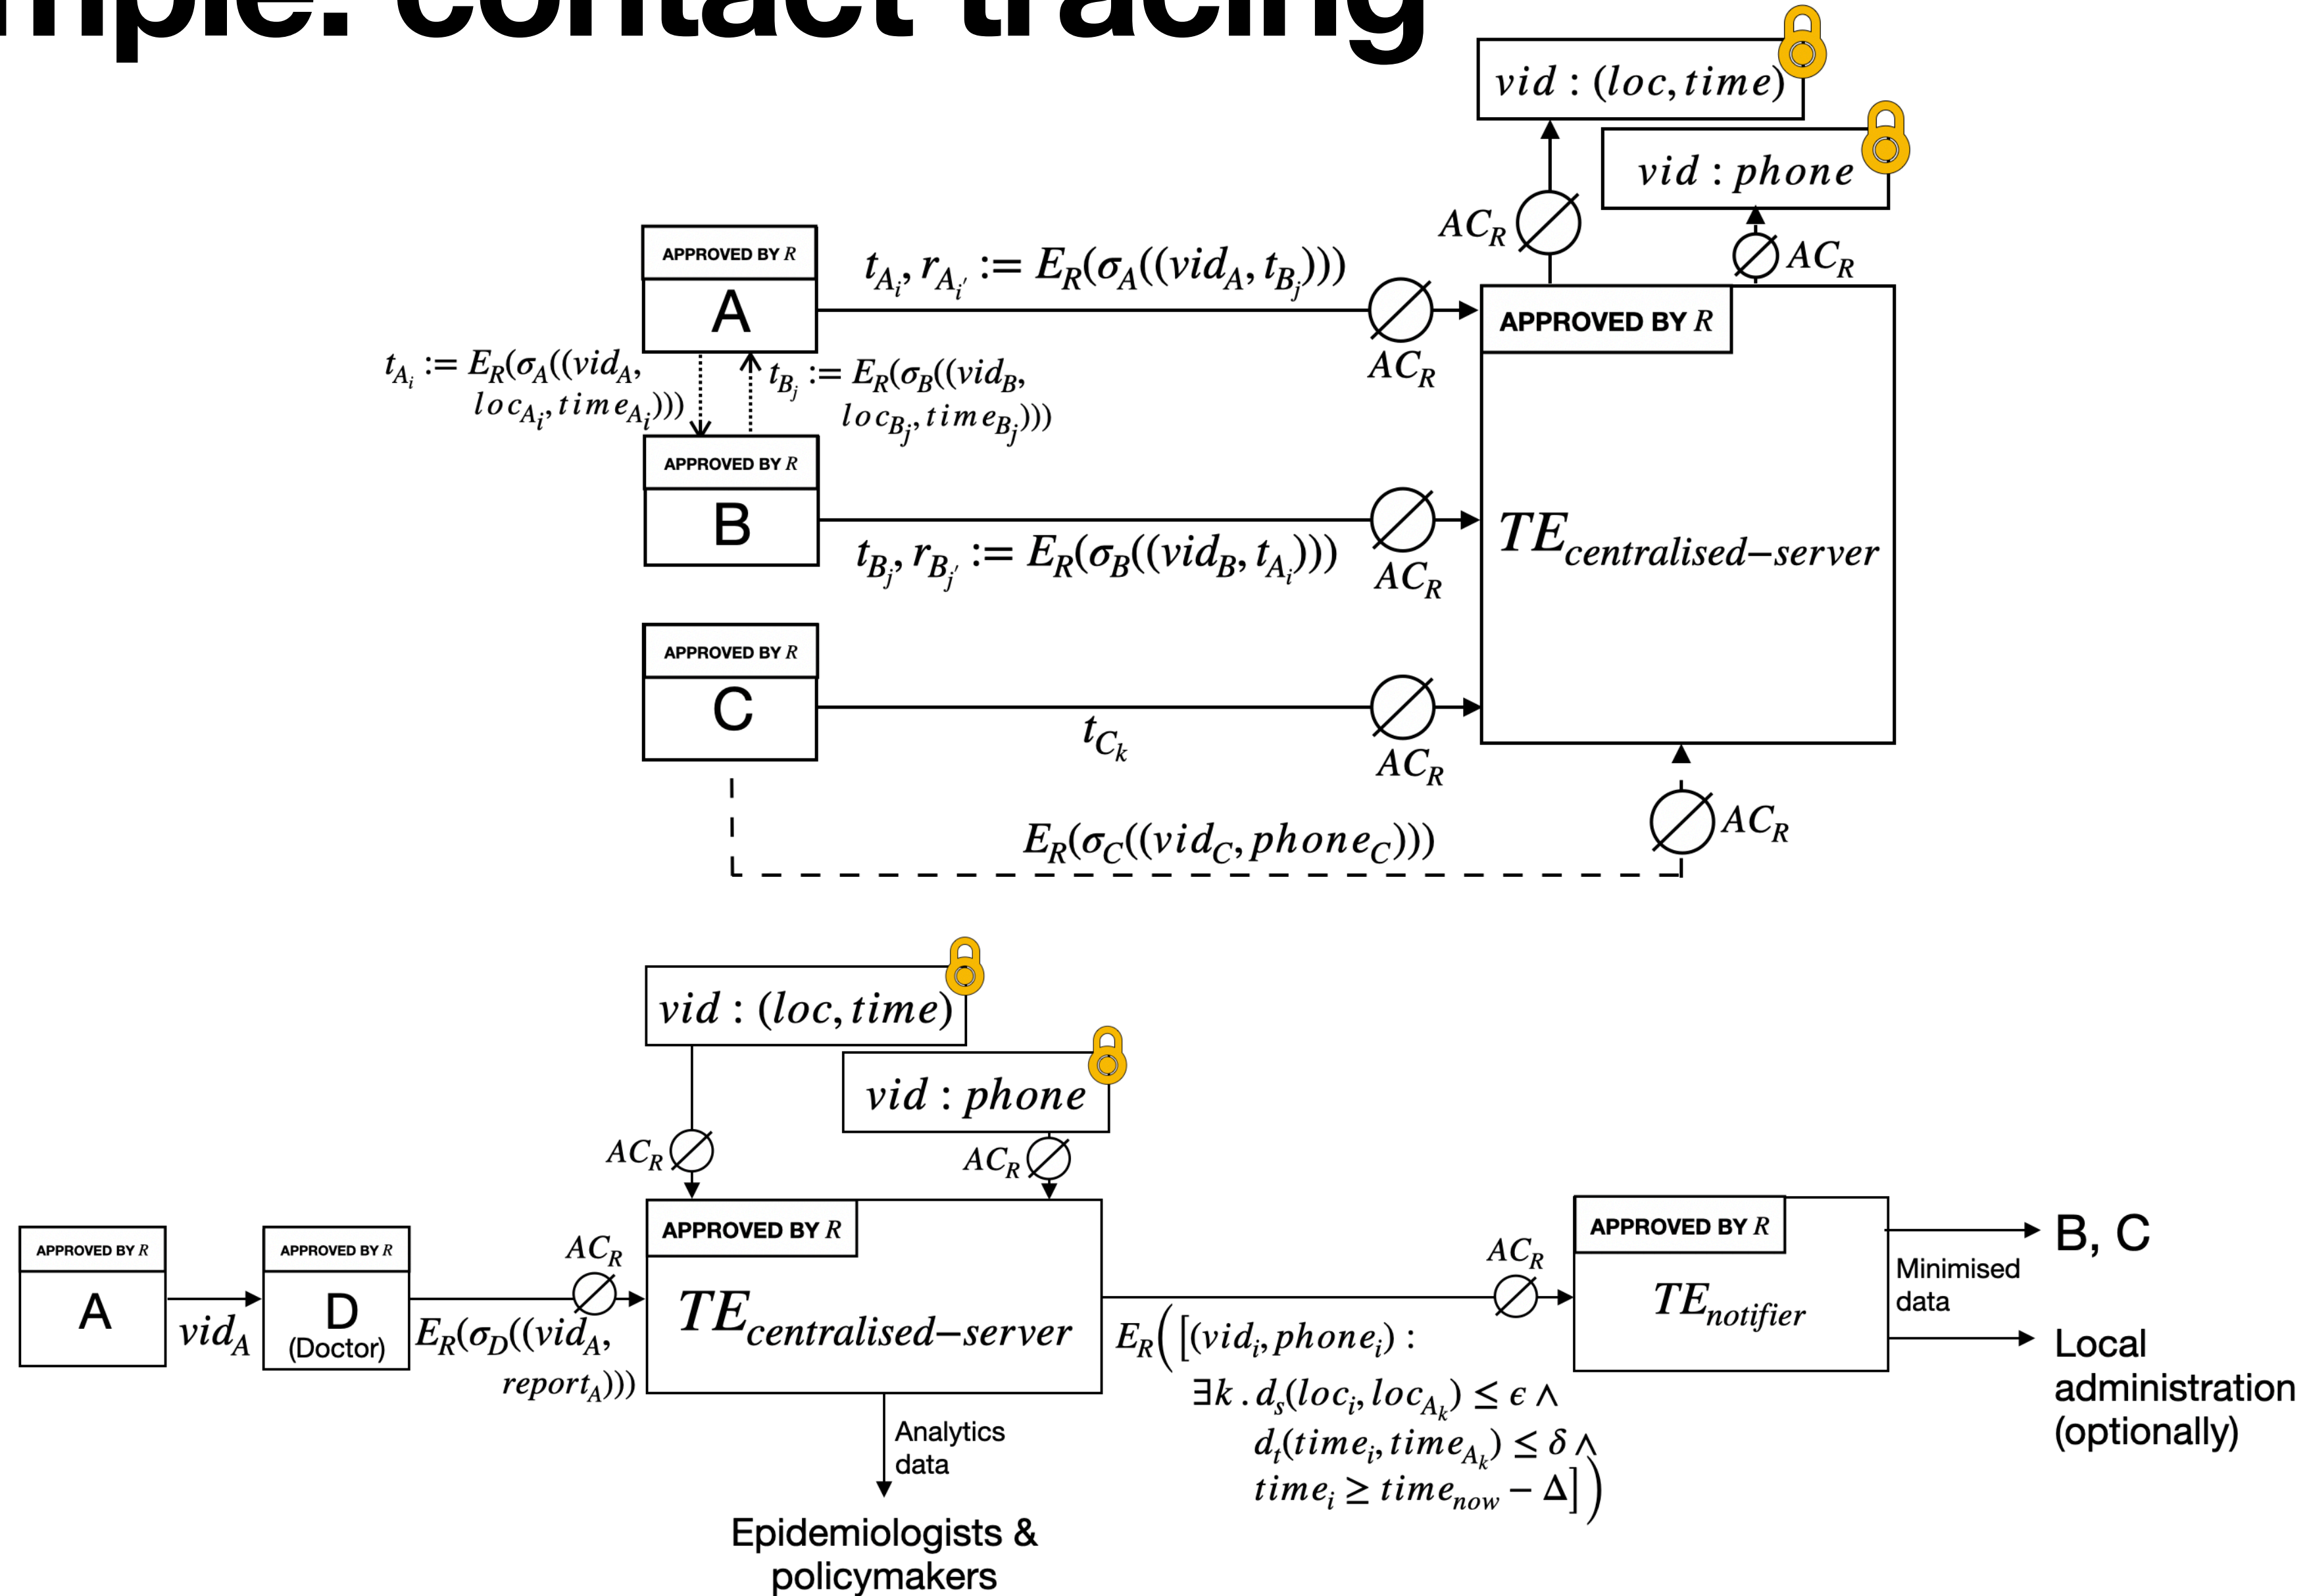
Resolving the tension between privacy, utility and usability

- Offload the responsibility of privacy management from individuals to regulatory authorities
- But keep the regulator accountable (Trust but verify)

Regulator's responsibilities

- **Identify a privacy policy** considering utility and privacy goals (should be backed by appropriate law). The policy should be able to express:
 - data minimisation goals (outputting only minimum information, preventing linking attacks)
 - dynamic and parametric access control (access control based on personalised context, revocation, etc.)
 - purpose limitation
- **Enforce compliance** to the privacy policy: Nothing except what is allowed by the policy should ever leak to anyone, not even an insider
- **Demonstrate to the general public** that the privacy has been upheld as per the policy and will always be upheld.

Example: contact tracing



Example policy: contact tracing

- Individuals should be able to learn only about their own infection.
- Local authorities should be able to learn about the contact details of only high risk individuals.
- Only individuals with legitimately high risk should be classified as such.
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism).
- Individuals must give their consent and they can opt out of the service at any point of time.
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals.
- Only individuals with legitimately high risk should be classified as such.
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism).
- Individuals must give their consent and they can opt out of the service at any point of time.
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such.
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism).
- Individuals must give their consent and they can opt out of the service at any point of time.
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism).
- Individuals must give their consent and they can opt out of the service at any point of time.
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism). Purpose limitation, access control
- Individuals must give their consent and they can opt out of the service at any point of time.
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism). Purpose limitation, access control
- Individuals must give their consent and they can opt out of the service at any point of time. Consent and revocation of consent (conditional and dynamic access control)
- A doctor to which the individual wilfully visited should be able to fetch her medical data.
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism). Purpose limitation, access control
- Individuals must give their consent and they can opt out of the service at any point of time. Consent and revocation of consent (conditional and dynamic access control)
- A doctor to which the individual wilfully visited should be able to fetch her medical data. Access dependent on relationships between individuals
- No one should be able to obtain any additional information even with insider access.
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.

Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism). Purpose limitation, access control
- Individuals must give their consent and they can opt out of the service at any point of time. Consent and revocation of consent (conditional and dynamic access control)
- A doctor to which the individual wilfully visited should be able to fetch her medical data. Access dependent on relationships between individuals
- No one should be able to obtain any additional information even with insider access. Prevention of insider attacks
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data.


Example policy: contact tracing

- Individuals should be able to learn only about their own infection. Individual-level access control
- Local authorities should be able to learn about the contact details of only high risk individuals. Purpose limitation, access control
- Only individuals with legitimately high risk should be classified as such. Correctness (indirectly affects privacy; Kafkaesque)
- Epidemiologists should be able to learn about aggregate information only (maybe only via a differentially private mechanism). Purpose limitation, access control
- Individuals must give their consent and they can opt out of the service at any point of time. Consent and revocation of consent (conditional and dynamic access control)
- A doctor to which the individual wilfully visited should be able to fetch her medical data. Access dependent on relationships between individuals
- No one should be able to obtain any additional information even with insider access. Prevention of insider attacks
- No persistent identifiers should ever leak else they could be used to arbitrarily link individuals' data. Prevention of linking attacks

**Let's start by looking at our
current identity infrastructure**


Traditional PKI and the problem with it

- Encryption




Traditional PKI and the problem with it

- Encryption




- Signatures




Traditional PKI and the problem with it


- Encryption



- Signatures




- Credentials




Traditional PKI and the problem with it


- Encryption



- Signatures




- Credentials




Traditional PKI and the problem with it

- Encryption




- Signatures




Traditional PKI is designed to encrypt messages for public organisations and obtain signatures from public organisations

- Credentials



Virtual identities: an individual-centric notion


- Each individual a owns a master secret key sk_a
- Individual a can generate multiple unlinkable virtual identities using sk_a



- *Notation:* a_i denotes the i -th ever generated virtual identity by agent a
- *PKI is a special case:* All agents only ever generate a single virtual identity and use it everywhere. Thus the public key of agent a is the only virtual identity a_0 generated by a , and its secret key would be its master secret key sk_a .


Requirements from virtual identities

- Anonymous credentials




Requirements from virtual identities

- Anonymous credentials




- Anonymous signatures




Requirements from virtual identities


- Anonymous credentials



- Anonymous signatures




- Anonymous encryptions




Requirements from virtual identities

- Anonymous credentials




Envisioned by Chaum '85,
practically realised by
Camenisch & Lysyanskaya '01
(but partial solutions only)

- Anonymous signatures




- Anonymous encryptions




Requirements from virtual identities

- Anonymous credentials




Envisioned by Chaum '85,
practically realised by
Camenisch & Lysyanskaya '01
(but partial solutions only)

- Anonymous signatures



- Anonymous encryptions



Most likely open
problems

**For now, let's assume we have a solution
for virtual identities and move forward**

Access control

- Access control should be parametric with respect to each individual
 - “don’t release x ’s data if x has not given her consent.”
 - relationships among individuals: “ x can access y ’s data only if x is a family member of y ”
- Access control should be context-dependant
 - “access to x ’s data may only be allowed if a warrant against x can be produced.”
 - “administrators are allowed to access one’s contact only if an algorithm classifies them as high-risk.”
- Access control should be dynamic
 - what is allowed today may be revoked tomorrow due to revocation of consent, signs, etc.
 - “Public is Private”: just because something has been released to an agent once does not mean there is no need to prevent access to it in the future.

In contrast

Model	Sender	Recipient	Subject	Attributes	Past	Future	Combination
RBAC	Role	Identity	×	×	×	×	●
XACML	Flexible	Flexible	Flexible	○	×	○	●
EPAL	Fixed	Role	Fixed	●	×	○	×
P3P	Fixed	Role	Fixed	●	○	×	○
CI	Role	Role	Role	●	●	●	●

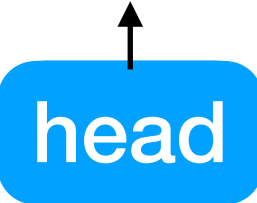
Figure 5. Comparison of various privacy languages. The symbol × indicates the feature is absent from the language, ○ indicates partial or limited functionality, and ● indicates the feature is fully functional. Note, [6] gives an extension of EPAL that is closed under combination.

([Barth et al. '06](#))

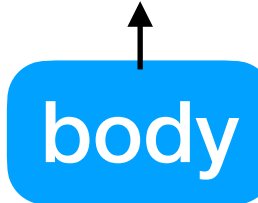
- Access control and privacy policy authoring languages are rather coarse:
 - Typically characterised by fixed roles, and often do not parametrise the data type with respect to the individuals
 - Very poorly handle dynamically changing context

A privacy policy language

- Logic programs (or, colloquially, Prolog programs):
 - Set of rules R of the form $p_0 \leftarrow p_1, p_2, \dots, p_n$, where each p_i is a first-order predicate containing variables (uppercase) and constants (lowercase):
 - e.g., say $R := \{ \text{allow}(z, X, \text{MedData}(Y)) \leftarrow \text{signed}(\text{consulted}(X), Y) \}$



↑
head



↑
body
 - If $n = 0$, i.e., a rule with only a head and no body, is called a *fact*, which is unconditionally true
 - All variables are implicitly universally quantified.

A privacy policy language (contd.)

- Given a logic program R and a (potentially first-order) predicate p , it can be answered in **polynomial time** whether p is implied by the rules R or not (we write this as $R \vdash p$)
- E.g., check if $R \vdash \text{allow}(z, \text{vid}(x_2), \text{MedData}(\text{vid}(y_3)))$?
- Key idea: **unification**. Find substitutions for variables that makes two terms syntactically identical.
 - e.g., $\text{unify}(p_1(X, y), p_1(x, Y)) := [X/x, Y/y]$

A privacy policy language (contd.)

- Basic access control procedure:
 - Check if $R \vdash \text{allow}(x, y, l)$. If yes, then allow x to send y data labelled with l .
 - Note: the rules in R have a head of the form $\text{allow}(X, Y, L)$, with all three potentially first-order variables
 - On unifying the queried allow predicate with the head, we get the concrete rule:
$$\text{allow}(z, \text{vid}(x_2), \text{MedData}(\text{vid}(y_3))) \leftarrow \text{signed}(\text{consulted}(\text{vid}(x_2)), \text{vid}(y_3))$$

Handling dynamic factors using signed predicates

- **Claim:** All reasonable dynamic factors can be expressed using a set of signed predicates of the form: $\text{signed}(P, A)$, denoting that A has expressed that predicate P is true.
 - *Consent/approvals:* $\text{signed}(\text{consent}(\text{vid}(b_3), \text{financeData}(\text{vid}(a_2))), \text{vid}(a_2))$
 - *Credentials by public authorities:* $\text{signed}(\text{passed}(\text{vid}(a_2)), \text{vid}(z_0))$. Transformable by a to be of the form $\text{signed}(\text{passed}(\text{vid}(a_3)), \text{vid}(z_0))$
 - *Credentials by private individuals:* $\text{signed}(\text{passed}(\text{vid}(a_2)), \text{vid}(b_3))$. Transformable by a to be of the form $\text{signed}(\text{passed}(\text{vid}(a_3)), \text{vid}(b_3))$.
 - *Machine-generated facts:* $\text{signed}(\text{currTime}(t), \text{timer})?$
 - *Revocation of previously granted access:* $\neg \text{signed}(\text{consented}(p), \text{via}(a_2))?$

Handling dynamic factors using signed predicates

- **Claim:** All reasonable dynamic factors can be expressed using a set of signed predicates of the form: $\text{signed}(P, A)$, denoting that A has expressed that predicate P is true.
 - *Consent/approvals:* $\text{signed}(\text{consent}(\text{vid}(b_3), \text{financeData}(\text{vid}(a_2))), \text{vid}(a_2))$
 - *Credentials by public authorities:* $\text{signed}(\text{passed}(\text{vid}(a_2)), \text{vid}(z_0))$. Transformable by a to be of the form $\text{signed}(\text{passed}(\text{vid}(a_3)), \text{vid}(z_0))$
 - *Credentials by private individuals:* $\text{signed}(\text{passed}(\text{vid}(a_2)), \text{vid}(b_3))$. Transformable by a to be of the form $\text{signed}(\text{passed}(\text{vid}(a_3)), \text{vid}(b_3))$.
 - *Machine-generated facts:* $\text{signed}(\text{currTime}(t), \text{timer})?$
 - *Revocation of previously granted access:* $\neg \text{signed}(\text{consented}(p), \text{via}(a_2))?$

Need to think about “identity”
and “signatures” of programs...

Need to think about contradictory
predicates in the logic program

Logic programs with exceptions

- Exceptions necessary to allow a rule override access given by the other rule
- In traditional logic, if you have a rule $\text{allow}(X, Y, L) \leftarrow B_1$ and a rule $\neg \text{allow}(X, Y, L) \leftarrow B_2$, then it leads to a *contradiction*.
- We need a kind of exception mechanism that gives priority to the negative rule.

$\text{allow}(X, Y, L) \leftarrow \text{not deny}(X, Y, L)$

not p is true if no rule of the form $p \leftarrow p_1, p_2, \dots, p_n$ exists (closed world assumption; different from $\neg p$)

Logic programs with exceptions


- Positive rules and exception rules, with exceptions taking priority:

Positive rules

$\text{allow}(\text{db}, Y, \text{MedData}(X)) \leftarrow \text{signed}(\text{visited}(Y), X)$

Exception rules

$\text{deny}(Z, Y, L) \leftarrow \text{contains}(L, X), \text{signed}(\text{revoked}(\text{consent}), X)$



Compiled to the following
logic program (to be
executed in Prolog)

$\text{allow}(\text{db}, Y, \text{MedData}(X)) \leftarrow \text{signed}(\text{visited}(Y), X), \text{not } \text{deny}(\text{db}, Y, \text{MedData}(X))$

$\text{deny}(Z, Y, L) \leftarrow \text{contains}(L, X), \text{signed}(\text{revoked}(\text{consent}), X)$

Purpose limitation

- At the time of data collection, purpose should be stated and it must be assured that purpose would not be violated.
- Reasoning about the *future*.
- Some preliminary formulations:
 - Purpose identified with the organisational role of the accessor ([Byun & Li '05](#))
 - Purpose identified with an action graph ([Jafari et al. '11](#))
 - Purpose identified with a “plan” ([Tschantz et al. '12](#))
- Common theme:
 - Either a poor proxy for purpose is chosen, or the enforcement mechanism is weak

Purpose limitation

- **Our notion:** Purpose identified by a “sandboxed program!”
- *Sandboxing:*
 - Program runs within a black-box.
 - No one can learn any intermediate execution information. Only official output is learnable.
 - No one can tamper with the execution of the program. Only official input can affect the execution.
 - *Cryptographic notions:* secure multiparty computation, functional encryption, ...
 - *System security notions:* hardware-based trusted execution environments, ...
- A sandboxed program can be assigned an identifier that defines its purpose!
- Can talk about signatures generated within the sandbox and messages encrypted for the sandbox.
 - E.g., `signed(currTime(t), timer)`

Some examples

$\text{allow}(\text{machine}(\text{EpidemicAnalysisProg}), \text{vid}(X), \text{hotspots}) \leftarrow \text{signed}(\text{isAnalyst}(\text{vid}(X)), \text{vid}(z_0))$

Sandboxing achieves data
minimisation, correctness of output!


$\text{allow}(\text{vid}(X), \text{machine}(\text{CancerAnalyzer}), \text{MedData}(X))$

Sandboxing achieves purpose
limitation, access control for writes

What does it mean to be compliant to this policy?

- Remember the *differential* notion in all the security definitions we have seen so far.
- Can we write a similar definition expressing that the data controller does not leak anything except what is allowed by this policy?
 - Yes, using ideas from secure multiparty computation

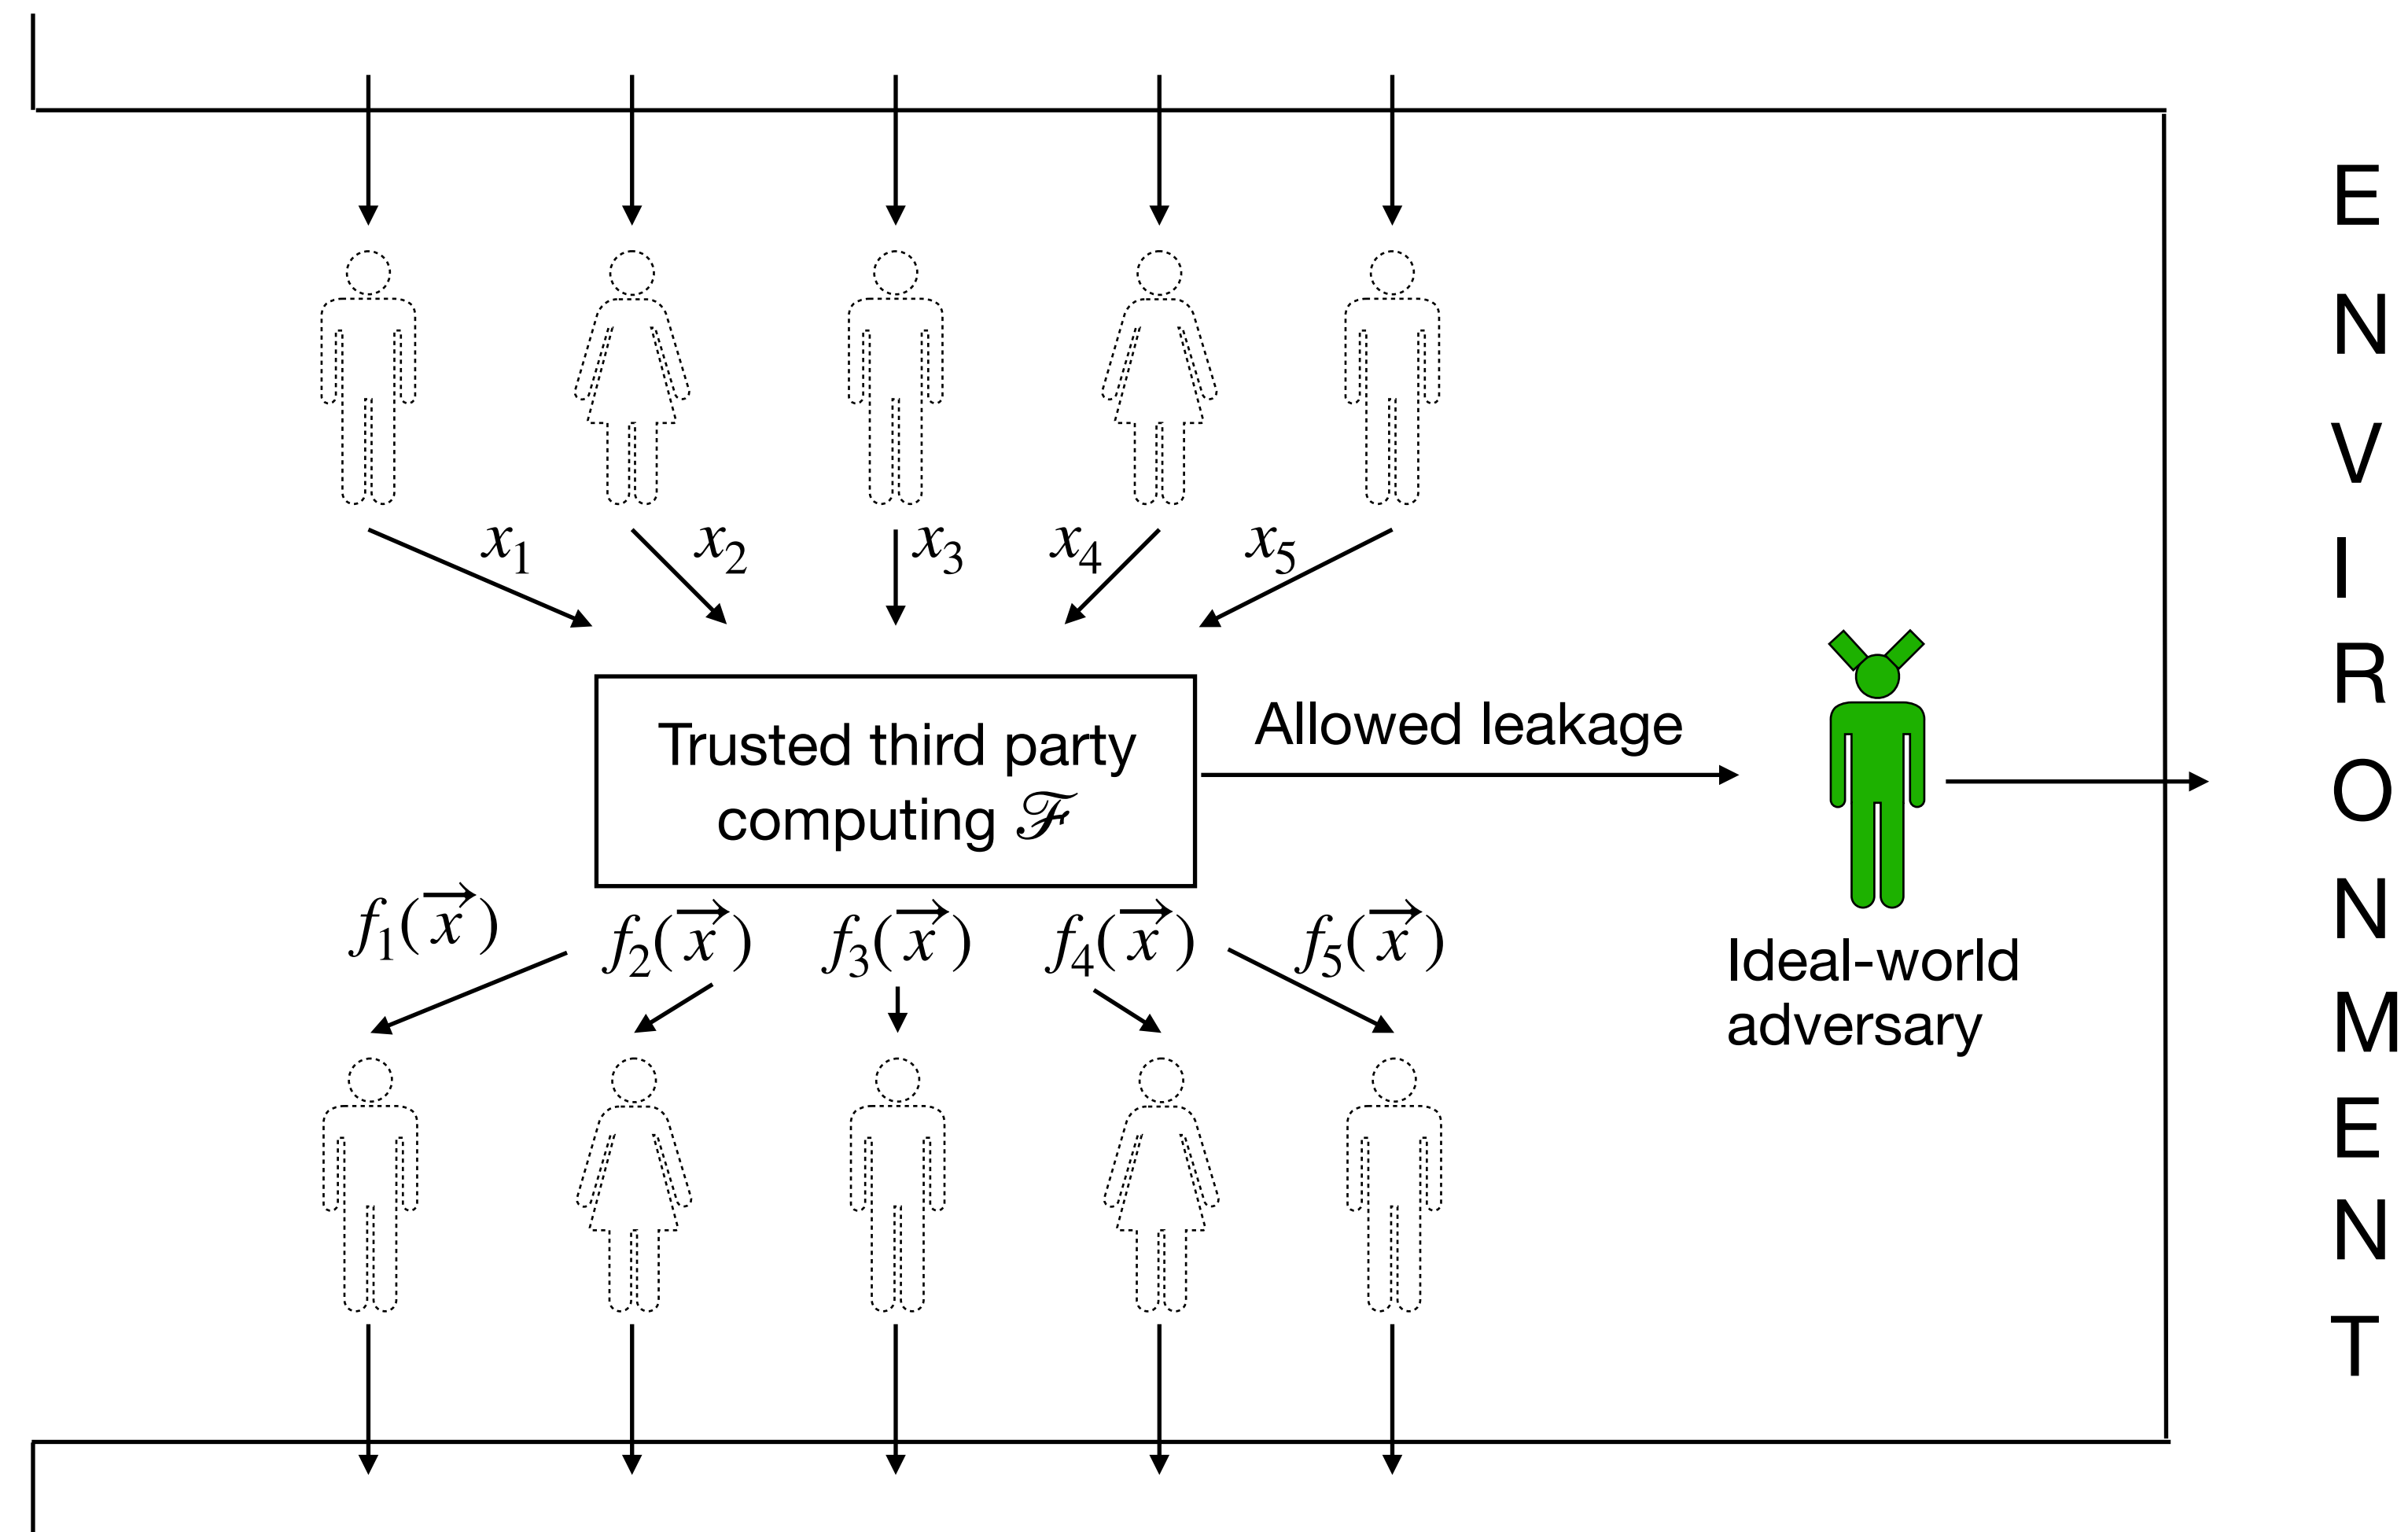
Secure multiparty computation



- All parties have private inputs x_i and wish to compute a joint function of each others' private inputs such that no party i learns anything other than $f_i(\vec{x})$.

Security of secure multiparty computation

**A hypothetical
ideal world
(secure by
definition):**



- A real protocol *securely computes functionality* \mathcal{F} if an ideal-world adversary and the dummy parties can *simulate* a view for the environment indistinguishable from its view when interacting with real parties and the real adversary.

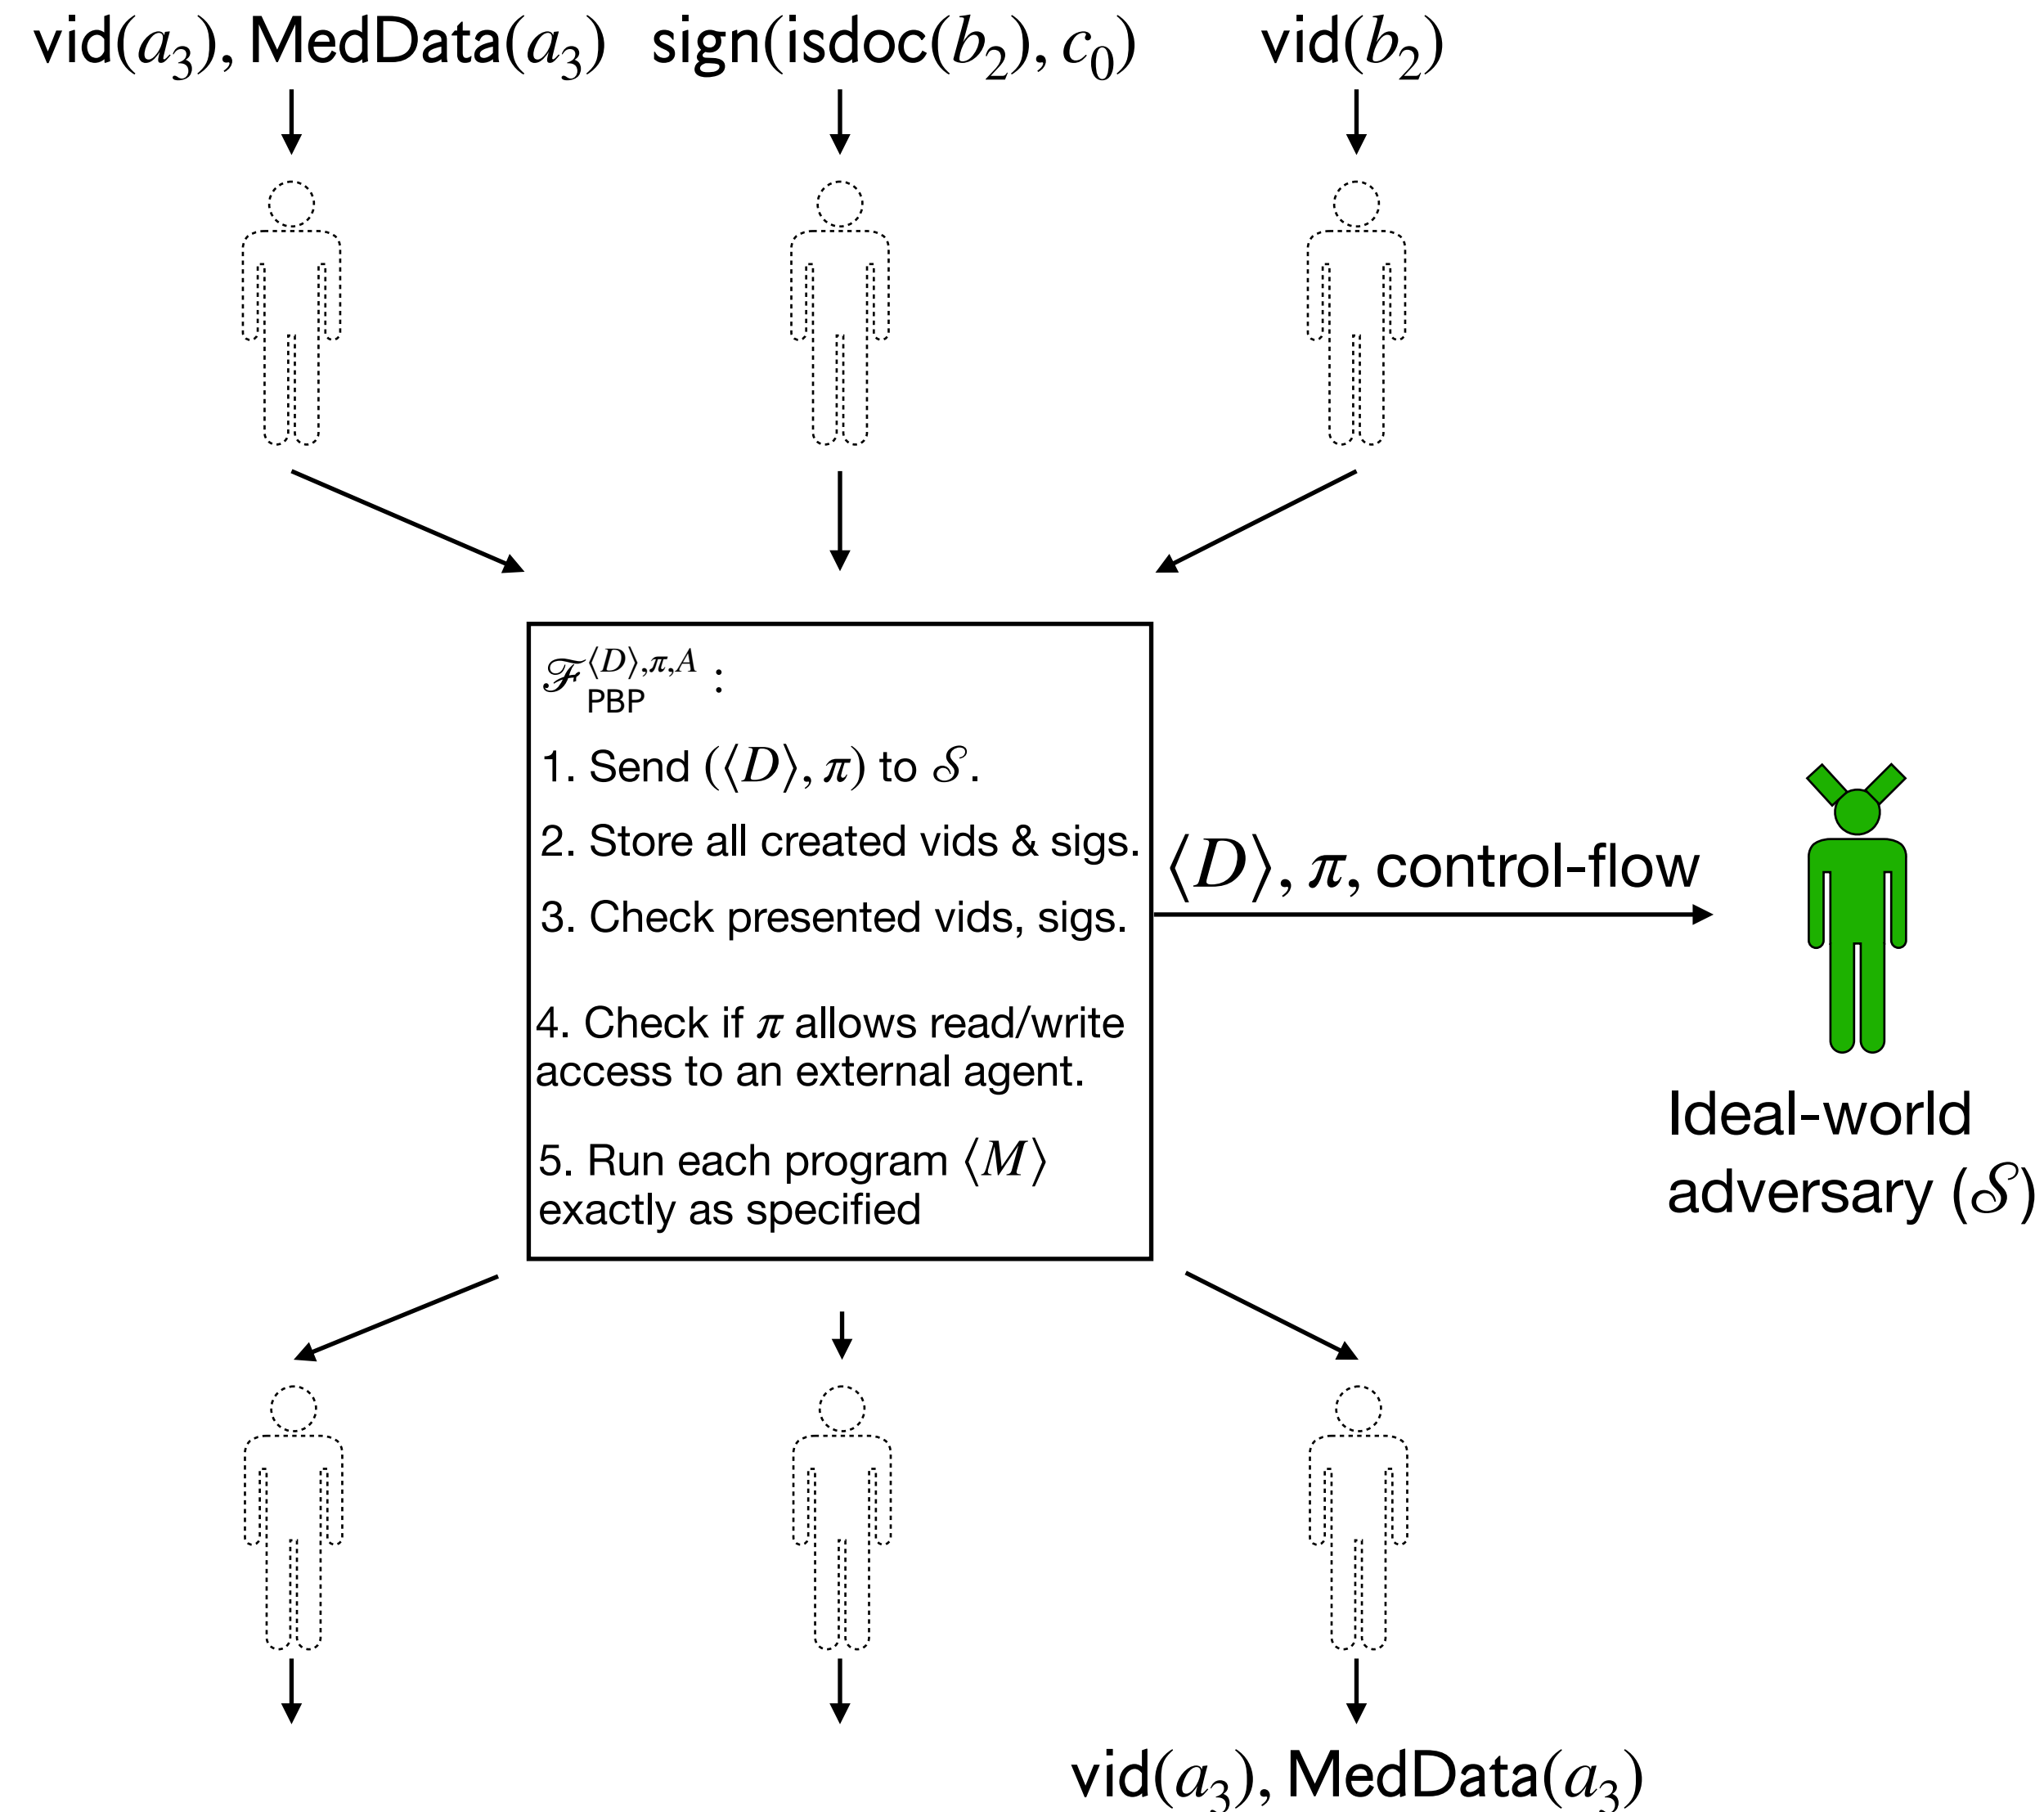
Privacy policy compliance as SMC

Policy $\pi := (\mathcal{R}, \mathcal{E})$

Data controller description

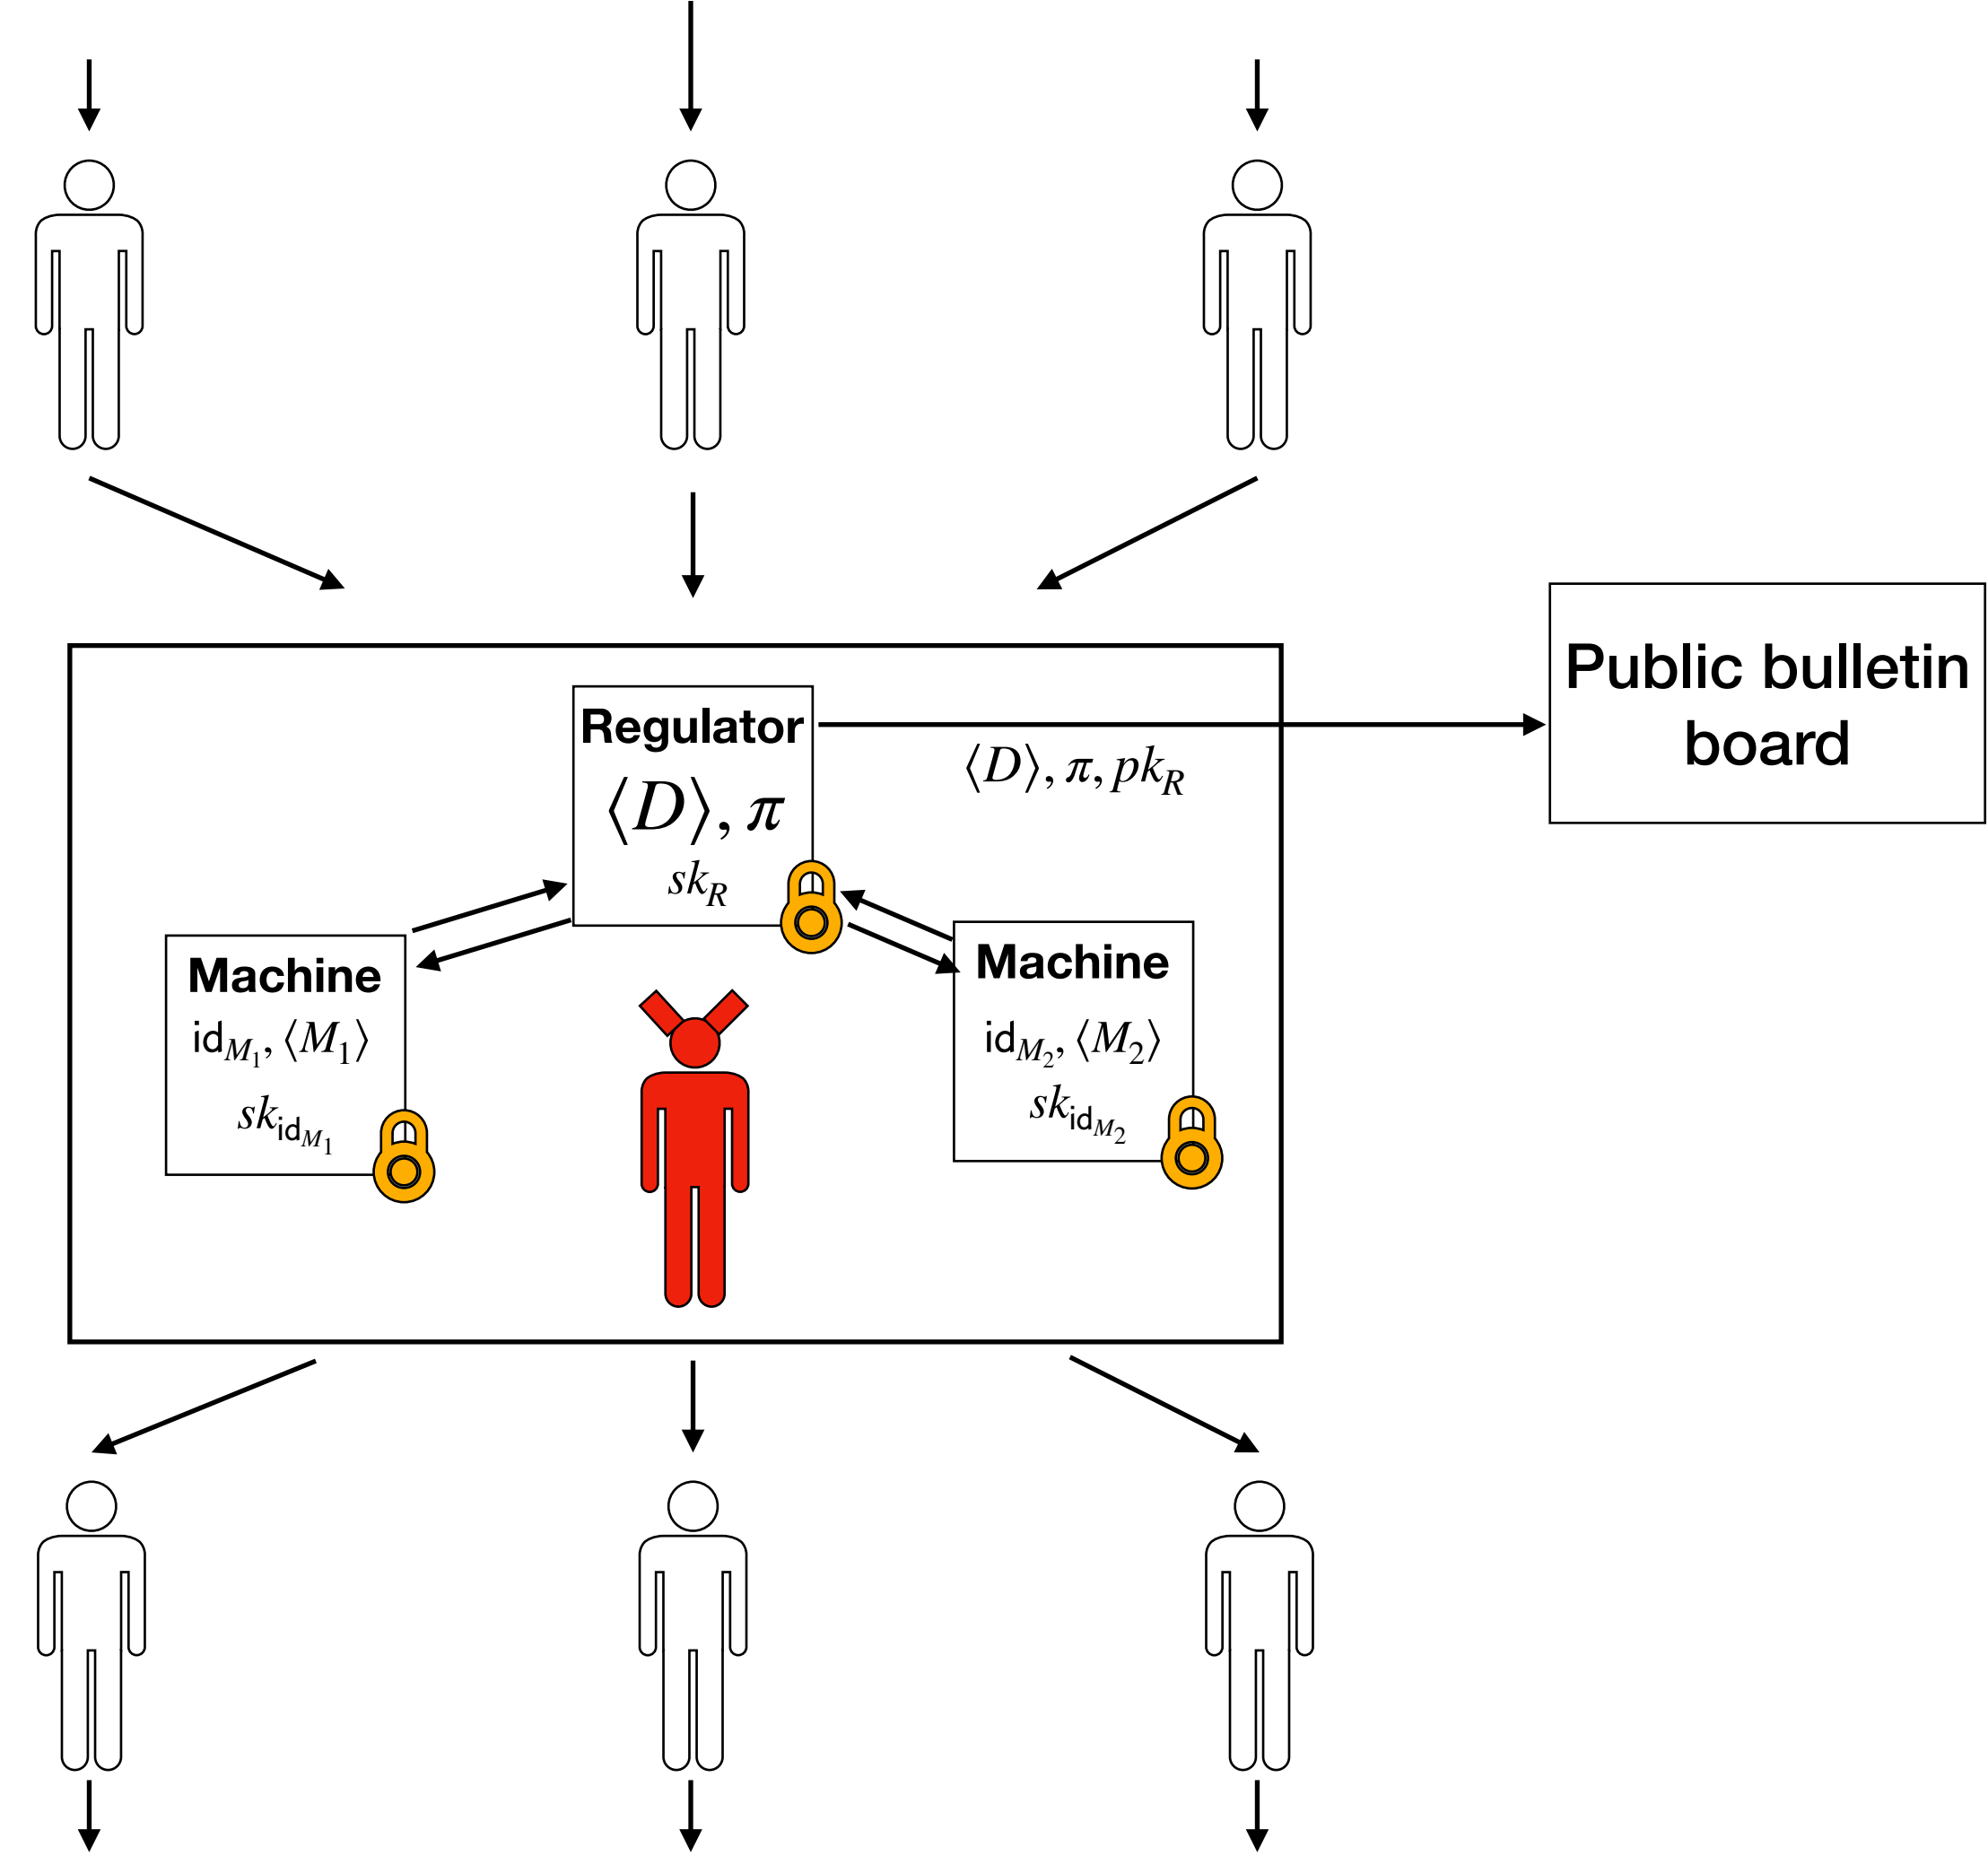
$\langle D \rangle := \{(\text{id}_M, \langle M \rangle) : M \in D\}$

**Ideal
world:**



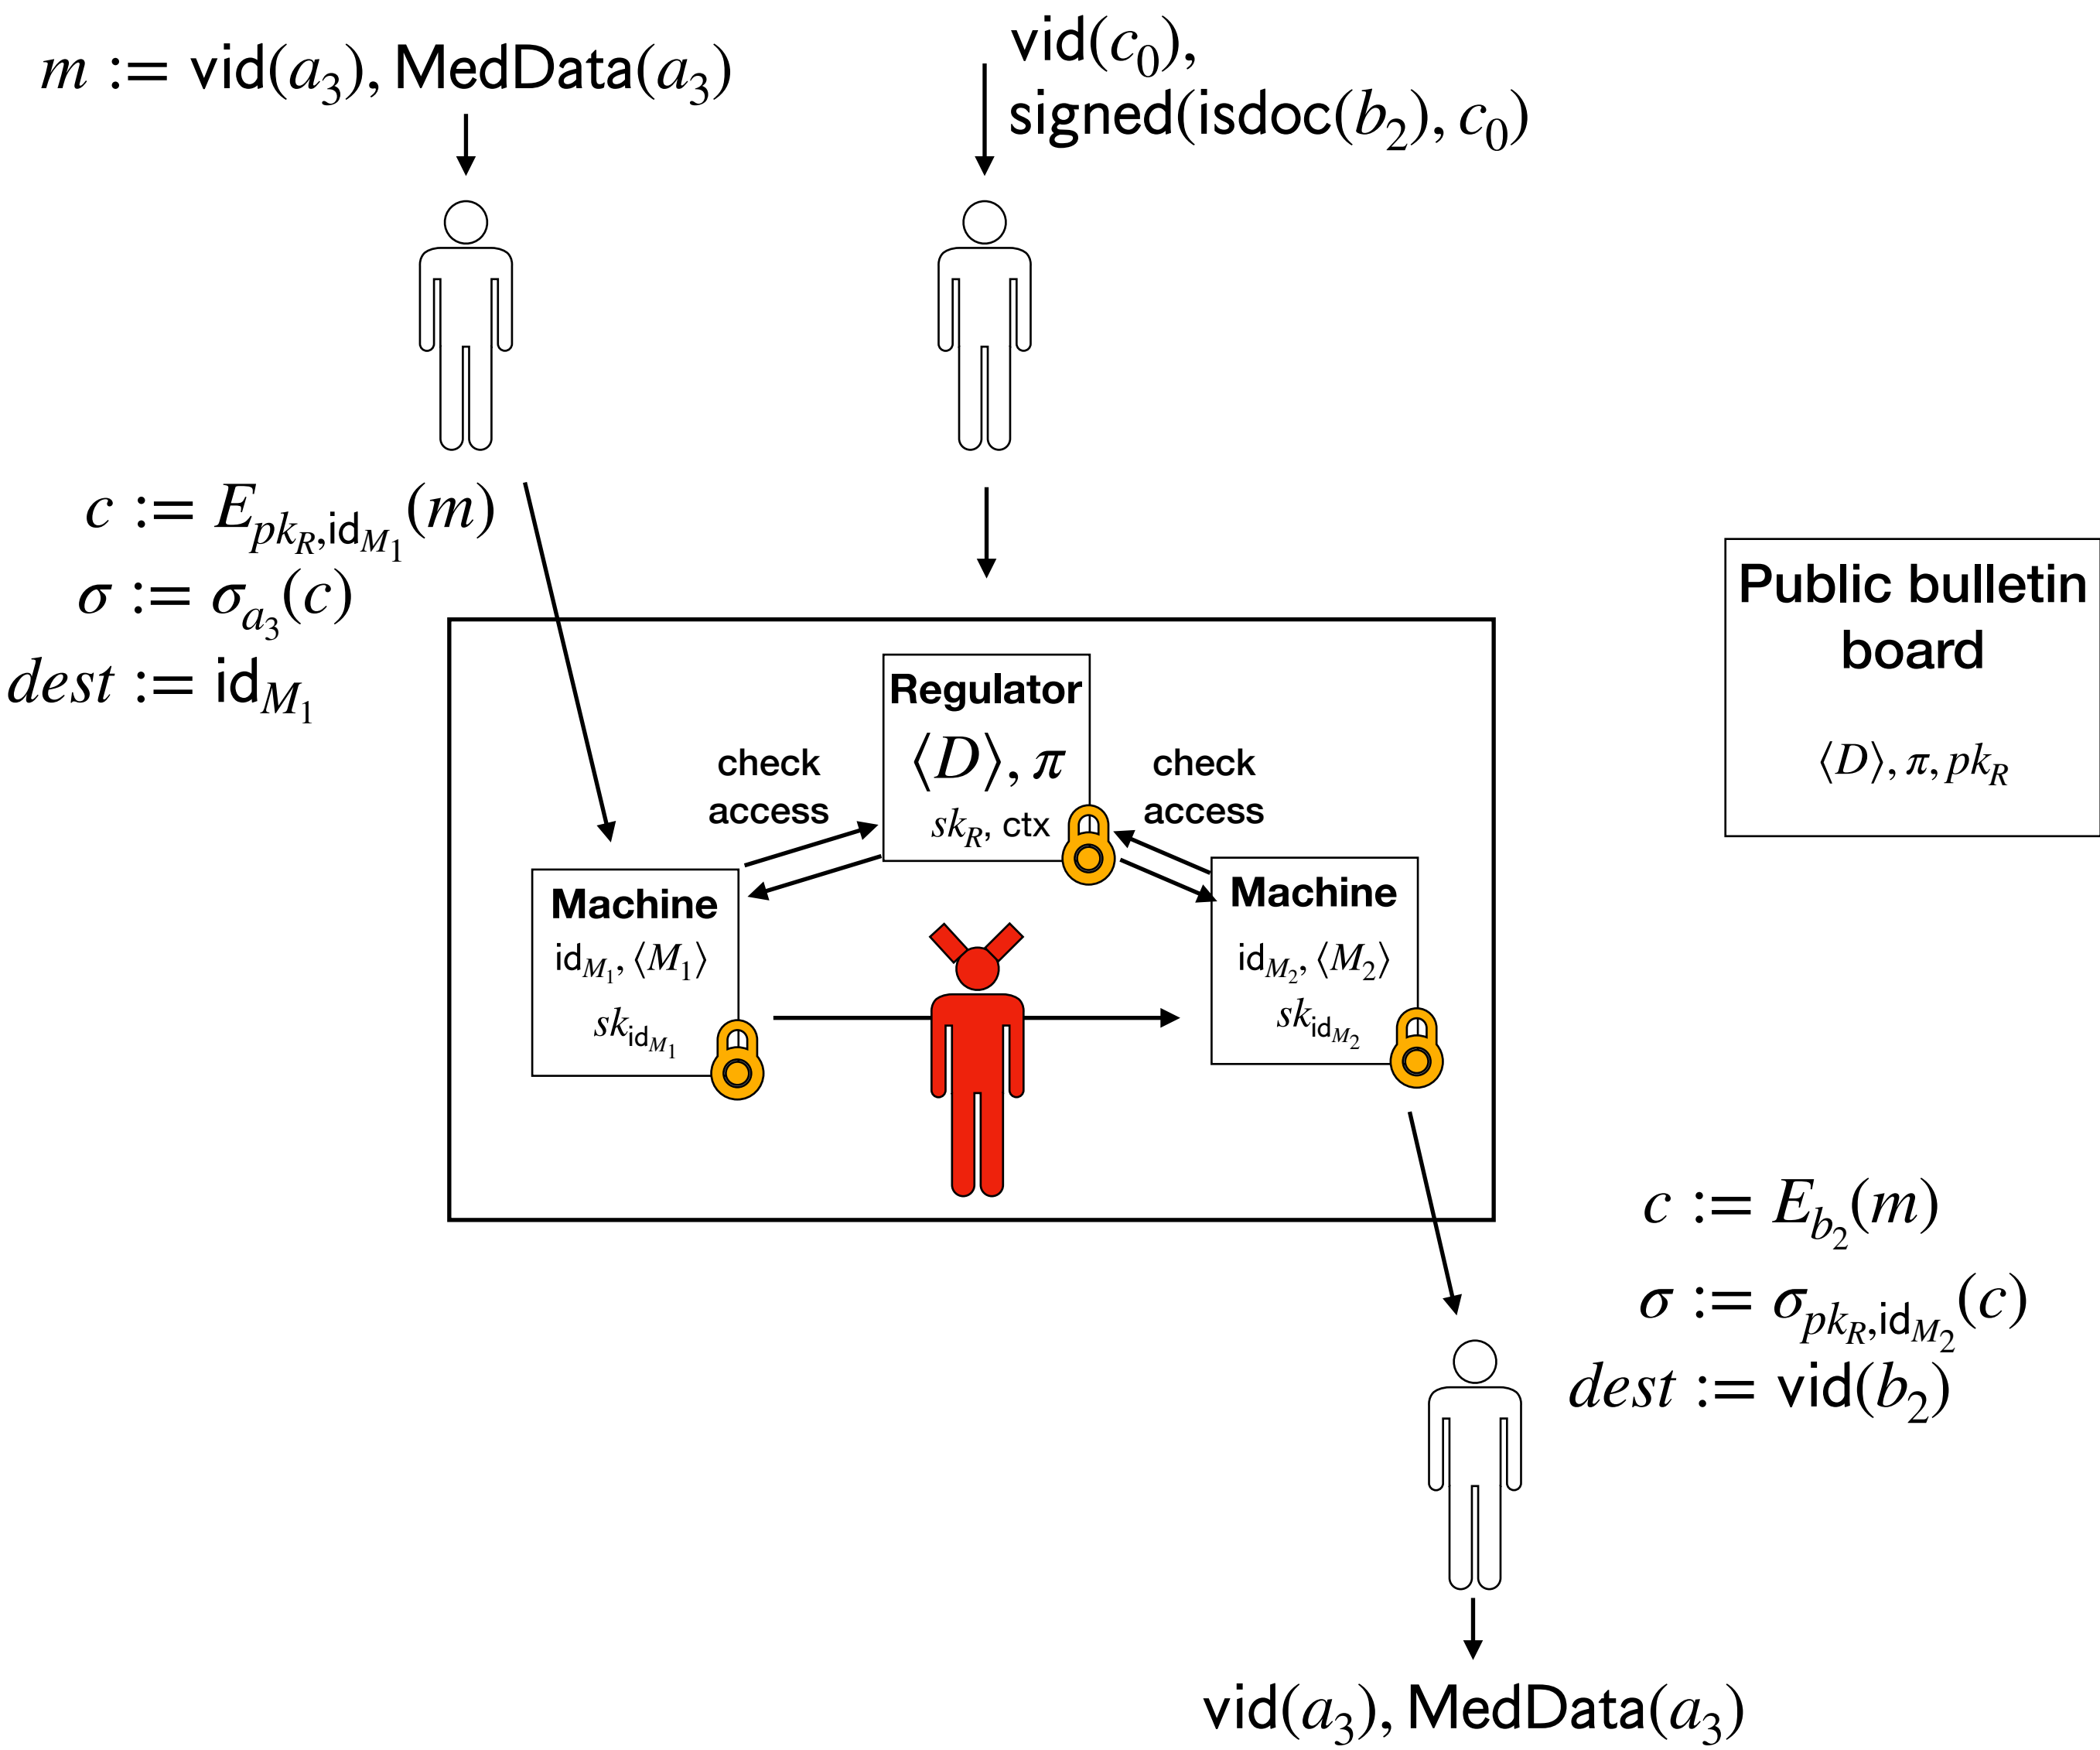
Overall architecture: initialisation

Real world:



Overall architecture: during runtime

Real
world:



Larger questions

- Regulatory capacity and will
- Identity infrastructure
- Performance issues
- Trust model
- Fast-moving private sector use-cases

References

- [Rivest et al. '78](#): *A Method for Obtaining Digital Signatures and Public-Key Cryptosystems*, CACM, Vol 21 No. 2, 1978
- [Diffie & Hellman '76](#): *New Directions in Cryptography*, IEEE Tr. Information Theory, Vol 22 No. 6, 1976
- [Goldwasser et al. '89](#): *The Knowledge Complexity of Interactive Proof Systems*, SIAM J. Computing, Vol 18 pp. 186-208, 1989
- [Goldreich et al. '91](#): *Proofs That Yield Nothing but Their Validity or All Languages in NP Have Zero-knowledge Proof Systems*, J. ACM, Vol 38 pp. 690-728, 1991
- [Pedersen '91](#): *Non-interactive and information-theoretic secure verifiable secret sharing*, CRYPTO, pp. 129-140, 1991
- [Byun & Li '06](#): *Purpose-based access control for privacy protection in relational database systems*, VLDB, 2006
- [Jafari et al. '11](#): *Towards defining semantic foundations for purpose-based privacy policies*, CODASPY, 2011
- [Tschantz et al. '11](#): *Formalizing and Enforcing Purpose Restrictions in Privacy Policies*, IEEE S&P, 2012
- [Chaum '85](#): *Security without Identification*, CACM, Vol 28 No. 10, 1985
- [Camenisch & Lysyanskaya '01](#): *An efficient system for non-transferable anonymous credentials with optional anonymity revocation*, EUROCRYPT, 2001
- [Barth et al. '06](#): *Privacy and contextual integrity: framework and applications*, IEEE S&P, 2006
- [Dinur & Nissim '03](#): *Revealing information while preserving privacy*, PODS, 2003
- [Dwork '05](#): *Differential privacy*, ICALP, 2006
- [Pinto & Santos '19](#): *Demystifying ARM Trustzone: A Comprehensive Survey*, ACM Computing Surveys, Vol 51 Issue 6, 2019