

Data Capsule: A New Paradigm for Automatic Compliance with Data Privacy Regulations

Lun Wang[‡], Joseph P. Near[†], Neel Somani[‡], Peng Gao[‡], Andrew Low[‡], David Dao[◊], and Dawn Song[‡]

[†]: University of Vermont [‡]: University of California, Berkeley [◊]: ETH Zurich

Abstract. The increasing pace of data collection has led to increasing awareness of privacy risks, resulting in new data privacy regulations like General data Protection Regulation (GDPR). Such regulations are an important step, but automatic compliance checking is challenging. In this work, we present a new paradigm, *Data Capsule*, for automatic compliance checking of data privacy regulations in heterogeneous data processing infrastructures. Our key insight is to pair up a data subject’s data with a policy governing how the data is processed. Specified in our formal policy language: PRIVPOLICY, the policy is created and provided by the data subject alongside the data, and is associated with the data throughout the life-cycle of data processing (e.g., data transformation by data processing systems, data aggregation of multiple data subjects’ data). We introduce a solution for static enforcement of privacy policies based on the concept of residual policies, and present a novel algorithm based on abstract interpretation for deriving residual policies in PRIVPOLICY. Our solution ensures compliance automatically, and is designed for deployment alongside existing infrastructure. We also design and develop PRIVGUARD, a reference data capsule manager that implements all the functionalities of *Data Capsule* paradigm.

Keywords: Data Privacy, GDPR, Formalism of Privacy Regulations, Compliance of Privacy Regulations

1 Introduction

The big data revolution has triggered an explosion in the collection and processing of our personal data, leading to numerous societal benefits and sparking brand-new fields of research. At the same time, this trend of ever-increasing data collection raises new concerns about data privacy. The prevalence of data breaches [1][2], insider attacks [3], and organizational abuses of collected data [4] indicates that these concerns are well-founded. Data privacy has thus become one of the foundational challenges of today’s technology landscape.

To address this growing challenge, governments have begun crafting data privacy regulations to protect us (the *data subjects*) from those who collect and process our personal data. Recent examples include the European Union’s *General Data Protection Regulation* (GDPR) [5], the *California Consumer Privacy Act* (CCPA) [6], the *Family Educational Rights & Privacy Act* [7], and the *Health Insurance Portability and Accountability Act* [8].

Unfortunately, compliance with data privacy regulations is extremely challenging with current data processing systems. The regulations are written in natural language, and thus are difficult to formalize for automatic enforcement. In addition, some of the systems currently used for data processing were designed and deployed before the existence of these privacy regulations, and their designs make the compliance even more difficult. For example, many existing data processing systems do not provide an option to delete data, since it was assumed that organizations would want to keep data forever [9]—but GDPR requires that a subject’s data be deleted on request. Even if the deletion is possible, its enforcement can be challenging: organizations often make multiple copies of data, without no systematic record of the copies, because each data processing platform requires its own data format; as a result, an organization may not even be able to *locate* all of the copies of a data subject’s data for deletion.

Compliance with data privacy regulations is therefore costly or impossible for many organizations. These challenges reduce the rate of compliance, resulting in harm to data subjects via privacy violations. Moreover, the cost of implementing compliance acts as a barrier to entry for small organizations, and serves to protect large organizations from new competition. Paradoxically, new data privacy regulations may actually *help* the large corporations whose abuses of data originally motivated those regulations.

This paper presents a new paradigm for *automatic compliance* with data privacy regulations in heterogeneous data processing infrastructures. Our approach is based on a new concept called the *data capsule*, which pairs up a data subject’s data with a *policy* governing how the data may be processed. The policy follows the data subject’s data *forever*, even when it is copied from one data processing system to another or mixed with data from other subjects. Our solution is designed for deployment *alongside* existing infrastructure, and requires only minimal changes to existing data processing systems. The approach is *automatic*, enabling compliance with minimal additional cost to organizations.

The Data Capsule Paradigm. We propose a new paradigm for collecting, managing, and processing sensitive personal data, called the Data Capsule Paradigm, which *automates* compliance with data privacy regulations. Our paradigm consists of three major components:

1. **Data capsule**, which contains sensitive personal data, a *policy* restricting how the data may be processed, and *metadata* relevant for data privacy concerns.
2. **Data capsule graph**, which tracks all data capsules, including data collected from data subjects and data derived (via processing) from the collected data.
3. **Data capsule manager**, which maintains the data capsule graph, registers new data capsules, enforces each capsule’s policy, and propagates metadata through the graph.

Principles of Data Privacy. To reach the design requirements for our solution, we examine four existing data privacy regulations (GDPR, CCPA, HIPAA, and FERPA). We propose five *principles of data privacy* which accurately represent common trends across these regulations: *transparency & auditing*, *consent*, *processing control*, *data portability*, and *guarantee against re-identification*. Our principles are designed to be flexible. A solution targeting these principles can be made compliant with current data privacy regulations, and is also capable of being extended to new regulations which may be proposed in the future.

PRIVPOLICY: a Formal Privacy Policy Language. To enforce the five principles of data privacy outlined above, we introduce PRIVPOLICY: a novel *formal policy language* designed around these principles, and capable of encoding the formalizable subset of recent data privacy regulations. By formalizable subset, we filter out requirements like “legitimate business purpose” in GDPR, which is almost impossible to formalize and have to rely on auditing to enforce requirements like this. We demonstrate the flexibility of PRIVPOLICY by encoding GDPR, CCPA, HIPAA, and FERPA.

PRIVPOLICY has a formal semantics, enabling a sound analysis to check whether a data processing program complies with the policy. To enforce these policies, we present a novel static analysis based on abstract interpretation. The data capsule graph enables *pipelines* of analysis programs which together satisfy a given policy. To enforce policies on these pipelines in a compositional way, we propose an approach which statically infers a *residual policy* based on an analysis program and an input policy; the residual policy encodes the policy requirements which remain to be satisfied by later programs in the pipeline, and is attached to the output data capsule of the program.

Our approach for policy enforcement is entirely static. It scales to datasets of arbitrary size, and is performed as a pre-processing step (independent of the *execution* of analysis programs). Our approach is therefore well-suited to the heterogeneous data processing infrastructures used in practice.

PRIVGUARD: a Data Capsule Manager. We design and implement PRIVGUARD, a reference data capsule manager. PRIVGUARD consists of components that manage the data capsule graph and perform static analysis of *analysis programs* which process data capsules. PRIVGUARD is designed to work with real data processing systems and introduces negligible performance overhead. Importantly, PRIVGUARD makes no changes to the format in which data is stored or the systems used to process it. Its static analysis occurs as a separate step from the processing itself, and can be performed in parallel. In a case study involving medical data, we demonstrate the use of PRIVGUARD to enforce HIPAA in the context of analysis programs for a research study.

Contributions. In summary, we make the following contributions.

- We propose five *principles of data privacy* which encompass the requirements of major data privacy regulations.
- We introduce PRIVPOLICY: a new and expressive formal language for privacy policies, which is capable of encoding policies for compliance with the formalizable subset of data privacy regulations.
- We propose the data capsule paradigm, an approach for ensuring compliance with privacy regulations encoded using PRIVPOLICY, and formalize the major components of the approach.
- We present the encoding of GDPR in PRIVPOLICY.
- We introduce a solution for static enforcement of privacy policies based on the concept of *residual policies*, and present a novel algorithm based on abstract interpretation for deriving residual policies in PRIVPOLICY.
- We design and develop PRIVGUARD, a reference data capsule manager that implements all the functionalities of data capsule paradigm.

2 Requirements of Data Privacy Regulations

Recent years have seen new efforts towards regulating data privacy, resulting in regulations like the European Union’s *General Data Protection Regulation* (GDPR). It joins more traditional regulations like the *Health Insurance Portability and Accountability Act* (HIPAA) and the *Family Educational Rights and Privacy Act* (FERPA).

2.1 Principles of Data Privacy

Historically, organizations have collected as much personal data as possible, and have not generally considered data privacy to be a high priority. The recent adoption of GDPR has forced a much wider set of organizations to consider solutions for ensuring compliance with data privacy regulations. Complying with regulations like GDPR is extremely difficult using existing systems, which generally are designed for *easy access to data* instead of strong data privacy protections. These regulations are even more difficult to satisfy when data is shared between organizational units or with third parties—yet the regulatory requirements apply even in these cases.

To address this challenge, we considered the *commonalities* between the three major data privacy regulations mentioned above to develop five *principles of data privacy*. These principles expose and generalize the fundamental ideas shared between regulations, and therefore are likely to also apply to future regulations. As described in Section 3, these five principles form the design criteria for our proposed data capsule paradigm.

In describing the five principles of data privacy, we use terminology from the GDPR. The term *data subject* refers to individuals whose data is being collected and processed, and the term *data controller* refers to organizations which collect and process data from data subjects. Briefly summarized, the five principles of data privacy are:

1. **Transparency & Auditing:** The data subject should be made aware of who has their data and how it is being processed.
2. **Consent:** The data subject should give explicit consent for the collection and processing of their data.
3. **Processing Control:** The data subject should have control over what types of processing are applied to their data.
4. **Data Portability:** The data subject should be able to obtain a copy of any data related to them.
5. **Guarantee Against Re-identification:** When possible, the results of processing should not permit the re-identification of any individual data subject.

2.2 Applying the Principles

The five principles described above represent the design criteria for our data capsule paradigm. They are specified specifically to be *at least as strong as* the common requirements of existing privacy regulations, to ensure that our approach is capable of expressing a large enough subset of existing and future regulations to ensure compliance. This section demonstrates how our five principles describe and subsume the requirements of the four major privacy regulations.

GDPR. The major pillars of GDPR fall squarely into the five requirement categories described by our principles. Articles 13 and 14 describe *transparency & auditing* requirements: the data controller must inform the data subject about the data being collected and who it is shared with. Article 4, 7 and 29WP requires *consent*: the controller must generally obtain consent from the data subject to collect or process their data. Note that there are also cases in GDPR when personal data can be used without consent, where some other “lawful basis for processing” applies, such as public interest, legal obligation, contract or the legitimate interest of the controller. However, these purposes are almost impossible to formalize so we have to rely on auditing to enforce them and omit them in the system. Articles 18 and 22 ensure *processing control*: the data subject may allow or disallow certain types of processing. Articles 15, 16, 17, and 20 require *data portability*: the data subject may obtain a copy of their data, fix errors in it, and request that it be deleted. Finally, Recital 26 and Article 29WP describes a *guarantee against re-identification*: data controllers are required to take steps to prevent the re-identification of data subjects.

CCPA. CCPA is broadly similar to GDPR, with some differences in the specifics. Like GDPR, the requirements of CCPA align well with our five principles of data privacy. Unlike GDPR, CCPA’s *consent* requirements focus on the *sale* of data, rather than its original collection. Its *access & portability* requirements focus on data deletion, and are more limited than those of GDPR. Like GDPR, CCPA ensures a *guarantee against re-identification* by allowing data subjects to sue if they are re-identified.

HIPAA. The HIPAA regulation is older than GDPR, and reflects a contemporaneously limited understanding of data privacy risks. HIPAA requires the data subject to be notified when their data is collected (*transparency & auditing*), and requires *consent* in some (but not all) cases. HIPAA requires organizations to store data in a way that prevents its unintentional release (partly ensuring a *guarantee against re-identification*), and its “safe harbor” provision specifies a specific set of data attributes which must be redacted before data is shared with other organizations (an attempt to ensure a *guarantee against re-identification*). HIPAA has only limited *processing control* and *data portability* requirements.

FERPA. The *Family Educational Rights and Privacy Act of 1974* (FERPA) is a federal law in the United States that protects the privacy of student education records. FERPA requires *consent* before a post-secondary institution shares information from a student’s education record. It also requires *access & portability*: students may inspect and review their records, and request amendments. In other respects, FERPA has fewer requirements than the other regulations described above.

3 The Data Capsule Paradigm

This section introduces the *data capsule paradigm*, an approach for addressing the five principles of data privacy described earlier. The data capsule paradigm comprises four major concepts:

- *Data capsules* combine sensitive data contributed by a data subject (or derived from such data) with a *policy* governing its use and *metadata* describing its properties.
- *Analysis programs* process the data stored inside data capsules; the input of an analysis program is a set of data capsules, and its output is a new data capsule.
- The *data capsule graph* tracks all data capsules and analysis programs, and contains edges between data capsules and the analysis programs which process them.
- The *data capsule manager* maintains the data capsule graph, propagates policies and metadata within the graph, and controls access to the data within each data capsule to ensure that capsule policies are never violated.

In Section 3.3, we demonstrate how these concepts are used to satisfy our five principles of data privacy. Section 4 describes PRIVGUARD, our proof-of-concept data capsule manager which implements the paradigm.

3.1 Life Cycle of a Data Capsule

The life cycle of a data capsule includes four phases:

1. **Data Ingestion.** *Data subjects* construct data capsules from their sensitive data via the *ingestion* process, which pairs the data with the policy which will govern its use. In our setting, the data subject is the original data owner, whose privacy we would like to protect.
2. **Analysis Program Submission.** *Analysts* who would like to process the data contained in data capsules may submit *analysis programs*, which are standard data analytics programs augmented with API calls to the data capsule manager to obtain raw data for processing.
3. **Data Processing.** Periodically, or at a time decided by the analyst, the data capsule manager may run an analysis program. At this time, the data capsule manager statically determines the set of input data capsules to the program, and performs static analysis to verify that the program would not violate the policies of any of its inputs. As part of this process, the data capsule manager computes a *residual policy*, which is the new policy to be attached to the program’s output. The data capsule manager then runs the program, and constructs a new data capsule by pairing up the program’s output with the residual policy computed earlier.
4. **Declassification.** A data capsule whose policy has been satisfied completely may be viewed by the analyst in a process called *declassification*. When an analyst requests that the data

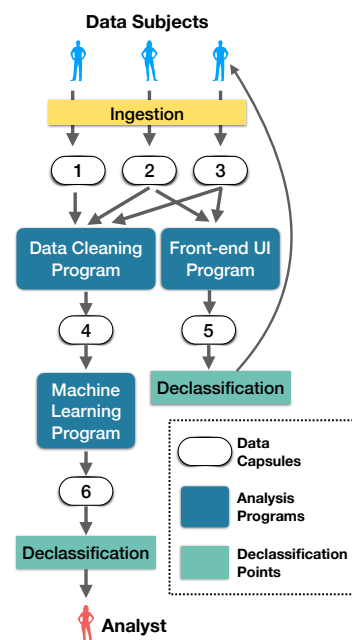


Fig. 1: Example Data Capsule Graph

capsule manager declassify a particular data capsule, the manager verifies that its policy has been satisfied, and that the analyst has the appropriate role, then sends the raw data to the analyst. Declassification is the only process by which data stored in a data capsule can be divorced from its policy.

3.2 The Data Capsule Manager

The data capsule life cycle is supported by a system implementing the functionality of the data capsule manager. The primary responsibility of the data capsule manager is to maintain the data capsule graph and maintain its invariants—namely, that no data capsule’s policy is violated, that new data capsules resulting from analysis programs have the right policies, and that metadata is propagated correctly. We describe our reference implementation, PRIVGUARD, in Section 4.

Figure 1 contains a global view of an example data capsule graph. This graph contains two different organizations representing data controllers, and data subjects associated with each one. Both organizations use analysis programs which combine and clean data from multiple data subjects into a single data capsule; a third analyst uses data capsules from *both* organizations to perform marketing research. Such a situation is allowed under privacy regulations like GDPR, as long as the policies specified by the data subjects allow it. This example therefore demonstrates the ability of the data capsule paradigm to break down data silos while at the same time maintaining privacy for data subjects—a key benefit of the paradigm.

In this example, the policy attached to each data subject’s capsule is likely to be a formal representation of GDPR. The data capsule paradigm requires a formal encoding of policies with the ability to efficiently compute residual policies; we describe our solution to this challenge in Section 5.

Note that the data capsules containing the data subjects’ combined data (capsules 1, 2, 3, and 4) cannot be viewed by *anyone*, since their policies have not been satisfied. This is a common situation in the data capsule paradigm, and it allows implementing useful patterns such as extract-transform-load (ETL) style pipelines [10]. In such cases, analysts may submit analysis programs whose primary purpose is to prepare data for *other* analysis programs; after being processed by some (potentially long) pipeline of analysis programs, the final output has satisfied all of the input policies and may be declassified. The intermediate results of such pipelines can never be viewed by the analyst.

3.3 Satisfying the Principles of Data Privacy

The data capsule paradigm is designed specifically to enable systems which satisfy the principles of data privacy laid out in Section 2.

Transparency & Auditing. The data capsule manager satisfies transparency & auditing by consulting the data capsule graph. The global view of the graph (as seen in Figure 1) can be restricted to contain only the elements reachable from the ingested data capsules of a single data subject, and the resulting sub-graph represents all of the data collected about or derived from the subject, plus all of the processing tasks performed on that data.

Consent. The data capsule manager tracks consent given by the data subject as metadata for each data capsule. Data subjects can be prompted to give consent when new analysis programs are submitted, or when they are executed.

Processing Control. The formal policies attached to data capsules can restrict the processing of the data stored in those capsules. These policies typically encode the restrictions present in data privacy regulations, and the data capsule manager employs a static analysis to verify that submitted analysis programs do not violate the relevant policies. This process is described in Section 5.

Data Portability. To satisfy the data portability principle, the data capsule manager allows each data subject to download his or her data capsules. The data capsule manager can also provide data capsules *derived* from the subject’s capsules, since these are reachable capsules in the data capsule graph. However, the derived data returned to the data subject must *not* include data derived from the capsules of *other subjects*, so a one-to-one mapping must exist between rows in the input and output capsules for each analysis program involved. We formalize this process in Section 5.

The same mechanism is used for data deletion. When a data subject wishes to delete a capsule, the set of capsules derived from that capsule is calculated, and these derived capsules are re-computed *without* the deleted capsule included in the input.

Guarantee Against Re-identification. To provide a robust formal guarantee against re-identification, the data capsule manager supports the use of various techniques for anonymization, including both informal techniques

(e.g. removing “personal health information” to satisfy HIPAA) and formal techniques (e.g. k -anonymity, ℓ -diversity, and differential privacy). A data capsule’s policy may require that analysis programs apply one of these techniques to protect against re-identification attacks.

4 PRIVGUARD: a Data Capsule Manager

We have designed and implemented a reference data capsule manager, called PRIVGUARD. The PRIVGUARD system manages the data capsule graph, propagates policies and metadata, and uses static analysis to calculate residual policies on analysis programs.

Figure 2 summarizes the architecture of PRIVGUARD. The two major components of the system are the data capsule manager itself, which maintains the data capsule graph, and the static analyzer, which analyzes policies and analysis programs to compute residual policies. We describe the data capsule manager here, and formalize the static analyzer in Section 5.

Finally, `outputCapsule` defines an output data capsule of the analysis program. The analyst specifies a dataframe containing the output data, and PRIVGUARD automatically attaches the correct residual policy. This process is formalized in Section 5.

Deployment & Integration. The data capsule paradigm is intended to be integrated with existing heterogeneous data processing infrastructures, like the ones already in place for data analysis at many organizations, and PRIVGUARD is designed to facilitate such deployments. These infrastructures leverage a variety of data stores, including SQL databases [11], key-value stores like MongoDB [12], distributed filesystems like HDFS [13], and short-term publish/subscribe systems like Cassandra [14]. They employ many different techniques for processing the data, including SQL engines and distributed systems like MapReduce [15], Hadoop [16], and Spark [17].

To work successfully in such a heterogeneous environment, PRIVGUARD is deployed *alongside* the existing infrastructure. As shown in Figure 2, policies and metadata are stored separately from the data itself, and the data can remain in the most efficient format for processing (e.g. stored in CSV files, in HDFS, or in a SQL database).

Similarly, PRIVGUARD’s static analyzer uses a common representation to encode the semantics of many different kinds of analysis programs, so it works for many programming languages and platforms. The only platform-specific code is the small PRIVGUARD API, which allows analysis programs to interact with the data capsule manager. Our static analysis is based on abstract interpretation, a concept which extends to all common programming paradigms. Section 5 formalizes the analysis for dataflow-based systems which are close to relational algebra (e.g. SQL, Pandas, Hadoop, Spark); extending it to functional programs or traditional imperative or object-oriented programs is straightforward.

5 Policies & Policy Enforcement

This section describes our formal language for specifying policies on data capsules, and our static approach for enforcing these policies when an analytics program is registered with the system. We describe each of the four major components of this approach:

- Our policy specification language: PRIVPOLICY (§ 5.1).
- A set of attribute definitions suitable for encoding policies like GDPR and HIPAA, which are more expressive than the corresponding attributes proposed in previous work (§ 5.2).
- A flexible approach for deriving the *policy effects* of an analysis program via abstract interpretation (§ 5.3).
- A formal procedure for determining the *residual policy* on the output of an analysis program (§ 5.4).

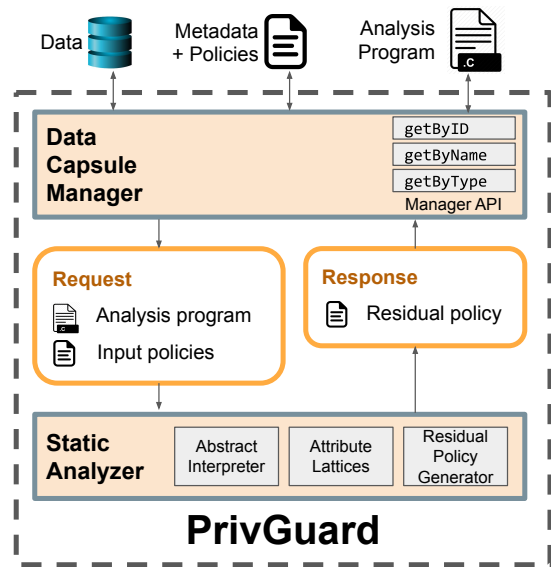


Fig. 2: The Architecture of PRIVGUARD.

5.1 PRIVPOLICY: Policy Specification Language

$A \in \text{attribute} \quad ::= \text{attrName attrValue}$ $C \in \text{policy clause} ::= A \mid A \text{ AND } C \mid A \text{ OR } C$ $P \in \text{policy} \quad ::= (\text{ALLOW } C)^+$	$A ::= \text{attrName attrValue}$ $C_{DNF} \subseteq \mathcal{P}(A)$ $P_{DNF} \subseteq \mathcal{P}(C_{DNF})$
--	---

(1) PRIVPOLICY surface syntax.

(2) PRIVPOLICY disjunctive normal form.

Fig. 3: Surface Syntax & Normal Form.

Our policy specification language: PRIVPOLICY is inspired by the LEGALEASE language [18], with small changes to surface syntax to account for our more expressive attribute lattices and ability to compute residual policies.

```

1 ALLOW SCHEMA NotPII
2 AND NOTIFICATION_REQUIRED
3 AND (ROLE $user.id
4 OR (CONSENT_REQUIRED
5 AND DECLASS DP 1 0.000001))

```

Fig. 4: A subset of GDPR using PRIVPOLICY.

of the policy, and *SCHEMA*, *NOTIFICATION_REQUIRED*, *ROLE*, *CONSENT_REQUIRED*, and *DECLASS* are attributes. This subset includes only a single clause, which says that information which is not personally identifiable may be processed by the data controller, as long as the data subject is notified, and either the results are only viewed by the data subject, or the data subject gives consent and differential privacy is used to prevent re-identification based on the results.

Conversion to Disjunctive Normal Form. Our first step in policy enforcement is to convert the policy to *disjunctive normal form* (DNF), a common conversion in constraint solving. Conversion to DNF requires removing OR expressions from each clause of the policy; we accomplish this by distributing conjunction over disjunction and then splitting the top-level disjuncts within each clause into separate clauses. After converting to DNF, we can eliminate the explicit uses of **AND** and **OR**, and represent the policy as a *set of clauses*, each of which is a *set of attributes* as shown in Figure 3 (2). The disjunctive normal form of our running example policy is shown in Figure 5. Note that the disjunctive normal form of our example contains two clauses, due to the use of **OR** in the original policy.

5.2 Policy Attributes

LEGALEASE [18] organizes attribute values into *concept lattices* [19], and these lattices give policies their semantics. Instead of concept lattices, PRIVPOLICY leverages *abstract domains* inspired by work on abstract interpretation of programs [20]. This novel approach enables more expressive attributes (for example, the *FILTER* attribute) and also formalizes the connection between the semantics of policies and the semantics of analysis programs.

We require each attribute domain to define the standard lattice operations required of an abstract domain: a partial order (\sqsubseteq), join (\sqcup), and meet (\sqcap), as well as top and bottom elements \top and \perp . Many of these can be defined in terms of the corresponding operations of an existing abstract domain from the abstract interpretation literature.

Filter Attributes. One example of our expressive attribute domains is the one for the *FILTER* attribute, which filters data based on integer-valued fields. The attribute domain for *FILTER* is defined in terms of an *interval* abstract domain [20]. We say *filter* $: f : i$ when the value of column f lies in the interval i . Then, we define the following operations on *FILTER* attributes, completing its attribute domain:

The grammar for the surface syntax of PRIVPOLICY is given in Figure 3 (1). The language allows specifying an arbitrary number of *clauses*, each of which encodes a formula containing conjunctions and disjunctions over attribute values. Effectively, each clause of a policy in our language encodes one way to satisfy the overall policy.

Example. Figure 4 specifies a subset of GDPR using PRIVPOLICY. Each **ALLOW** keyword denotes a clause

```

{{SCHEMA NotPII, NOTIFICATION_REQUIRED,
  ROLE $user.id}
 {SCHEMA NotPII, NOTIFICATION_REQUIRED,
  CONSENT_REQUIRED, DECLASS DP(1.0)}}

```

Fig. 5: Disjunctive normal form of the example policy.

$$\begin{aligned}
\text{filter} : f : i_1 \sqcup \text{filter} : f : i_2 &= \text{filter} : f : i_1 \sqcup i_2 \\
\text{filter} : f : i_1 \sqcap \text{filter} : f : i_2 &= \text{filter} : f : i_1 \sqcap i_2 \\
\text{filter} : f : i_1 \sqsubseteq \text{filter} : f : i_2 &= : i_1 \sqsubseteq i_2
\end{aligned}$$

Schema Attributes. The schema attribute leverages a *set* abstract domain, in which containment is defined in terms of an underlying (finite) lattice of datatypes:

$$\begin{aligned}
\text{schema} : S_1 \sqcup \text{schema} : S_2 &= \text{schema} : \{s' \mid s_1 \in S_1 \wedge s_2 \in S_2 \wedge s' = s_1 \sqcup s_2\} \\
\text{schema} : S_1 \sqcap \text{schema} : S_2 &= \text{schema} : \{s' \mid s_1 \in S_1 \wedge s_2 \in S_2 \wedge s' = s_1 \sqcap s_2\} \\
\text{schema} : S_1 \sqsubseteq \text{schema} : S_2 &= \forall s_1 \in S_1, s_2 \in S_2 . s_1 \sqsubseteq s_2
\end{aligned}$$

Other Attributes. In PRIVPOLICY, as in LEGALEASE, the partial ordering for analyst roles is typically finite. It encodes the important properties of each analyst (e.g. for GDPR, the government typically has more authority to analyze data than members of the public). The `role`, `declass`, and `redact` attributes are defined by finite lattices. We omit the details here.

5.3 Abstract Interpretation of Analysis Programs

$$\begin{aligned}
f \in \text{field} \quad m \in \text{int} \quad s \in \text{schema} \quad x \in \text{data capsules} \\
\delta \in \text{filter} \quad ::= f < m \mid f > m \\
e \in \text{expression} ::= \text{getDC}(x) \mid \text{filter}(\varphi, e) \mid \text{project}(s, e) \mid \text{redact}(a, e) \\
\quad \quad \quad \mid \text{join}(e, e) \mid \text{union}(e, e) \mid \text{dpCount}(\epsilon, \delta, e)
\end{aligned}$$

Fig. 6: Program Surface Syntax

We next describe the use of abstract interpretation to determine the *policy effect* of an analysis program. We introduce this concept using a simple dataflow-oriented language, similar to relational algebra, Pandas, or Spark, presented in Figure 6. We write an abstract data capsule with schema s and policy effect ψ as $\mathbb{D}[s, \psi]$. A data capsule environment Δ maps data capsule IDs to their schemas (i.e. $\Delta : id \rightarrow s$).

$$\begin{array}{c}
\frac{\Delta(id) = s}{\Delta \vdash \text{getDC}(id) : \mathbb{D}[s, \emptyset]} \text{GETDC} \qquad \frac{\Delta \vdash e : \mathbb{D}[s, \psi] \quad \varphi \rightsquigarrow_s a : v}{\Delta \vdash \text{filter}(\varphi, e) : \mathbb{D}[s, \psi + \text{filter} : a : v]} \text{FILTER} \\
\\
\frac{\Delta \vdash e : \mathbb{D}[s, \psi] \quad s' \sqsubseteq s}{\Delta \vdash \text{project}(s', e) : \mathbb{D}[s', \psi + \text{schema} : s']} \text{PROJECT} \qquad \frac{\Delta \vdash e : \mathbb{D}[s, \psi] \quad a \in s \quad e_r \rightsquigarrow_s v}{\Delta \vdash \text{redact}(a, e_r, e) : \mathbb{D}[s, \psi + \text{redact} : a : v]} \text{REDACT} \\
\\
\frac{\Delta \vdash e_1 : \mathbb{D}[s_1, \psi_1] \quad \Delta \vdash e_2 : \mathbb{D}[s_2, \psi_2]}{\Delta \vdash \text{join}(e_1, e_2) : \mathbb{D}[s_1 \cup s_2, \psi_1 \cup \psi_2]} \text{JOIN} \qquad \frac{\Delta \vdash e_1 : \mathbb{D}[s, \psi_1] \quad \Delta \vdash e_2 : \mathbb{D}[s, \psi_2]}{\Delta \vdash \text{union}(e_1, e_2) : \mathbb{D}[s, \emptyset]} \text{UNION} \\
\\
\frac{\Delta \vdash e : \mathbb{D}[s, \psi]}{\Delta \vdash \text{dpCount}(\epsilon, \delta, e) : \mathbb{D}[s, \psi + \text{declass} : \text{DP}(\epsilon, \delta)]} \text{DPCOUNT}
\end{array}$$

Fig. 7: Sample rules implementing an abstract interpreter for the data capsule expressions in the language presented in Figure 6.

We present the *abstract interpreter* [20] for PRIVPOLICY in Figure 7. If we can use the semantics to build a derivation tree of the form $\Delta \vdash e : \mathbb{D}[s, \psi]$, then we know that the program is guaranteed to satisfy the policy clause ψ (or any clause which is less restrictive than ψ).

5.4 Computing Residual Policies

Let $\mathcal{Y}(id)$ be the policy of the data capsule with ID id . The free variables of a program e , written $fv(e)$, are the data capsule IDs it uses.

We define the *input policy* of a program e to be the least upper bound of the policies of its free variables:

$$\mathcal{Y}_{in(e)} = \bigsqcup_{id \in fv(e)} \{\mathcal{Y}(id)\}$$

This semantics means that the input policy will be *at least as restrictive* as the *most restrictive* policy on an input data capsule. It is computable as follows, because the disjunctive normal form of a policy is a set of sets:

$$p_1 \sqcup p_2 = \{c_1 \cup c_2 \mid c_1 \in p_1 \wedge c_2 \in p_2\}$$

The *residual policy* applied to the output data capsule is computed by considering each clause in the input policy, and computing its residual based on the policy effect of the program. The residual policy is computed using the following rule:

$$\frac{\vdash e : \mathbb{D}[s, \psi]}{\mathcal{Y}_{out(e)} = \{c' \mid c \in \mathcal{Y}_{in(e)} \wedge residual(c, \psi) = c'\}} \text{RP}$$

where

$$residual(c, \psi) = c - \{k : p \mid k : p \in c \wedge satisfies(k : p, \psi)\}$$

$$satisfies(k : p, \psi) = \exists k : p' \in \psi.p \sqsubseteq p'$$

Here, the *satisfies* relation holds for an attribute $k : p$ in the policy when there exists an attribute $k : p'$ in the policy effect of the program, such that p (from the policy) is less restrictive than p' (the guarantee made by the program). Essentially, we compute the residual policy from the input policy by *removing* all attributes for which *satisfies* holds.

6 Formally Encoding GDPR

This section describes our formal encoding of a subset of GDPR, which is intended to ensure automated compliance with the regulation. We are in the process of developing similar encodings for other regulations, including HIPAA, FERPA.

Figure 8 contains our formal encoding. The first clause (lines 1-5) allows the use of data for any purpose, as long as it is protected against re-identification and subject to consent by the data subject. The third clause (lines 11-13) allows the use of personal information by organizations affiliated with the data subject—a relationship which we encode as a metafunction. The final two clauses specify specific public interest exceptions, for public health (lines 15-17) and for judicial purposes (lines 19-21).

GDPR is designed specifically to be simple and easy for users to understand, and its requirements are well-aligned with

```

1 ALLOW SCHEMA NotPII
2 AND NOTIFICATION_REQUIRED
3 AND (ROLE $user_id
4   OR (CONSENT_REQUIRED
5     AND DECLASS DifferentialPrivacy 1 0.000001))
6
7 # Definitions mainly given in Article 6. Also Article 4, 25, 32
8 ALLOW SCHEMA PersonalInformation (Article 9)
9   AND CONSENT_REQUIRED (Article 4, 6)
10
11 ALLOW SCHEMA PersonalInformation (Article 9)
12   AND ROLE UserAffiliatedOrganizations($user_id)
13   AND SCHEMA HasAppropriateSafeguards (article 25, 32, 46)
14
15 ALLOW SCHEMA PersonalInformation
16   AND ROLE SupervisoryAuthority OR ROLE
17   HealthcareOrganization
18   AND PURPOSE PublicInterest LegalObligation PublicHealth
19
20 ALLOW SCHEMA PersonalInformation
21   AND ROLE LegalAuthority
   AND PURPOSE PublicInterest ForJudicialPurposes

```

Fig. 8: Formal encoding of GDPR.

our five principles of data privacy. Our formal encoding is therefore correspondingly simple. We expect that most uses of data will fall under the third clause (for “business uses” of data, e.g. displaying Tweets to a Twitter user) or the first clause (for other purposes, e.g. marketing).

Data subjects who wish to modify this policy will generally specify more rigorous settings for the technologies used to prevent their re-identification. For example, a privacy-conscious data subject may require differential privacy in the first clause, instead of allowing any available de-identification approach.

7 Performance Evaluation

Recall that one of the goals of PRIVGUARD is to easily work with existing heterogeneous data processing systems and incur smallest additional overhead to the analysis itself. In order to achieve this goal, PRIVGUARD must have good scalability when the number of users are large. In this section, we first conduct an end-to-end evaluation to figure the bottleneck of the system scalability. Then we conduct several experimental evaluation to test the scalability of the bottleneck component. Specifically, we want to answer the following question: how is the scalability of PRIVGUARD and how much overhead PRIVGUARD will incur into the original data processing system.

7.1 Experimental Design & Setup

We first conduct an end-to-end evaluation to determine the sources of overhead when PRIVGUARD is used in a complete analysis. We then focus on the performance of policy ingestion as described in Section 3.1, which turns out to be the largest source of overhead in PRIVGUARD. We vary the number of data capsules from 2 to 1024 (with a log interval 2) and the policies are random subset of GDPR, HIPAA, FERPA or CCPA following a Gaussian distribution. The experiments are run on a single thread on top of an Ubuntu 16.04 LTS server with 32 AMD Opteron Processors. The experiments are run for 10 iterations to reach relatively stable results. Performance of PRIVGUARD does not depend on the data, so a real deployment will behave just like the simulation if the policies are similar.

7.2 Evaluation Results

In the following, we show and summarize the experiment results. In addition, we discuss and analyze the reasons for our findings.

End-to-end evaluation. We first conduct an end-to-end evaluation to figure out the most time-consuming component in the execution path of PRIVGUARD. We evaluate the time of each component in the execution path in a system with 1024 clients with random subsets of HIPAA.

Operation	Parsing	Ingestion	Residual Policy
Time (ms)	83	9769	11

Table 1: End-to-end evaluation

The results are summarized in 1. The “Parsing” column represents the time parsing the analysis program. The “Ingestion” column represents the time ingesting the input policies. The “Residual Policy” column represents the time computing residual policy given the ingested input policy and the analysis program. Note that the input policy ingestion take up almost all of the running time, which indicates this operation as the bottleneck of the system.

The results demonstrate that the performance overhead of PRIVGUARD is negligible for these programs when the policy guard approach is used, and the bottleneck of PRIVGUARD is the ingestion operation.

Policy ingestion evaluation : Next, we perform a targeted microbenchmark to evaluate the scalability of the policy ingestion operation. Figure 9 contains the results. As we can observe in the figure, the running




Fig. 9: Scalability of policy analysis: ingestion operation.

time exhibits a polynomial growth (recall that both the x-axis and y-axis are in log scale) at first and then keep stable after the policy number reaches some threshold. The reason is that the ingestion is implemented using a least upper bound (LUB) operation, and the LUB operation in PRIVGUARD is composed of two sub-operations: (1) unique with $\mathcal{O}(n \log n)$ complexity, and (2) reduce with $\mathcal{O}(n')$ complexity (n' is the number of policies after unique operation). Because policies are a random subset of some complete policy (GDPR, HIPAA, FERPA and CCPA in this case), if the number of policies are large enough, n' will become a constant. Furthermore the unique operation is $\mathcal{O}(n \log n)$ with a small coefficient so this part is negligible compared to the reduce operation. All these factors result in the trend we observe in Figure 9. This indicates excellent scalability of PRIVGUARD in terms of the number of data capsules.

8 Related Work

Recently, there are some research efforts on bootstrapping privacy compliance in big data systems. Technically, the works in this area can be categorized into two directions - (1) summarize the issues in privacy regulations to guide deployment; (2) formalise privacy regulations in a strict programming language flavor; (3) enforce privacy policies in data processing systems. Our work falls into all three categories. In the following, we briefly describe these research works and discuss why these existing approaches do not fully solve the problems in our setting.

Issues in Deploying Privacy Regulations. Gruschka et al. [21] summarize privacy issues in GDPR. Renaud et al. [22] synthesize the GDPR requirements into a checklist-type format, derive a list of usability design guidelines and providing a usable and GDPR-compliant privacy policy template for the benefit of policy writers. Politou et al. [23] review all controversies around the new stringent definitions of consent revocation and the right to be forgotten in GDPR and evaluate existing methods, architectures and state-of-the-art technologies in terms of fulfilling the technical practicalities for the implementation and effective integration of the new requirements into current computing infrastructures. Tom et al. [24] present the current state of a model of the GDPR that provides a concise visual overview of the associations between entities defined in the legislation and their constraints. In this work, our research goal is to summarize and formalize general-purpose privacy principles and design a lightweight paradigm for easy deployment in heterogeneous data processing systems. As a result, these discussions can serve as a good guidance to our work but not actually solve the problem we aim to tackle.

Privacy Regulation Formalism. In [25], Hanson et al. present a data-purpose algebra that can be used to model these kinds of restrictions in various different domains. To formalise purpose restrictions in privacy policies, Tschantz et al. [26] provide a semantics using a formalism based on planning modeled using a modified version of Markov Decision Processes. Chowdhury [27] present a policy specification language based on a restricted subset of first order temporal logic (FOTL) which can capture the privacy requirements of HIPAA. Lam et al. [28] prove that for any privacy policy that conforms to patterns evident in HIPAA, there exists a finite representative hospital database that illustrates how the law applies in all possible hospitals. However, because of these works' specific focus on purpose restriction or HIPAA, the above two approaches do not generalize to other regulations like GDPR. Gerl et al. [29] introduce LPL, an extensible Layered Privacy Language that allows to express and enforce these new privacy properties such as personal privacy, user consent, data provenance, and retention management. Sen et al. introduce LEGALEASE [18], a language composed of (alternating) ALLOW and DENY clauses where each clause relaxes or constricts the enclosing clause. LEGALEASE is compositional and specifies formal semantics in attribute lattices. These characteristics are useful in general-purpose description of privacy regulations and are inherited in PRIVPOLICY. However, compared with LEGALEASE, PRIVPOLICY supports much more expressive attributes to represent abstract domains for static analysis which allows us to encode more complicated privacy regulations. Other work (e.g. Becker et al. [14]) focuses on the access control issues related to compliance with data privacy regulations, but such approaches do not restrict *how* the data is processed—a key component of recent regulations like GDPR.

Privacy Regulation Compliance Enforcement. Going beyond formalism of privacy regulations, recent research also explores techniques to enforce these formalised privacy regulations in real-world data processing systems. Chowdhury et al. [30] propose to use temporal model-checking for run-time monitoring of privacy policies. While Chowdhury demonstrates the effectiveness of this approach in online monitoring of privacy policies, it does not provide the capability of static analysis to decide if a analytic program satisfies a privacy policy and can only report privacy violation after it happens. Sen et al. [18] introduce GROK, a data inventory for Map-Reduce-like big data systems. Although working perfectly in Map-Reduce-like systems, GROK lacks adaptability to non-Map-Reduce-like data processing systems.

9 Conclusion & Future Work

In this paper, we have proposed the data capsule paradigm, a new paradigm for collecting, managing, and processing sensitive personal data. The data capsule paradigm has the potential to break down data silos and make data more useful, while at the same time reducing the prevalence of data privacy violations and making compliance with privacy regulations easier for organizations. We implemented PRIVGUARD, a reference platform for the new paradigm.

We are currently in the preliminary stages of a collaborative case study to apply the data capsule paradigm to enforce HIPAA in a medical study of menstrual data collected via mobile app. The goal of this study [31] and similar work [32, 33] is to demonstrate the use of mobile apps to assess menstrual health and fertility. Data capsules will allow study participants to submit their sensitive data in the context of a policy which protects its use. As part of this effort, we are in the process of encoding the requirements of HIPAA using PRIVPOLICY and applying PRIVGUARD to the analysis programs written by the study’s designers.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. This work was supported by DARPA & SPAWAR under contract N66001-15-C-4066. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

1. “The 18 biggest data breaches of the 21st century,” <https://www.csoonline.com/article/2130877/the-biggest-data-breaches-of-the-21st-century.html>, 2019, online; accessed 23 May 2019.
2. D. J. Solove and D. K. Citron, “Risk and anxiety: A theory of data-breach harms,” *Tex. L. Rev.*, vol. 96, p. 737, 2017.
3. “Insider threat 2018 report,” <https://www.ca.com/content/dam/ca/us/files/ebook/insider-threat-report.pdf>, 2019, online; accessed 23 May 2019.
4. L. E. Murdock, “The use and abuse of computerized information: Striking a balance between personal privacy interests and organizational information needs,” *Alb. L. Rev.*, vol. 44, p. 589, 1979.
5. “The eu general data protection regulation (gdpr),” <https://eugdpr.org/>, 2019, online; accessed 16 April 2019.
6. “California consumer privacy act (ccpa),” <https://www.caprivacy.org/>, 2019, online; accessed 16 April 2019.
7. “The family educational rights and privacy act of 1974 (ferpa),” <https://www.colorado.edu/registrar/students/records/ferpa>, 2019, online; accessed 16 April 2019.
8. “Health insurance portability and accountability act (hipaa),” <https://searchhealthit.techtarget.com/definition/HIPAA>, 2019, online; accessed 16 April 2019.
9. “Google keeps your data forever - unlocking the future transparency of your past,” <https://www.siliconvalleywatcher.com/google-keeps-your-data-forever--unlocking-the-future-transparency-of-your-past/>, 2019, online; accessed 30 May 2019.
10. “Extract, transform, load,” https://en.wikipedia.org/wiki/Extract,_transform,_load, 2019, online; accessed 30 May 2019.
11. E. F. Codd, “A relational model of data for large shared data banks,” *Communications of the ACM*, vol. 13, no. 6, pp. 377–387, 1970.
12. K. Chodorow, *MongoDB: the definitive guide: powerful and scalable data storage*. ” O’Reilly Media, Inc.”, 2013.
13. K. Shvachko, H. Kuang, S. Radia, R. Chansler *et al.*, “The hadoop distributed file system.” in *MSST*, vol. 10, 2010, pp. 1–10.
14. A. Lakshman and P. Malik, “Cassandra: a decentralized structured storage system,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 2, pp. 35–40, 2010.
15. J. Dean and S. Ghemawat, “Mapreduce: simplified data processing on large clusters,” *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
16. K. Shvachko, H. Kuang, S. Radia, R. Chansler *et al.*, “The hadoop distributed file system.” in *MSST*, vol. 10, 2010, pp. 1–10.
17. M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, “Spark: Cluster computing with working sets.” *HotCloud*, vol. 10, no. 10-10, p. 95, 2010.
18. S. Sen, S. Guha, A. Datta, S. K. Rajamani, J. Tsai, and J. M. Wing, “Bootstrapping privacy compliance in big data systems,” in *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 327–342.
19. “Formal concept analysis,” https://en.wikipedia.org/wiki/Formal_concept_analysis, 2019, online; accessed 30 May 2019.
20. F. Nielson, H. R. Nielson, and C. Hankin, *Principles of program analysis*. Springer, 2015.

21. N. Gruschka, V. Mavroeidis, K. Vishi, and M. Jensen, "Privacy issues and data protection in big data: A case study analysis under gdpr," in *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 2018, pp. 5027–5033.
22. K. Renaud and L. A. Shepherd, "How to make privacy policies both gdpr-compliant and usable," in *2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA)*. IEEE, 2018, pp. 1–8.
23. E. Politou, E. Alepis, and C. Patsakis, "Forgetting personal data and revoking consent under the gdpr: Challenges and proposed solutions," *Journal of Cybersecurity*, vol. 4, no. 1, p. tty001, 2018.
24. J. Tom, E. Sing, and R. Matulevičius, "Conceptual representation of the gdpr: Model and application directions," in *International Conference on Business Informatics Research*. Springer, 2018, pp. 18–28.
25. C. Hanson, T. Berners-Lee, L. Kagal, G. J. Sussman, and D. Weitzner, "Data-purpose algebra: Modeling data usage policies," in *Eighth IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'07)*. IEEE, 2007, pp. 173–177.
26. M. C. Tschantz, A. Datta, and J. M. Wing, "Formalizing and enforcing purpose restrictions in privacy policies," in *2012 IEEE Symposium on Security and Privacy*. IEEE, 2012, pp. 176–190.
27. O. Chowdhury, A. Gampe, J. Niu, J. von Ronne, J. Bennatt, A. Datta, L. Jia, and W. H. Winsborough, "Privacy promises that can be kept: a policy analysis method with application to the hipaa privacy rule," in *Proceedings of the 18th ACM symposium on Access control models and technologies*. ACM, 2013, pp. 3–14.
28. P. E. Lam, J. C. Mitchell, A. Scedrov, S. Sundaram, and F. Wang, "Declarative privacy policy: finite models and attribute-based encryption," in *Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium*. ACM, 2012, pp. 323–332.
29. A. Gerl, N. Bannani, H. Kosch, and L. Brunie, "Lpl, towards a gdpr-compliant privacy language: Formal definition and usage," in *Transactions on Large-Scale Data-and Knowledge-Centered Systems XXXVII*. Springer, 2018, pp. 41–80.
30. O. Chowdhury, L. Jia, D. Garg, and A. Datta, "Temporal mode-checking for runtime monitoring of privacy policies," in *International Conference on Computer Aided Verification*. Springer, 2014, pp. 131–149.
31. L. Symul, K. Wac, P. Hillard, and M. Salathe, "Assessment of menstrual health status and evolution through mobile apps for fertility awareness," *bioRxiv*, 2019. [Online]. Available: <https://www.biorxiv.org/content/early/2019/01/28/385054>
32. B. Liu, S. Shi, Y. Wu, D. Thomas, L. Symul, E. Pierson, and J. Leskovec, "Predicting pregnancy using large-scale data from a women's health tracking mobile application," *arXiv preprint arXiv:1812.02222*, 2018.
33. A. Alvergne, M. Vlajic Wheeler, and V. Hgqvist Tabor, "Do sexually transmitted infections exacerbate negative premenstrual symptoms? Insights from digital health," *Evolution, Medicine, and Public Health*, vol. 2018, no. 1, pp. 138–150, 07 2018. [Online]. Available: <https://doi.org/10.1093/emph/eoy018>