

Modular Verification of Differential Privacy in Probabilistic Higher-Order Separation Logic (including Appendix)

PHILIPP G. HASELWARTER, Aarhus University, Denmark

ALEJANDRO AGUIRRE, Aarhus University, Denmark

SIMON ODDERSHEDDE GREGERSEN, New York University, USA

JOSEPH TASSAROTTI, New York University, USA

LARS BIRKEDAL, Aarhus University, Denmark

Differential privacy is the standard method for privacy-preserving data analysis. The importance of having strong guarantees on the reliability of implementations of differentially private algorithms is widely recognized and has sparked fruitful research on formal methods. However, the design patterns and language features used in modern DP libraries as well as the classes of guarantees that the library designers wish to establish often fall outside of the scope of previous verification approaches.

We introduce a program logic suitable for verifying differentially private implementations written in complex, general-purpose programming languages. Our logic has first-class support for reasoning about privacy budgets as a separation logic resource. The expressiveness of the logic and the target language allow our approach to handle common programming patterns used in the implementation of libraries for differential privacy, such as privacy filters and caching. While previous work has focused on developing guarantees for programs written in domain-specific languages or for privacy mechanisms in isolation, our logic can reason modularly about primitives, higher-order combinators, and interactive algorithms.

We demonstrate the applicability of our approach by implementing a verified library of differential privacy mechanisms, including an online version of the Sparse Vector Technique, as well as a privacy filter inspired by the popular Python library OpenDP, which crucially relies on our ability to handle the combination of randomization, local state, and higher-order functions. We demonstrate that our specifications are general and reusable by instantiating them to verify clients of our library. All of our results have been foundationally verified in the Rocq Prover.

1 Introduction

Differential privacy [15, 16] (DP) is a collection of programming techniques to release aggregate information from a database while providing statistical guarantees about the privacy of individual user data. DP has been used widely in government and industrial applications to protect critical personal information (medical, financial, demographic, behavioral). The correctness of the implementations of DP is therefore crucial: logic- or implementation-level bugs can lead to catastrophic failure of the promised privacy guarantees.

The importance of having trustworthy implementations of DP is widely recognized. On the one hand, it has led to significant research in programming language and program verification. On the other, industrial developments of DP are routinely accompanied by pen-and-paper proofs of their claimed privacy properties. The OpenDP collaboration¹ even goes as far as to collect a proof for each element of the library, which are checked by a “privacy proof review board”.

DP provides strong *statistical* guarantees that the data contributed by an individual does not influence the result of the analysis by too much, and hence cannot be recovered by observing the outcome. This is achieved by adding a small amount of random noise to each step of data analysis that could leak private information. This way, even if the input data changes by a small amount

¹OpenDP was initiated as a collaboration between Harvard University and Microsoft.

Authors' Contact Information: [Philipp G. Haselwarter](#), Aarhus University, Denmark, pgh@cs.au.dk; [Alejandro Aguirre](#), Aarhus University, Denmark, alejandro@cs.au.dk; [Simon Oddershede Gregersen](#), New York University, USA, s.gregersen@nyu.edu; [Joseph Tassarotti](#), New York University, USA, jt4767@cs.nyu.edu; [Lars Birkedal](#), Aarhus University, Denmark, birkedal@cs.au.dk.

(say, the data of one user in a database), we can expect the output of the noisy program output to be similar, and an observer will not be able to tell whether the change comes from the noise or a change in inputs. The strength of the privacy guarantee of differentially private (dp) programs is governed by a parameter $\epsilon \in \mathbb{R}_{\geq 0}$, that controls the overall probability that any user's privacy is compromised; a lower ϵ means a more private program.

Managing the privacy budget. A central idea of programming with DP is to think of the parameter ϵ that controls the noise distribution as the “privacy budget”. To ensure that the whole program P , written as a sequence of computations $P = c_0; c_1; \dots; c_n$, is ϵ -differentially private (ϵ -dp), we can “allocate” some part ϵ_i of our budget to each part c_i of the program and check that the probability that c_i compromises privacy is controlled according to ϵ_i . So long as the sum of the error terms ϵ_i is below the global budget ϵ , the program as a whole is ϵ -dp. This principle is known as the sequential composition theorem of DP, and it justifies our budget intuition of the privacy parameter: ϵ represents a *resource* that can be *split* according to the structure of the program, and is *consumed* by computing noisy results from a database.

The most widely used deployments of DP are based on *dynamic* techniques for tracking the privacy budget via a trusted library such as OpenDP (written in Python, Rust) or Google's Differential Privacy library (C++, Go, Java). Rather than specifying exactly how much ϵ_i each part of a program consumes, the budget consumption of the completed computations $c_0; \dots; c_{i-1}$ is tracked, and a runtime check ensures that enough budget for c_i remains.

1.1 The Challenges of Verifying DP Frameworks

Most practical DP frameworks use “advanced” programming language features like higher-order functions (e.g., in the form of classes or function pointers) and dynamically allocated, local state (e.g., via class-private attributes). This is necessary to enable the *modular* construction of private programs through an API, where a library client does not have to concern themselves with the management of the privacy budget and correct application of noise to the results. To be able to reason compositionally, privacy proofs of a library API should likewise support *modular specifications*. We consider three representative challenges that arise from reasoning about libraries for DP:

- (i) encapsulating dynamic, fine-grained budget accounting,
- (ii) interactive or “online” data analysis,
- (iii) budget minimization via caching.

As we shall see shortly, all three of these techniques rely on stateful randomized higher-order functions to implement features that are essential to the modular construction of dp programs. To verify modular API specifications we must thus support higher-order reasoning about local state.

We illustrate the programming patterns via Python code snippets inspired by OpenDP, but the challenges also arise in, e.g., the C++ implementation of Google's Differential Privacy library.² Although the verification approach developed in this paper is for an ML-like core language (§2.3) rather than Python, it can faithfully represent the salient aspects of these Python programs.

1.1.1 Dynamic budget accounting. An essential functionality of these libraries is that they offer *privacy filters* [36] that encapsulate the intricate reasoning about the privacy budget in a core, trusted API. A privacy filter tracks the privacy cost that a program has incurred up to the current execution point $c_0; \dots; c_{i-1}$, and only executes c_i if enough budget remains. The example implementation of a simplified PrivacyFilter in Figure 1 is initialized with some budget and provides a single `try_run(cost, f)` method, which only runs `f` if there is indeed at least enough

²See, e.g., <https://github.com/google/differential-privacy/blob/0a3b05a65f7ed8de/cc/accounting/accountant.cc#L46>

budget left to cover its cost. So long as each mechanism f correctly reports its budget consumption (cost), the filter will ensure that the total privacy cost will never exceed the global budget.

```
class PrivacyFilter:
    def __init__(self, budget):
        self.budget_left = budget

    def try_run(self, cost, f):
        if self.budget_left < cost :
            return None
        self.budget_left -= cost
        return f()
```

Fig. 1. PrivacyFilter ensures that no client can exceed the privacy budget.

Centralized dynamic budget management simplifies the privacy analysis of programs and allows for scaling to industrial applications. Importantly, dynamically computed bounds enable a tighter analysis of the privacy cost compared to static type checks [18, 27]. Recall that the usual sequential composition theorem of DP requires that all ϵ_i be chosen upfront. The adaptive composition theorem [36] for DP lifts this restriction and allows the analyst to *adaptively* chose each ϵ_i depending on the results of previous (private) computations. Despite the successes of type- and program-logic-based analyses of DP, none of the existing systems can be used to specify and verify the correctness of implementations of privacy filters.

1.1.2 Interactive data analysis. *Interactivity* is central to real-world data analysis under DP. Given a new dataset, a data analyst does not just issue a static set of queries. Instead, they may compute some summary statistics to explore what ranges of values or categories of responses are of interest, and construct further queries based on those observations. This interactive or “online” style of analysis is crucial for practical utility [12, 16, 25]. A standard way to represent interactive computations is as streams: given a stream of queries, a private interactive mechanism should produce a stream of results.

The `AboveThreshold` mechanism in Figure 2 constructs a stream of booleans from a stream of queries as an iterator. The n -th boolean indicates whether q_n was above the threshold T , after suitable Laplace noise was added. Crucially, the queries themselves are provided as an iterator rather than as a precomputed list, and q_n is only evaluated in the n -th call to `__next__`. The `AboveThreshold` mechanism is ϵ -dp because the stream stops producing results after the first time `True` is returned; this is a standard result in DP but challenging to prove formally because it does *not* follow simply from the sequential composition theorem.

```
class AboveThreshold:
    def __init__(self, eps, T, queries, db):
        self.eps = eps
        self.T = Laplace(T, eps/2)
        self.queries = queries
        self.db = db
        self.halted = False

    def __iter__(self):
        return self

    def __next__(self):
        if self.halted:
            raise StopIteration
        q = next(self.queries)
        v = Laplace(q(self.db), self.eps/4)
        b = (self.T <= v)
        self.halted = b
        return b
```

Fig. 2. The classic Above Threshold interactive mechanism.

Any realistic verification framework for DP must therefore capture not only isolated, static mechanisms but also their behavior as interactive, stateful processes that maintain internal state and privacy budget across calls.

1.1.3 Budget minimization via caching. Certain results (e.g., the number of non-zero entries) have to be calculated many times over as part of different analysis passes on a dataset. This can be wasteful in practical applications, since each private computation consumes some of the privacy budget, even if the same result has already been computed elsewhere in the program. Caching provides a solution to this problem: if the noisy results of queries are cached, then a repeated query can simply reuse the prior noisy result without incurring any privacy cost. Recent implementations of privacy frameworks have studied increasingly sophisticated caching strategies for DP workloads [25, 26, 30, 34], but even simple implementations of caching via memoization can lead to substantial savings in privacy budget in practical workloads [34][25, Fig. 3].

The implementation of memoization `mk_query_cache` in Figure 3 locally allocates a cache associated to a mechanism `add_noise` and a dataset `db` and returns a closure `f` which can be used subsequently to privately evaluate queries on `db`. Although it is a simple general-purpose caching mechanism, studying it formally requires reasoning about local state and higher-order functions. This places it outside of the scope of existing systems.

```
def mk_query_cache(add_noise, db):
    cache = {}
    def f(query):
        if query in cache:
            return cache[query]
        v = add_noise(query(db))
        cache[query] = v
        return v
    return f
```

Fig. 3. A generic caching mechanism.

1.2 Formal Guarantees for Differential Privacy

DP lends itself well to the analysis via standard PL methods due to its compositional nature. We distinguish two main classes of approaches. Type-based approaches [1, 8, 13, 18, 27, 28, 33, 35, 40–42] use a static typing discipline to ensure that the programs accepted by the system are dp. These approaches benefit from high degrees of automation. Most type-based approaches, however, do not handle mutable state. More generally, the complex nature of the type systems required (dependent-, linear-, contextual-, or refinement types, or combinations thereof), hinders their integration with mainstream programming languages.

On the other hand, relational probabilistic Hoare logics such as apRHL [6, 7, 9], are more expressive but require more user effort in the form of manual or interactive proofs. These methods works particularly well when the program logic is defined with respect to a relatively simple denotational semantics of programs as subdistributions. This, however, generally restricts the applicability of the method to first-order programs. HO-RPL [2] extends the previous approach to support higher-order functions by constructing a more sophisticated semantic model, but general recursion or dynamic allocation and higher-order state are still unsupported. In a similar vein, some projects verified the privacy of sampling algorithms and simple mechanisms directly in the denotational semantics of programs [14, 39]. In summary, these approaches work well for the verification of algorithms in isolation, but do not focus on supporting *modular* program verification.

1.3 Modular Verification for DP Libraries in Clutch-DP

To address the challenges arising from the verification of DP libraries, we introduce the relational approximate probabilistic higher-order separation logic Clutch-DP. Clutch-DP supports higher-order functions such as `PrivacyFilter` and `mk_query_cache` via quantification over specifications, and

references storing closures as they occur in the implementation of iterators (e.g., `AboveThreshold`) via impredicative invariants à la Iris [24]. Building on previous relational separation logics [20, 21], we internalize the privacy budget of DP in Clutch-DP as a first-class separation logic resource. These *privacy credits* can be tracked in invariants and interact with the heap and with higher-order functions that consume a privacy budget in flexible ways. We prove that the usual rules for (sequential) relational separation logic [17] are sound in the probabilistic setting. To reason about privacy of primitives, we prove novel relational sampling rules for the Laplacian; in particular, we internalize a proof technique based on choice couplings [4] as a logical rule, which, in particular, is used in the privacy proof of `AboveThreshold`. We apply Clutch-DP to a number of case studies inspired by the challenge problems described above. Our specifications of, e.g., `AboveThreshold` are compositional and reusable: we derive the privacy of clients and more complex mechanisms from the specifications of the building blocks, without referring to implementation details.

Contributions. To summarize, we make the following contributions.

- (1) *A higher-order separation logic for DP* which internalizes privacy credits as first-class, composable logical resources and which supports heap allocation and higher-order closures.
- (2) *New probabilistic sampling rules*, including a Laplacian rule that enables selective recovery of privacy credits, which can be used to verify examples whose privacy analysis goes beyond composition theorems.
- (3) *A library of reusable, abstract specifications* for common DP primitives (Laplace, `AboveThreshold`, Sparse Vector Technique, Report-Noisy-Max, Privacy Filters, Caching via Memoization) that cleanly separate mechanism proof from client reasoning.
- (4) *Client case studies* that highlight both expressiveness and reusability of specifications and that demonstrate that our approach successfully addresses the challenges outlined above (dynamic budget accounting, interactive data analysis, and budget minimization via caching).
- (5) *Mechanized proofs*: a foundational formalization of the logic and all case studies in the Rocq Prover together with an adequacy theorem connecting our logic to standard DP.

Outline. In §2 we briefly recall the basic notions of DP and define the RandML programming language used in the remainder of the paper. In §3 we define the Clutch-DP logic and explain how it relates to DP. In §4 we illustrate reasoning in Clutch-DP via a number of case studies addressing the challenges set out in §1.1. The soundness of Clutch-DP is addressed in §5. Finally, in §6 we survey related work, before concluding in §7.

2 Differential Privacy and Programming Language Preliminaries

We briefly recall the elements of probability theory §2.1 and DP §2.2 we need to refer to, and formally define the programming language §2.3 used throughout the paper.

2.1 Probability Theory

Since we do not assume a priori that all programs we study terminate, we allow programs to “lose mass” on diverging runs and define the operational semantics using probability *sub*-distributions.

DEFINITION 2.1. *A discrete subdistribution (henceforth simply distribution) on a countable set A is a function $\mu : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \mu(a) \leq 1$. The distributions on A are denoted by $\mathcal{D}(A)$.*

In §2.3, we will define the operational semantics of RandML in terms of the distribution monad.

LEMMA 2.2. *The discrete distribution monad induced by \mathcal{D} has operations*

$$\begin{array}{ll} \text{ret} : A \rightarrow \mathcal{D}(A) & \text{bind} : (A \rightarrow \mathcal{D}(B)) \rightarrow \mathcal{D}(A) \rightarrow \mathcal{D}(B) \\ \text{ret}(a)(a') \triangleq \begin{cases} 1 & \text{if } a = a' \\ 0 & \text{otherwise} \end{cases} & \text{bind}(f, \mu)(b) \triangleq \sum_{a \in A} \mu(a) \cdot f(a)(b) \end{array}$$

We write $(\mu \gg f)$ for $\text{bind}(f, \mu)$.

2.2 Differential Privacy

For background information on DP see, e.g., Dwork and Roth [16] and Cowan et al. [12].

Defining privacy. The definition of DP captures the intuition that it is hard to reconstruct information about any individual in a database from the output of a dp program.

DEFINITION 2.3. A function $f : \mathcal{DB} \rightarrow \mathcal{D}(X)$ is (ϵ, δ) -differentially private (short: “ f is (ϵ, δ) -dp”) if $\Pr_{f(x)}[\phi] \leq e^\epsilon \cdot \Pr_{f(y)}[\phi] + \delta$ for all adjacent $x, y : \mathcal{DB}$ and all predicates $\phi \subseteq X$. If f is $(\epsilon, 0)$ -dp we simply say that f is ϵ -dp.

A function is thus (ϵ, δ) -dp if it amplifies the probability of *any* observation ϕ by at most e^ϵ , or by more than that with probability at most δ . A small value for ϵ and δ thus means strong privacy guarantees. The definition of DP is parametrized by a type of databases DB and an adjacency relation. Clutch-DP can work with any database type and adjacency relation, but a common choice is to think of a database as a list of rows where each row is a tuple of a fixed size. If the type of databases comes with a notion of distance $d_{\text{DB}} : \text{DB} \rightarrow \mathbb{R}_{\geq 0}$ we say that two databases x, y are adjacent if $d_{\text{DB}}(x, y) \leq 1$. For instance, if $d_{\text{DB}}(x, y)$ is the number of rows where x and y differ, then adjacency means they only differ in one row.

Adding noise: the Laplace mechanism. To prove that any program is dp, we need primitives that add random noise. The prototypical example of a noise mechanism that achieves DP is the Laplacian distribution. However, since the Laplacian is a continuous distribution on \mathbb{R} , an implementation of a Laplacian sampler would have to work with exact real arithmetic, since implementations using floats lead to well-known privacy bugs [31].

We therefore work with the *discrete Laplacian*, the distribution on \mathbb{Z} obtained by discretizing the continuous Laplacian. The discrete Laplacian with scale parameter ε and mean m has as probability mass function³ [22, 23]:

$$\mathcal{L}_\varepsilon^m(v) \triangleq \frac{1}{W} \cdot e^{-\varepsilon \cdot |v - m|} \quad \text{where } W \triangleq \sum_{z \in \mathcal{Z}} e^{-\varepsilon \cdot |z|} \quad (1)$$

Interpreting adjacency on \mathbb{Z} as being at distance at most c , we have the following result.

THEOREM 2.4 ([19]). $(\lambda m. \mathcal{L}_{\varepsilon/c}^m) : \mathbb{Z} \rightarrow \mathcal{D}(\mathbb{Z})$ is ε -dp.

As a direct consequence, given any function $f : \mathcal{D}\mathcal{B} \rightarrow \mathbb{Z}$ such that $|f(x) - f(y)| \leq 1$ for adjacent datasets $x, y : \mathcal{D}\mathcal{B}$, the function $(\lambda m. \mathcal{L}_\varepsilon^m) \circ f : \mathcal{D}\mathcal{B} \rightarrow \mathcal{D}(\mathbb{Z})$ is ε -dp.

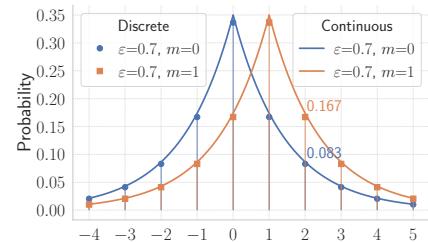


Fig. 4. Continuous and discrete Laplacian, with $\mathcal{L}_{0,7}^m(v)$ for $v=2$ and $m \in \{0, 1\}$, demonstrating 0.7-DP: $\mathcal{L}_{0,7}^1(2) \leq e^{0.7} \cdot \mathcal{L}_{0,7}^0(2)$.

³NB: The weight W is a geometric series over \mathbb{Z} (hence why \mathcal{L}_ε is also called the (two-sided) ε -geometric) with the closed form $\mathcal{L}_\varepsilon^m(v) = \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \cdot e^{-\varepsilon \cdot |v - m|}$ [11]. We adopt the convention that $\mathcal{L}_\varepsilon^m = \text{ret}(m)$ if $\varepsilon \leq 0$.

Composing DP. Differentially private functions satisfy a number of composition laws that can help structure and simplify the privacy verification of larger programs.

LEMMA 2.5. *DP is stable under post-processing: if $f : A \rightarrow \mathcal{D}(B)$ is (ϵ, δ) -dp then for any $g : B \rightarrow C$, the function $\lambda x. (f(x) \gg \lambda y. \text{ret } g(y)) : A \rightarrow \mathcal{D}(C)$ is (ϵ, δ) -dp.*

When two dp functions are composed sequentially, their privacy parameters add up. Due to the post-processing property this holds even if the later computations can see the results of earlier computations. We only show the case for two functions, but the lemma directly generalizes to arbitrary k -fold composition for $k \in \mathbb{N}$.

LEMMA 2.6 (SEQUENTIAL COMPOSITION). *Let $f : \mathcal{D}(B) \rightarrow \mathcal{D}(C)$ be (ϵ_1, δ_1) -dp and let g be a function $g : \mathcal{D}(B) \times B \rightarrow \mathcal{D}(C)$ such that $(\lambda x. g(x, b)) : \mathcal{D}(B) \rightarrow \mathcal{D}(C)$ is (ϵ_2, δ_2) -dp for all $b \in B$. Then $\lambda x. (f(x) \gg \lambda b. g(x, b))$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -dp.*

Another useful composition law holds for functions which increase the distance between inputs by at most a fixed amount c in the following sense.

DEFINITION 2.7. *We say that $f : A \rightarrow B$ is c -sensitive (also: “ c -stable”) if for all $x, y \in A$, the bound $d_B(f(x), f(y)) \leq c \cdot d_A(x, y)$ holds, where the distances d_A, d_B are taken with respect to a metric space structure on A and B .*

Note that if $f : A \rightarrow B$ is c -sensitive and $g : B \rightarrow C$ is d -sensitive then $g \circ f$ is $(c \cdot d)$ -sensitive. The following *metric composition law* then generalizes the remark following [Theorem 2.4](#).

LEMMA 2.8. *If $f : \mathcal{D}(B) \rightarrow \mathcal{D}(C)$ is c -sensitive then $(\lambda m. \mathcal{L}_{\epsilon/c}^m) \circ f : \mathcal{D}(B) \rightarrow \mathcal{D}(C)$ is ϵ -dp.*

2.3 The Language: Randomized ML

The RandML language that we consider is an ML-like language with higher-order recursive functions and higher-order state that we extend with an operator `Laplace` $a b m$ that samples from the Laplacian with scale (a/b) and mean m . The syntax is defined by the grammar below.

$$\begin{aligned}
 v, w \in \text{Val} &::= z \in \mathbb{Z} \mid b \in \mathbb{B} \mid () \mid \ell \in \text{Loc} \mid \text{rec } f x = e \mid (v, w) \mid \text{inl } v \mid \text{inr } v \\
 e \in \text{Expr} &::= v \mid x \mid \text{rec } f x = e \mid e_1 e_2 \mid e_1 + e_2 \mid e_1 - e_2 \mid \dots \mid \text{if } e \text{ then } e_1 \text{ else } e_2 \mid (e_1, e_2) \mid \text{fst } e \mid \dots \\
 &\quad \text{ref } e_1 \mid !e \mid e_1 \leftarrow e_2 \mid \text{Laplace } e_1 e_2 e_3 \mid \dots \\
 K \in \text{Ectx} &::= - \mid e K \mid K v \mid \text{ref } K \mid !K \mid e \leftarrow K \mid K \leftarrow v \mid \text{Laplace } e_1 e_2 K \mid \text{Laplace } e K v \mid \dots \\
 \sigma \in \text{State} &\triangleq \text{Loc} \xrightarrow{\text{fin}} \text{Val} \quad \rho \in \text{Cfg} \triangleq \text{Expr} \times \text{State}
 \end{aligned}$$

In RandML, `ref` e_1 allocates a new reference containing the value returned by e_1 , `!e` dereferences the location e evaluates to, and $e_1 \leftarrow e_2$ evaluates e_2 and assigns the result to the location that e_1 evaluates to. We may refer to a recursive function value `rec` $f x = e$ by its local name f . The heap is represented as a (partial) finite map from locations to values, and evaluation happens right to left as indicated by the evaluation context grammar `Ectx`.

The expression `Laplace` $a b m$ samples from the discrete Laplacian with scale $\epsilon = a/b$ and mean m . To avoid unnecessary complications with adding real numbers to the programming language, we require the scale ϵ to be a rational number. Formally, in RandML, `Laplace` takes three integers as input, but to keep our notation free from clutter we will simply write `Laplace` ϵm instead of `Laplace` $a b m$ with $\epsilon = a/b$.

Operational Semantics. Program execution is defined by iterating $\text{step} : \text{Cfg} \rightarrow \mathcal{D}(\text{Cfg})$, where $\text{step}(\rho)$ is the distribution induced by the single step reduction of the configuration ρ . The semantics is mostly standard. We first define head reduction and then lift it to reduction in

an evaluation context K . All non-probabilistic constructs reduce deterministically as usual, e.g., $\text{step}((\lambda x. e) v, \sigma) = \text{ret}(e[v/x], \sigma)$. We write $e \rightsquigarrow e'$ if the evaluation is deterministic and holds independently of the state, e.g., $(\lambda x. e) v \rightsquigarrow e[v/x]$ and $\text{fst}(v_1, v_2) \rightsquigarrow v_1$. The sampling operator $\text{Laplace } a b m$ reduces according to the Laplacian with scale a/b and mean m , i.e.,

$$\text{step}(\text{Laplace } a b m, \sigma)(v, \sigma) \triangleq \begin{cases} \mathcal{L}_{a/b}^m(v) & \text{for } v \in \mathbb{Z}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

With the single step reduction $\text{step} : \text{Cfg} \rightarrow \mathcal{D}(\text{Cfg})$ defined, we next define a step-stratified execution probability $\text{exec}_n : \text{Cfg} \rightarrow \mathcal{D}(\text{Val})$ by induction on n :

$$\begin{aligned} \text{exec}_0(e, \sigma)(v) &\triangleq \begin{cases} 1 & \text{if } e \in \text{Val} \wedge e = v, \\ 0 & \text{otherwise.} \end{cases} \\ \text{exec}_{n+1}(e, \sigma)(v) &\triangleq \begin{cases} 1 & \text{if } e \in \text{Val} \wedge e = v, \\ \sum_{(e', \sigma') \in \text{Expr} \times \text{State}} \text{step}(e, \sigma)(e', \sigma') \cdot \text{exec}_n(e', \sigma')(v) & \text{otherwise.} \end{cases} \end{aligned}$$

That is, $\text{exec}_n(e, \sigma)(v)$ is the probability of stepping from the configuration (e, σ) to a value v in at most n steps. The probability that an execution, starting from configuration ρ , reaches a value v is taken as the limit of its stratified approximations, which exists by monotonicity and boundedness:

$$\text{exec}(\rho)(v) \triangleq \lim_{n \rightarrow \infty} \text{exec}_n(\rho)(v)$$

The interpretation of programs as distributions induces a natural notion of (ϵ, δ) -DP for RandML programs. Concretely, if $f \in \text{Expr}$ is a RandML function then for adjacent databases x, y it should be the case that for all states σ ,

$$\text{Pr}_{\text{exec}(f \text{ inj}(x), \sigma)}[\phi] \leq e^\epsilon \cdot \text{Pr}_{\text{exec}(f \text{ inj}(y), \sigma)}[\phi] + \delta \quad (3)$$

where $\text{inj} : \text{DB} \rightarrow \text{Val}$ embeds the type of databases into RandML values (we usually omit inj).

Note that in particular $f = (\lambda x. \text{Laplace } \epsilon x)$ is ϵ -dp in the sense of (3) by the definition of the operational semantics (2) and [Theorem 2.4](#).

3 Program Logic

In this section, we introduce the Clutch-DP logic, the soundness theorem of Clutch-DP, which connects it to DP, and the new logical connectives and rules pertaining to the privacy budget reasoning. On the surface, our logic looks similar to Approxix [21], but we remark that the underlying model is different, and that it is designed to support our novel reasoning principles for DP.

The logical connectives. Clutch-DP is built on top of the Iris separation logic framework [24] and inherits many of Iris's logical connectives, a selection of which is shown below. Most of the propositions are standard, such as separating conjunction $P * Q$ and separating implication $P \multimap Q$.

$$\begin{aligned} P, Q \in i\text{Prop} ::= & \text{True} \mid \text{False} \mid P \wedge Q \mid P \vee Q \mid P \Rightarrow Q \mid \forall x. P \mid \exists x. P \mid P * Q \mid P \multimap Q \mid \\ & \ell \mapsto v \mid \ell \mapsto_s v \mid \mathbb{F}^\times(\epsilon) \mid \mathbb{F}^+(\delta) \mid \{P\} e \preceq e' \{v v'. Q\} \mid \dots \end{aligned}$$

Since DP is a relational property, Clutch-DP is a relational program logic that proves properties about the execution of two programs e and e' . The central logical connective capturing the relationship between e and e' is the *Hoare quadruple*:

$$\{P\} e \preceq e' \{v v'. Q\} \quad (4)$$

Intuitively, this quadruple asserts that under the precondition P , if the two programs e and e' evaluate to results v and v' respectively, then the postcondition Q holds. Since the pre- and postcondition

range over arbitrary Clutch-DP assertions, they can refer to the state or contain nested Hoare quadruples, which will be useful to write specifications about higher-order functions.

Logical assertions about the heap come in two versions, one for e and one for e' . The heap points-to assertion that denotes ownership of location ℓ for the left-hand side program e is written as $\ell \mapsto v$, and $\ell' \mapsto_s v'$ denotes ownership of ℓ' for the right-hand side program e' .

Clutch-DP defines two kinds of resources that are novel with respect to standard Iris and which are inspired by Approxix [21]. The proposition $\mathbb{F}^{\times}(\varepsilon)$ asserts ownership of e *multiplicative privacy credits* and $\mathbb{F}^{+}(\delta)$ similarly asserts ownership of δ *additive privacy credits*. These credits are a logical representation of the privacy budget of DP. Just as reasoning about the physical state of a program is logically captured by the points-to connective, so is reasoning about the privacy budget (ε, δ) logically expressed via privacy credits.⁴

Internalizing privacy. We now have all of the ingredients to define DP internal to Clutch-DP.

DEFINITION 3.1. *A RandML function f is internally (ε, δ) -dp if the following quadruple holds:*

$$\forall db \ db'. \{Adj(db, db') * \mathbb{F}^{\times}(\varepsilon) * \mathbb{F}^{+}(\delta)\} f db \preceq f db' \{v, v'. v = v'\} \quad (5)$$

As with the usual (“external”) definition of DP, this notion is parametrized by an internal adjacency relation $Adj(db, db')$ on inputs. For instance, we could represent the databases as lists, where each element represents one entry, and say two databases are adjacent if they differ only in one entry (when considered as multisets). We abbreviate the Hoare quadruple (5) as (ε, δ) -iDP(f) when the relation is clear from context, and simply ε -iDP(f) for the case where $\delta = 0$.

The precise meaning of Hoare quadruples and privacy credits can be understood through the following theorem which connects Clutch-DP to DP in the sense of (3).

THEOREM 3.2 (SOUNDNESS). *If f is internally (ε, δ) -dp then f is also externally (ε, δ) -dp.*

This theorem is a consequence of the adequacy of the semantic model of Clutch-DP (see §5).

3.1 Relational Separation Logic Rules

The standard rules of relational separation logic as used in, e.g., [17] are also available in Clutch-DP. For instance, the **FRAME** rule enables local reasoning by framing out R , and the **BIND** rule allows us to focus on sub-programs in evaluation contexts. The usual load and store rules for heap locations come in a left- and a right version, requiring ownership of the corresponding points-to connective.

$$\begin{array}{c} \frac{\{P\} e \preceq e' \{Q\}}{\{P * R\} e \preceq e' \{Q * R\}} \text{ FRAME} \quad \frac{\{P\} e \preceq e' \{v v'. R\} \quad \forall v v'. \{R\} K[v] \preceq K'[v'] \{Q\}}{\{P\} K[e] \preceq K'[e'] \{Q\}} \text{ BIND} \\ \frac{\{\ell \mapsto w\} w \preceq e' \{Q\}}{\{\ell \mapsto w\} !\ell \preceq e' \{Q\}} \text{ LOAD-L} \quad \frac{\{\ell' \mapsto_s w\} e \preceq w \{Q\}}{\{\ell' \mapsto_s w\} e \preceq !\ell' \{Q\}} \text{ LOAD-R} \quad \frac{\{\ell \mapsto w\} () \preceq e' \{Q\}}{\{\ell \mapsto v\} \ell \leftarrow w \preceq e' \{Q\}} \text{ STORE-L} \end{array}$$

Fig. 5. Excerpt of the non-probabilistic rules of Clutch-DP.

It is worth noting that the rules in Figure 5 do *not* mention distributions, despite the fact that the operational semantics of RandML is probabilistic in general and references can, for instance, store randomly sampled values. Clutch-DP thus provides a convenient basis for program verification where reasoning about privacy is integrated as an *orthogonal* feature while preserving the familiar rules of relational higher-order separation logic.

⁴Contrary to heaps, the privacy budget pertains to a pair of program executions rather than to the left or right program, and hence there is no left- and right-hand version of credits; instead, they are shared.

3.2 Privacy Credit Laws

The privacy credit resources logically track two non-negative real numbers ε and δ corresponding to the privacy budget. In light of this intuition, one would expect that they support laws such as sequential composition (Lemma 2.6). Indeed, we can derive an internal version of this law from the primitive rules pertaining to privacy credits show in Figure 6 together with the structural rules for RandML in Figure 5.

$$\mathbb{F}^{\times}(\varepsilon_1 + \varepsilon_2) \dashv\vdash \mathbb{F}^{\times}(\varepsilon_1) * \mathbb{F}^{\times}(\varepsilon_2) \quad \mathbb{F}^{+}(\delta_1 + \delta_2) \dashv\vdash \mathbb{F}^{+}(\delta_1) * \mathbb{F}^{+}(\delta_2) \quad \mathbb{F}^{+}(1) \vdash \text{False}$$

Fig. 6. Privacy credit laws.

Just as in Lemma 2.6, both multiplicative and additive privacy credits can be split into their summands⁵ (and recombined). The last rule states that from 1 additive error credit, we can derive False and hence anything. This makes sense if we recall that δ models a bound on an inequality between *probabilities* in the definition of DP (Def. 2.3), since any probability is bounded by 1, and hence any program is trivially $(\varepsilon, 1)$ -dp.

The fact that the privacy budget is treated like any other separation logic resource and being able to split the budget enables *flexible* reasoning about privacy. Privacy credits are animated by the rules that govern their interaction with the operational semantics of RandML via sampling.

3.3 Rules for Sampling Noise

Privacy credits are consumed by the rules that reason about sampling operations. We introduce a rule **LAPLACE-SHIFT** to reason about a pair of Laplacians, with the same parameter ε and means m, m' . By setting $k = 0$ (and thus, $|m - m'| \leq c$), the rule lets us spend $\mathbb{F}^{\times}(c \cdot \varepsilon)$ to ensure that both Laplacians return the same result. If, instead, $k > 0$ and $c = 0$ (and thus, $|m - m'| = k$), the rule lets us conclude, without consuming privacy, that we will get two samples at distance k . The rule combines the two reasoning principles into one:

$$\frac{|k + m - m'| \leq c}{\{\mathbb{F}^{\times}(c \cdot \varepsilon)\} \text{ Laplace } \varepsilon m \precsim \text{Laplace } \varepsilon m' \{z \ z' \cdot z' = z + k\}} \text{ LAPLACE-SHIFT}$$

A consequence of **LAPLACE-SHIFT** (when $k = 0$) is that **Laplace** is internally private, *i.e.*, the Clutch-DP statement $\varepsilon\text{-iDP}(\lambda x. \text{Laplace } \varepsilon x)$ is derivable. This internalizes DP of the Laplacian (Thm. 2.4).

Finally, we have the following novel rule for the Laplacian which can be used to *recover* some privacy credits. The idea is to partition the results into two groups of outcomes depending on some threshold T . If both Laplacians sample a result above T (in fact, above $T + 1$ for the right-hand side) the privacy credits are consumed. If, on the other hand, both results remain below their respective thresholds, then the privacy credits can be recovered. This yields a useful reasoning principle for the Laplacian that is applied, *e.g.*, in the verification of the *AboveThreshold* mechanism (§4.1.1).

$$\frac{|m - m'| \leq 1 \quad T \in \mathbb{Z}}{\{\mathbb{F}^{\times}(2\varepsilon)\} \text{ Laplace } \varepsilon m \precsim \text{Laplace } \varepsilon m' \left\{ z \ z' \cdot \begin{array}{l} (T \leq z \wedge T + 1 \leq z') \vee \\ (z < T \wedge z' < T + 1 * \mathbb{F}^{\times}(2\varepsilon)) \end{array} \right\}} \text{ LAPLACE-CHOICE}$$

Intuitively this rule is sound because it partitions the outcomes into disjoint events such that the total privacy budget for the two is bounded by the initial privacy budget. Despite its conceptual simplicity, constructing a sound model that validates **LAPLACE-CHOICE** required substantial new insights in the form of a new composition theorem (see Theorem A.5).

⁵NB: Both kinds of privacy credit are split into separating conjunctions by using addition, not multiplication, as the resource algebra operation. The name “multiplicative privacy credit” for $\mathbb{F}^{\times}(\varepsilon)$ derives from the multiplicative factor e^{ε} in Def. 2.3.

3.4 Internal Composition Laws

Working in a program logic makes it simple to internalize the various composition laws for DP, as presented in §2. First, we remark that the internal version of post-processing holds. This is stated in our logic through the following lemma:

LEMMA 3.3 (INTERNAL POST-PROCESSING). *Let f be internally (ϵ, δ) -dp and assume g is “safe to execute” in the sense that $\forall w. \{\text{True}\} g w \preceq g w \{v v'. v = v'\}$ holds. Then $g \circ f$ is internally (ϵ, δ) -dp.*

This lemma can be proven entirely within the logic as an immediate consequence of the **BIND** rule. We can also recast the notion of sensitivity (Def. 2.7) in Clutch-DP.

DEFINITION 3.4. *A RandML function f is internally c -sensitive ($c\text{-sens}(f)$) if the following holds:*

$$\forall x y : A. \{\text{True}\} f x \preceq f y \{v v'. d_B(v, v') \leq c \cdot d_A(x, y)\}$$

As before, the definition is parametrized by two distances at types A and B , which can then be internalized as distances between values. The internal equivalent of metric composition then follows.

LEMMA 3.5. $\forall f c \epsilon. c\text{-sens}(f) \rightarrow \epsilon\text{-iDP}(\lambda x. \text{Laplace } (\epsilon/c) (f x))$ is derivable in Clutch-DP.

4 Reusable Specs for Privacy Mechanisms

Privacy *mechanisms* are the primitive building blocks of DP. Mechanisms sample the appropriate noise for a given data processing task. We give specifications for some widely used mechanisms and demonstrate that clients can be verified based on these abstract specs.

Our first example will be the Above Threshold (AT) mechanism, and we will see how to (1) prove that it satisfies an abstract specification capturing its privacy, (2) build the Sparse Vector Technique (SVT) from it, and (3) use AT to calculate the clipping bounds required to privately compute averages over a dataset. Both (2) and (3) use the same abstract specification of AT.

Besides the case studies presented hereafter, we also verified the privacy of the Report Noisy Max mechanism, which is of note because—as with AT—its privacy does not follow from composition laws but requires careful manipulation of the privacy budget and the use of budget- and state-dependent invariants. Details can be found in §B. We will present the most interesting ideas for each proof. Full proofs from the rules of Clutch-DP are available in our Rocq formalization.

4.1 Sparse Vector Technique

Suppose we have an incoming sequence of (1-sensitive) queries on a database and a fixed privacy budget. The Sparse Vector Technique (SVT) allows us to fix a threshold T and release, in a private manner, whether the result of each query exceeds T or not. The benefit of SVT is that one only has to spend privacy budget on the “successful” queries which do indeed exceed T and are released; the results of the queries that do not exceed the threshold can be computed (and discarded) without incurring a privacy cost. We can thus set in advance a maximum number N of successful queries to be released, and keep answering incoming queries interactively until N is reached. The SVT is usually implemented in terms of the Above Threshold mechanism (AT), which finds a *single* query above T . SVT then simply runs N iterations of AT. The privacy cost of SVT is N times the cost of finding one query above T .

The SVT is of interest for verification because several buggy privacy proofs have been published (see Lyu et al. [29] for a survey). As the survey explains, the SVT is particularly interesting in the interactive setting; in the non-interactive setting, one can use the Exponential Mechanism instead and get more accurate results. We explore two subtleties of SVT: how to perform the privacy analysis of AT, and how to build an interactive algorithm out of AT. The privacy analysis of AT

(and SVT) is challenging because it requires fine-grained reasoning about the privacy budget that *cannot* be justified by the sequential composition law alone.

4.1.1 Above Threshold. The Above Threshold mechanism can be used to evaluate queries on a database until one query returns a result that exceeds a specified threshold T . The implementation of `above_threshold` in Figure 7 initializes the noisy threshold \hat{T} and returns a function to run queries interactively. This function receives a query, computes its result, adds additional noise to it and checks whether the noisy result exceeds \hat{T} .

```
let above_threshold ε T = let  $\hat{T} = \text{Laplace}(\epsilon / 2) T$  in
  let  $f q db = \text{let } x = q db \text{ in let } y = \text{Laplace}(\epsilon / 4) x \text{ in } \hat{T} \leq y$ 
  in  $f$ 
```

Fig. 7. The Above Threshold mechanism.

Example 4.1. Suppose we want to privately compute the number of even numbers in a list and check whether it exceeds a threshold of 3. If we run AT with a low privacy budget (e.g., $\epsilon = 10^{-3}$) there is very little noise added and we are very likely to observe the true result, *i.e.*

<code>above_threshold 10⁻³ 3 (List.count (λ x. x mod 2 = 0))</code>	<code>[1, 2, 3, 4, 5]</code>	\rightarrow	<code>false</code>
<code>above_threshold 10⁻³ 3 (List.count (λ x. x mod 2 = 0))</code>	<code>[1, 2, 3, 4, 5, 6]</code>	\rightarrow	<code>true</code>

As the value of ϵ increases, we are more likely to observe `true` in the first query or `false` in the second, *i.e.*, privacy improves.

Our specification (6) captures the idea that `above_threshold` is *interactive*: after initialization, the function f can be used to compare a query to the (noisy) threshold \hat{T} until a result above \hat{T} is found, but queries can be supplied and *chosen* one by one after observing the result of previous queries.

$$\left\{ \begin{array}{l} \{ \mathcal{F}^X(\epsilon) \} \\ \text{above_threshold } \epsilon T \precsim \text{above_threshold } \epsilon T \\ \exists \text{AUTH. } \text{AUTH} * \\ \left\{ \begin{array}{l} \{ \text{AUTH} * \text{Adj}(db, db') * \text{1-sens}(q) \} \\ f f'. \quad \forall db db' q. \quad f q db \precsim f' q db' \\ \{ b b'. b = b' * \text{if not } b \text{ then } \text{AUTH} \} \end{array} \right\} \end{array} \right\} \quad (6)$$

Let us unpack the specification piece by piece. After initializing the mechanism with a privacy budget of ϵ , we obtain a pair of functions f and f' , where both functions represent the same computation but with a priori different randomly sampled values of \hat{T} , as well as an abstract “authorization token” AUTH . The Hoare quadruple for f, f' in the postcondition indicates that so long as we have the AUTH token, running the functions on a 1-sensitive query q and adjacent datasets (1) produces the same result for both $q db$ and $q db'$, *i.e.*, the computation is private, and (2) only consumes the AUTH token if $q db$ is above \hat{T} . If the result of the comparison is `false`, AUTH can be recovered in the postcondition and we can continue to privately look for a query that exceeds \hat{T} . Note that the initial budget $\mathcal{F}^X(\epsilon)$ is only spent once at initialization, but we can still privately run as many queries as it takes to get a result above the threshold.

Proving Privacy. The proof of the specification (6) proceeds as follows:

- (1) We split $\mathcal{F}^X(\epsilon)$ into $\mathcal{F}^X(\epsilon/2) * \mathcal{F}^X(\epsilon/2)$ and use the first half to pay for the **LAPLACE-SHIFT** rule with parameters $m = m' = T, c = 1, k = 1$. We thus obtain related results \hat{T}, \hat{T}' for the left- and right-hand program such that $\hat{T}' = \hat{T} + 1$. Forcing the two noisy thresholds to be at distance 1 rather than equal is a standard “trick” in the analysis of Above Threshold.

(2) We pick $AUTH \triangleq \mathcal{L}^\times(\varepsilon/2)$ and use the remaining budget to provide the initial $AUTH$ token.
 (3) We now have to show the specification for

$$\begin{aligned} f \ q \ db &\triangleq \text{let } x = q \ db \text{ in let } y = \text{Laplace } (\varepsilon / 4) \ x \text{ in } \hat{T} \leq y \\ f' \ q \ db' &\triangleq \text{let } x' = q \ db' \text{ in let } y' = \text{Laplace } (\varepsilon / 4) \ x' \text{ in } \hat{T}' \leq y' \end{aligned}$$

(4) By sensitivity of q and adjacency of the databases, x and x' are at distance at most 1.
 (5) By the definition of $AUTH$, the precondition of the refinement confers us a privacy budget of $\mathcal{L}^\times(\varepsilon/2)$. We use this budget to apply the **LAPLACE-CHOICE** rule and partition the outcomes (y, y') of the remaining Laplace sampling into two mutually exclusive cases:

- $\hat{T} \leq y$ and $\hat{T}' \leq y'$. In this case, both of the comparisons return **true**.
- $y < \hat{T}$ and $y' < \hat{T}'$. In this case, both comparisons return **false**, the rule does not consume the privacy budget and we can return $AUTH$.

Either way, f and f' return the same result b , and if $b = \text{false}$ then $AUTH$ is returned too. \square

To see that the specification (6) is indeed useful, we will now use it to verify the privacy of the interactive sparse vector technique.

4.1.2 An interactive Sparse Vector Technique. We prove privacy of the interactive SVT directly from the abstract specification of the AT mechanism (6). SVT orchestrates repeated invocations of `above_threshold` in order to identify the first N of the queries that exceed the threshold T . Unlike a purely batch-style algorithm, the SVT exposes a streaming interface that allows queries to be submitted interactively, *i.e.*, depending on the results of earlier queries. This makes the verification of privacy significantly more challenging, as the set of queries cannot be fixed in advance but may depend on previously released (noisy) information. In Clutch-DP however, the proof that `sparse_vector` is `dp` is a relatively straightforward consequence of the privacy of `above_threshold`.

The implementation in §4.1.2 works as follows. It maintains two pieces of mutable state: a reference AT to the current `Above Threshold` instance and a counter that tracks how many additional **true** results (*i.e.*, queries that exceed the threshold) may still be released. Each invocation of the returned function f runs the current `above_threshold` instance on a new query q and database db , producing a boolean result b . If b is **true** and the counter has not yet reached zero, the mechanism consumes ε privacy credits and reinitializes AT with a fresh `Above Threshold` instance. Otherwise, the counter and function reference remain unchanged. The caller may then use the result b to decide which query to issue next.

```

let sparse_vector ε T N =
  let AT = ref (above_threshold ε T) in
  let counter = ref (N - 1) in
  let f q db = let b = (! AT) q db in
    if !counter > 0 && b then
      (counter ← (! counter - 1));
      AT ← above_threshold ε T ;
    b
  in f

let SVT_stream ε T N QS db =
  let f = sparse_vector ε T N in
  let rec iter i bs =
    if i = N then List.reverse bs
    else let q = QS bs in
      let b = f q db in
      iter (if b then (i + 1) else i) (b :: bs)
  in iter 0 []

```

Fig. 8. The Sparse Vector Technique and a streaming client.

The specification (7) formalizes the intuition that the interactive SVT behaves like a private state machine that can be queried multiple times while consuming a fixed privacy budget.

$$\begin{aligned}
 & \{ \mathbb{Z}^X(N \cdot \varepsilon) \} \\
 & \text{sparse_vector } \varepsilon \text{ } T \text{ } N \precsim \text{sparse_vector } \varepsilon \text{ } T \text{ } N \\
 & \left\{ \begin{array}{l} \exists iSVT. iSVT(N) * \\ \quad \{ iSVT(n+1) * \text{Adj}(db, db') * 1\text{-sens}(q) \} \\ f f'. \quad \forall db \text{ } db' \text{ } q \text{ } n. \quad f \text{ } q \text{ } db \precsim f' \text{ } q \text{ } db' \\ \quad \{ b \text{ } b'. \text{ } b = b' * iSVT(\text{if } b \text{ then } n \text{ else } n+1) \} \end{array} \right\} \quad (7)
 \end{aligned}$$

Initially, the mechanism owns a total privacy resource of $\mathbb{Z}^X(N \cdot \varepsilon)$, corresponding to N possible above- T releases. The abstract token $iSVT(K)$ in the postcondition tracks the remaining number K of `true` results we can release. The returned function f satisfies the nested quadruple: if we own at least $iSVT(1)$, we can run a 1-sensitive query on adjacent databases and ensure we get the same result; if the result is true (i.e., the query exceeds the threshold), we decrease $iSVT$ by 1, otherwise we get back our initial token. This specification thus captures both the privacy accounting and the interactive behavior of SVT: the mechanism remains private for any adaptively chosen sequence of 1-sensitive queries until the allotted number N of positive releases has been exhausted.

Notably, the proof of this specification treats AT abstractly and relies only on its specification (6), not its implementation, which gives us a more modular analysis. The proof relies on keeping a simple invariant over the counter, the AT reference and the remaining privacy budget.

To summarize, this implementation and our verification of it has several noteworthy features:

- the SVT is a client of AT via an abstract specification,
- its interactive interface is specified through nested Hoare quadruples, and
- it requires storing a private higher-order function in a reference.

Next, we demonstrate that our specifications are expressive enough for reuse by different clients.

4.1.3 SVT Client: Streams of Queries. The standard textbook account of SVT [16] presents it as an algorithm that takes in a *stream* of queries QS and produces a list of booleans bs where b_i indicates whether the i -th query was above the noisy threshold. The stream is represented as a (possibly stateful) function that produces a new query on each invocation, and interactivity is modeled by the fact that each time a new query is requested, QS gets access to the booleans bs resulting from the preceding queries.

We can directly prove that the implementation `SVT_stream` (§4.1.2) is private by applying the generic specification for SVT. A user of `SVT_stream` only has to satisfy the textbook assumption that all of the queries are indeed 1-sensitive. The privacy reasoning is encapsulated in (7).

4.1.4 Above Threshold Client: `auto_avg`. As a last application of the Above Threshold mechanism, we analyze the `auto_avg` client that privately computes the average of a dataset. This example is taken from the online textbook [32, Chapter 10].

The fact that both `sparse_vector` and `auto_avg` can both be verified against the same abstract interface for `above_threshold` is good evidence that our specifications are indeed reusable and can be used to verify libraries without having to worry about implementation details.

To privately compute the average of a dataset it is not enough to first compute the average and then add ε Laplacian noise for a fixed ε , as this may leak information about the size of the dataset. The noise has to be calibrated to the largest element—but that value in itself is private information!

The solution adopted in the implementation of `auto_avg` in Figure 9 is to “clip” the elements of the database to lie in a bounded range $[0, B]$. If two adjacent databases are clipped to the same bound, their sum can differ by at most B . In other words, we can prove that `clip_sum B` is internally

```

let auto_avg bnds ε db =
  let bound = get_clip_bound bnds ε db in
  let sum = clip_sum bound db in
  let sum_noisy = Laplace (ε / bound) sum in
  let count_noisy = Laplace ε (List.length db) in
  sum_noisy / count_noisy

let get_clip_bound bnds ε db =
  let qs = List.map (λ b. (b, mk_query b)) bnds in
  let (bound, _) = AT_list ε 0 db qs in
  bound

let mk_query b db =
  (clip_sum b db) – (clip_sum (b + 1) db)

let clip_sum bound db =
  List.sum (List.clip bound db)

let AT_list ε T db qs =
  let AT = above_threshold ε T in
  List.find (λ (bound, q). AT q db) qs

```

Fig. 9. Privately computing the average of a list of data.

B -sensitive. We can apply internal metric composition (Lem. 3.5) to show that computing $\text{sum}_{\text{noisy}}$ by adding to the clipped sum Laplacian noise with scale (ε/B) is ε -private. Therefore, `auto_avg` achieves $(3 \cdot \varepsilon)$ -DP, where the budget is divided equally between the call to `get_clip_bound` and the two calls to `Laplace`:

$$\{\mathcal{F}^{\times}(3 \cdot \varepsilon) * \text{Adj}(db, db')\} \text{auto_avg bnds } \varepsilon \text{ db} \preceq \text{auto_avg bnds } \varepsilon \text{ db}' \{x \ x'. x = x'\}$$

The utility of `auto_avg` stems from carefully choosing B . Given a list of candidate bounds `bnds` we can do this privately via the `AT` mechanism. The function `get_clip_bound` to finds the first value b in `bnds` such that the sum of elements in `db` stops increasing if the clipping bound is relaxed from b to $b + 1$. Testing this for all values in `bnds` via `mk_query` is 1-sensitive. Therefore we can directly apply the specification (6) for `above_threshold` to derive that `get_clip_bound` is ε -dp.

4.2 Privacy Filters

A common implementation technique for DP in general-purpose programming languages is to explicitly track the remaining privacy budget as a program variable. At the beginning of a data analysis, this variable is initialized to the global privacy budget ε , and it must remain non-negative throughout the program execution. So long as care is taken to decrement the budget every time (noisy) data is released, the whole data analysis is ε -private. To ensure that these rules are respected, the management of the privacy budget is commonly encapsulated in a *privacy filter*, a higher-order function that runs a computation only if there is sufficient budget for it. This programming pattern provides a separation of concerns: if the privacy analysis of the individual computations is correct, and the filter is correctly implemented, then the entire computation is private. This allows us to verify the different components (privacy filter and mechanisms) modularly.

```

let privacy_filter ε_budget =
  let ε_rem = ref ε_budget in
  let try_run ε_cost f =
    if !ε_rem < ε_cost then
      None
    else
      ε_rem ← !ε_rem – ε_cost ;
      Some (f ())
  in try_run

let adaptive_count ε_coarse ε_precise T ε_budget predicates db =
  let try_run = privacy_filter ε_budget in
  List.map (λ pred. let n_exact = List.count pred db in
    let g _ = Laplace ε_precise n_exact in
    let f _ = let n_coarse = Laplace ε_coarse n_exact in
      let n_precise = if T < n_coarse
        then try_run ε_precise g
        else None in
      (n_coarse, n_precise)
    in try_run ε_coarse f
  predicates

```

Fig. 10. Implementations of a Privacy Filter and Adaptive Count.

The implementation `privacy_filter` in [Figure 10](#) works as follows. Upon initialization it allocates a reference that tracks the remaining privacy budget and returns a closure `try_run` that can be used to run private computations in an interactive manner. If a client tries to run a computation with cost exceeding the remaining budget, `try_run` does not run the computation, otherwise it decreases the budget by the cost, runs the computation, and returns the result.

For the sake of simplicity we present a privacy filter that only tracks the ε -budget, but the method directly generalizes to (ε, δ) -privacy filters.

4.2.1 Proving Privacy of Privacy Filters. The high-level intuition for the privacy filter is that it should never exceed the budget that was initially set up so long as any client that calls `try_run` $\varepsilon_{\text{cost}} f$ ensures that f is indeed $\varepsilon_{\text{cost}}\text{-dp}$. This intuition is captured by the following specification.

$$\begin{aligned} & \{ \mathcal{F}^X(\varepsilon_{\text{budget}}) \} \\ \forall \varepsilon_{\text{budget}}. \quad & \text{privacy_filter } \varepsilon_{\text{budget}} \preceq \text{privacy_filter } \varepsilon_{\text{budget}} \\ & \{ \text{try_run } \text{try_run'}. \exists iPF. iPF * \text{try_run-spec} \} \end{aligned} \quad (8)$$

where `try_run-spec` is defined as

$$\begin{aligned} & \left\{ \begin{array}{l} iPF * \text{Inv}_f \\ * \{ \mathcal{F}^X(\varepsilon_{\text{cost}}) * iPF * \text{Inv}_f \} f () \preceq f' () \{ v v'. v = v' * iPF * \text{Inv}_f \} \end{array} \right\} \\ \forall \varepsilon_{\text{cost}} f f' \text{Inv}_f. \quad & \text{try_run } \varepsilon_{\text{cost}} f \preceq \text{try_run' } \varepsilon_{\text{cost}} f' \\ & \{ b b'. b = b' * iPF * \text{Inv}_f \} \end{aligned} \quad (9)$$

The existentially quantified iPF token in the postcondition of (8) represents a client's ability to execute computations privately via `try_run`. The specification (9) defines the behavior of `try_run` $\varepsilon_{\text{cost}} f$. Assuming that for a privacy cost of $\varepsilon_{\text{cost}}$ the functions f and f' produce equal results and maintain an invariant Inv_f , calling them through `try_run` (and `try_run'` respectively) satisfies the same invariant regardless of whether there actually is enough budget left to execute f . Since f may itself contain calls to `try_run`, the specification for (f, f') has access to the iPF token. We give a simple application of this expressivity in the form of nested calls to the privacy filter in [§4.2.2](#).

The proof of the specification in Clutch-DP is straightforward. We define the token iPF as

$$\exists \varepsilon. \mathcal{F}^X(\varepsilon) * \varepsilon_{\text{rem}} \mapsto \varepsilon * \varepsilon'_{\text{rem}} \mapsto_s \varepsilon.$$

The link between the logical resource representing ownership of the error budget and the program state tracking the remaining budget allows us to conclude that the call to `Some(f())` in the definition of `try_run` is only executed when sufficient privacy budget remains to satisfy the precondition of (f, f') in (9), and hence the invariant Inv_f is satisfied. In case `try_run` decides that the budget is insufficient for f , the invariant is trivially preserved.

The power of this specification for `privacy_filter` lies in the fact that a client does not have to perform *any* privacy accounting or reasoning whatsoever for `try_run` $\varepsilon_{\text{cost}} f$! We can conveniently combine library functions to build a private f and run it without having to worry whether we still have enough budget: the filter ensures that the initial budget is never exceeded.

4.2.2 Client: Adaptive Counting. Another advantage of implementing DP through privacy filters is that it allows the data analyst to decide *dynamically* where the privacy budget should be spent, *i.e.*, the way the budget is spent can adapt to the results of prior analyses. This is especially useful in exploratory data analysis when it is unclear, *a priori*, what values the dataset ranges over.

The example `adaptive_count` in [Figure 10](#) employs a form of adaptivity to privately count the number of elements of db that satisfy each of the tests in the list of predicates. First, a cheaper but less precise count is performed, consuming $\varepsilon_{\text{coarse}}$ privacy credits. Only if this yields a promising result that exceeds a threshold T , a more precise analysis is performed for an additional larger

budget of $\epsilon_{\text{precise}}$. The result of `adaptive_count` is thus a numeric estimate for each predicate, with a more precise value for a few “important” candidates.

A conservative privacy analysis would have to assume a worst-case cost of $\text{length}(\text{predicates}) \cdot (\epsilon_{\text{coarse}} + \epsilon_{\text{precise}})$ even if many of the coarse counts may in practice not exceed the threshold. By employing a privacy filter, we can instead fix a budget ϵ_{budget} that we want to allocate to this analysis task and try to run the analysis so long as the filter has enough budget left. If only a few of the data entries exceed T , this allows to count many more predicates than the conservative analysis without a privacy filter.

In Clutch-DP, we can prove that `adaptive_count` is ϵ_{budget} -dp from the specification (8) because `f` and `g` meet the precondition of `try_run`, as they consume a budget of ϵ_{coarse} and $\epsilon_{\text{precise}}$ respectively.

4.3 Caching Techniques for DP

Interactive analysis with data-dependent queries is common in real-world workloads for DP. This poses a challenge for DP frameworks because it makes it impossible to statically avoid repeated evaluation of certain queries, say, by refactoring code that requires the same result, and hence repeated queries inflate the privacy cost unnecessarily. This problem can be solved with a query cache that memoizes the results of a query on first execution and reuses this result upon repetition. One would hope that reusing a noisy result in a repeated query should be “for free” and consume no privacy budget. However, the privacy analysis of such a memoization method is subtle, because the privacy cost of a query depends on the history of queries, which is highly non-local information.

```
let mk_query_cache add_noise db =
  let cache = Map.init () in
  let run_cache q =
    match Map.get cache q with
    | Some x => x
    | None  => let x = add_noise q db in
      Map.set cache q x;
      x
  in run_cache

let map_cache add_noise qs db =
  let run_cache =
    mk_query_cache add_noise db in
  List.map run_cache qs
```

Fig. 11. Implementations of a cache and a client.

In this section, we reason about the privacy of the caching method introduced in Figure 3, which we implement in RandML through the algorithm shown in Figure 11. Our formalization crucially relies on our logic being able to support higher-order functions, local state, and a resource-based representation of the privacy budget.

4.3.1 Cache spec: repeated queries are free. Upon initialization, `mk_query_cache` allocates a mutable map cache and returns a closure `run_cache` that stores and looks up noisy query results in cache. This is reflected in (10) as the existentially quantified $\text{iC}(M_{\text{cache}})$ resource. Initially, the map M_{cache} is empty, but it can be updated via `run_cache` as we will see next.

$$\begin{aligned} & \{\text{Adj}(db, db')\} \\ & \quad \text{mk_query_cache add_noise db} \preceq \text{mk_query_cache add_noise db}' \\ & \quad \{\text{run_cache run_cache}'. \exists \text{iC}. \text{iC}(\text{Map.empty}) * \text{spec-cached} * \text{spec-fresh}\} \end{aligned} \tag{10}$$

where `spec-fresh` is defined as

$$\begin{aligned} & \{q \notin \text{dom}(M_{\text{cache}}) * \mathcal{L}^X(\epsilon) * \mathcal{L}^+(\delta) * (\epsilon, \delta)\text{-iDP}(\text{add_noise } q) * \text{iC}(M_{\text{cache}})\} \\ & \forall M_{\text{cache}}. q. \quad \text{run_cache } q \preceq \text{run_cache}' q \\ & \quad \{v v'. v = v' * \text{iC}(M_{\text{cache}}[q \mapsto v])\} \end{aligned} \tag{11}$$

and *spec-cached* is defined as

$$\begin{aligned} & \{q \in \text{dom}(M_{\text{cache}}) * \text{iC}(M_{\text{cache}})\} \\ \forall M_{\text{cache}} \ q. \quad & \text{run_cache } q \precsim \text{run_cache}' q \\ & \{v \ v'. \ v = v' * \text{iC}(M_{\text{cache}}) * M_{\text{cache}}[q] = v\} \end{aligned} \quad (12)$$

The specification (11) describes the behavior of $(\text{run_cache } q)$ on a query that has not been memoized yet. It requires that add_noise should indeed run q under (ε, δ) -iDP and assumes ownership of enough privacy credits to pay for this execution. Furthermore, it requires ownership of $\text{iC}(M_{\text{cache}})$ for the current internal state of the cache. In the postcondition, we recover iC where M_{cache} is updated with the result of the noisy query.

The intuition that repeated queries should be free is formalized in (12): if q is in the cache then no privacy credits are consumed for executing it under run_cache .

4.3.2 A cache client. A simple application is map_cache (Figure 11), which employs the cache to run add_noise on a list of queries qs . We can prove that the privacy cost of map_cache is $(k\varepsilon, k\delta)$ where (ε, δ) is the privacy cost of the add_noise mechanism and $k = |\{q \in qs\}|$ is the number of unique queries in qs . The proof of (13) follows directly from the abstract specification (10).

$$(\forall q \in qs. (\varepsilon, \delta)\text{-iDP}(\text{add_noise } q)) \multimap (k\varepsilon, k\delta)\text{-iDP}(\text{map_cache add_noise } qs) \quad (13)$$

Without caching, the privacy cost would have to be multiplied by $\text{List.length}(qs)$ instead of k .

5 Soundness: A Model of Clutch-DP

In this section we give an overview of the model behind Clutch-DP and its adequacy theorem. A detailed account can be found in §A.

Our program logic is based around the notion of (ε, δ) -approximate coupling [7, 37]:

DEFINITION 5.1. Let A, B be countable types, and $\Phi \subseteq A \times B$ a relation. Given two real-valued parameters $0 \leq \varepsilon, \delta$, we say that there is an (ε, δ) -approximate Φ -coupling between distributions $\mu_1: \mathcal{D}(A), \mu_2: \mathcal{D}(B)$ if, for any real-valued random variables $f: A \rightarrow [0, 1], g: B \rightarrow [0, 1]$ such that $\forall (a, b) \in \Phi, f(a) \leq g(b)$, the following holds: $\mathbb{E}_{\mu_1}[f] \leq \exp(\varepsilon) \cdot \mathbb{E}_{\mu_2}[g] + \delta$. We denote the existence of such a coupling by $\mu_1 \Phi^{(\varepsilon, \delta)} \mu_2$.

The model is similar in spirit to that of Approxix [21], which can be seen as based on a notion of $(0, \delta)$ -approximate coupling. There is a tight connection between couplings and DP: a function $f: DB \rightarrow \mathcal{D}(X)$ is (ε, δ) -dp iff for any adjacent inputs $b, b': DB$, we have $f b \stackrel{(\varepsilon, \delta)}{=} f b'$.

With this definition in mind, we can now state the adequacy theorem of Clutch-DP.

THEOREM 5.2. Let f, f' be two RandML functions and $\Phi, \Psi: \text{Val} \times \text{Val} \rightarrow \text{Prop}$. If

$$\{\Phi(w, w') * \mathcal{L}^+(\delta) * \mathcal{L}^-(\varepsilon)\} f w \precsim f' w' \{\Psi\}$$

holds in Clutch-DP then, for any initial states σ, σ' , we have $\text{exec}(f w, \sigma) \Psi^{(\varepsilon, \delta)} \text{exec}(f' w', \sigma')$.

By instantiating Φ with adjacency and Ψ with equality, we recover Theorem 3.2 as a corollary. Internally, the Hoare quadruples are defined in terms of a primitive, unary notion of weakest precondition (WP), where the right-hand side program is represented as a separation logic resource [17, 20, 21]. Validity of the WP is defined by guarded induction on the program execution, establishing an approximate coupling at each step, and finally composing all the couplings into a coupling for the full execution. Each program logic rule is then proven sound w.r.t. the definition of the WP. In particular, all standard separation logic rules for the deterministic fragment of RandML can be re-established. We refer the reader to the supplementary material for more details.

	Methodology	State	Higher-order functions		DP definition	Verified	Beyond comp.	Adaptive	Interactive
			global, ● local	△ first-order, ■ HO					
Fuzz [35]	lin. types	○	●	ε	○	○	○	○	○
Fuzz ^{ε,δ} [13]	lin. typ., path adj.	○	●	(ε, δ)	○	○	○	○	○
DFuzz [18]	lin. dep. types	○	●	ε	○	○	○	○	○
AFuzz [41]	dyn. typ. + priv. filter	○	●	(ε, δ)	○	○	●	○	○
Fuzzi [43]	Fuzzi-like + apRHL	○	○	(ε, δ)	○	●	○	○	○
Duet [33]	sens. + priv. types	○	●	(ε, δ), RDP, (z/t)CDP	○	○	○	○	○
Jazz [40]	ctxt'l sens. + priv. typ.	○	●	(ε, δ), RDP, zCDP	○	○	○	○	○
DPella [27]	eDSL dep. types	○	●	ε	○	○	○	○	○
Solo [1]	eDSL dep. types	○	●	(ε, δ), RDP	○	○	○	○	○
Spar [28]	eDSL dep. types	○	●	ε	○	○	○	○	○
HOARe ² [8]	rel. refinement types	○	●	(ε, δ)	○	○	○	○	○
LightDP [42]	types + annot.	○	△	ε	○	●	○	○	○
SampCert [14]	semantic	○	●	(ε, δ), zCDP	●	●	○	○	○
DP/Isabelle [39]	semantic	○	○	(ε, δ)	●	●	○	○	○
apRHL+ [6, 7, 10]	rel. prob. Hoare logic	○	△	○	(ε, δ)	○	●	○	●
HO-RPL [2]	rel. prob. Hoare logic	○	△	●	(ε, δ)	○	○	○	●
Clutch-DP	rel. prob. sep. logic	●	■	●	(ε, δ)	●	●	●	●

Fig. 12. Comparison of DP systems. “Beyond comp.” = support for mechanisms whose DP goes beyond composition laws (e.g., SVT, RNM).

6 Related Work

Types for DP. A wide range of type systems ensuring DP have been developed. Fuzz [35] and its variants track function sensitivity via linear types and rely on metric composition to statically ensure pure DP for a probabilistic λ -calculus. DFuzz [18] integrated linear and dependent types to improve sensitivity analysis. Fuzz^{ε,δ} [13] extends Fuzz with support for (ϵ, δ) -DP. The two-level type system of Adaptive Fuzz [41] enhances static typing by integrating a trusted privacy filter into the language runtime. The system allows programming with adaptive composition by dynamically type-checking programs during execution and composing them according to the privacy filter. In Clutch-DP, privacy filters are just regular programs that can be verified in their own right.

Duet [33] is a linear type system supporting various notions of DP by a separation of the language into a sensitivity and a privacy layer which interact through bespoke composition laws that restrict rescaling, which limits the kinds of higher-order functions that can be type-checked. Jazz [40] lifts some of these restrictions by introducing latent contextual effect types. HOARe² [8] encodes sensitivity and privacy information in relational refinement types for a pure calculus.

The DPella system [27], Solo [1], and Spar [28] leverage extensions to Haskell’s type system to encode sensitivity (or distance) information via dependent types instead of linear types.

None of these systems support mutable state, or the verification of programs whose privacy requires advanced budget management instead of following directly from composition laws.

LightDP [42] employs a dependent relational type system to bound distances in program variables. With SMT-backed type inference, LightDP can verify DP for some mechanisms beyond composition laws (e.g., non-interactive SVT), but the method does not extend to advanced language features or flexible privacy budget analysis as used, e.g., in Report Noisy Max.

Logics for DP. The apRHL(+) program logics [6, 7, 10] can prove (ϵ, δ) -DP for programs written in a first-order while-language. Reasoning about basic mechanisms is well supported in apRHL, and it has been applied to advanced mechanisms such as SVT. Verification of interactive DP is supported

through a specialized rule for adversaries that can interact with mechanisms through fixed patterns. In apRHL, adversaries are programs subject to a number of side-conditions; in contrast, our model of interactive computations via higher-order functions does not come with syntactic restrictions and can be composed modularly. For instance, we can recover the apRHL model of SVT used by Barthe et al. [7, Fig. 1] through our streaming SVT (§4.1.3) by instantiating the stream of queries QS with the adversary \mathcal{A} of *loc. cit.* The privacy budget in apRHL is handled via a grading on judgments which offers less flexibility compared to Clutch-DP’s privacy credits.

A variant of the EasyCrypt prover supports apRHL but the implementation of the apRHL rules are part of the trusted code base, whereas the rules of Clutch-DP are proven sound in a proof assistant. During the development of our case studies we found a bug in the Laplace sampling rule of EasyCrypt⁶ which has accidentally been exploited in a user-contributed privacy proof of Report Noisy Max. Our foundational approach would have prevented us from introducing such an erroneous rule. Proof search for coupling-based proofs of DP for first-order programs was studied in [4]. Their definition of choice couplings inspired our **LAPLACE-CHOICE** rule.

Fuzzi [43] integrates an apRHL-style logic with a Fuzz-inspired sensitivity- and privacy-logic; working at the intersection of the languages of the two systems, it does not support mutable state or higher-order functions.

The HO-RPL logic extends the ideas from apRHL to support higher-order functions and continuous distributions by giving a denotational semantics of programs in Quasi-Borel Spaces, but it is not known how to extend this approach to other language features that our challenge problems require such as dynamic allocation or higher-order store.

The Isabelle/HOL formalization of DP [39] develops the mathematical theory of DP in the continuous setting. Working directly in the semantics, they prove privacy of Report Noisy Max, stating however that “the formal proof is quite long”. Our proof of RNM is about six times shorter, demonstrating the benefits of working in a program logic.

The SampCert project [14] formalized DP in the Lean prover by interpreting a shallow embedding of a while-like language in the same kind of denotational semantics that apRHL is based on. Rather than assuming that the language has a primitive that samples from the Laplacian, SampCert proves that an efficient implementation of a sampler realizes the Laplace distribution. Their approach is particularly well-suited to carrying out the low-level probabilistic reasoning and focuses less on building modular systems. SampCert formalizes non-interactive variants of AT and SVT.

Separation logic. Several probabilistic separation logics exist, but only Clutch-DP supports reasoning about DP. The Approxix [21] relational separation logic supports reasoning about approximate program equivalence. Their notion of α -approximate equivalence amounts to $(0, \alpha)$ -dp in our setting, *i.e.*, our additive privacy credits correspond to their “error credits”. Approxix can prove cryptographic security or correctness of samplers but cannot express DP. Bluebell [5] encodes coupling-based relational reasoning via a conditioning modality, but only supports exact program equivalences for terminating first-order programs.

7 Conclusion and Future Work

We have developed Clutch-DP, a probabilistic higher-order separation logic for DP. To demonstrate how Clutch-DP enables modular verification of DP libraries, we addressed three representative challenges and verified a wide range of case studies involving interactive mechanisms, privacy filters, and memoization. Clutch-DP is proven sound as a library for Iris in the Rocq Prover.

In future work, we would like to extend Clutch-DP to model concurrency to reason about local DP in a distributed setting. We would also like to integrate verified sampling mechanisms as developed

⁶We have disclosed the bug to the EasyCrypt team.

by, *e.g.*, de Medeiros et al. [14], to provide end-to-end DP guarantees by adapting the techniques of Aguirre et al. [3] to verify rejection samplers to the relational setting. Finally, it would be interesting to integrate other divergences [38] with relational separation logics to model, *e.g.*, Rényi-DP.

Data Availability Statement

The Rocq formalization accompanying this work is available on GitHub at github.com/logsem/clutch.

Acknowledgments

This work was supported in part by the National Science Foundation, grant no. 2338317, the Carlsberg Foundation, grant no. CF23-0791, a Villum Investigator grant, no. 25804, Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation, and the European Union (ERC, CHORDS, 101096090). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

A Semantic Model

In this section we dive into the model behind Clutch-DP and show how our adequacy theorem is proven. We begin by recalling some notions of probability theory.

A.1 Probability Theory and Probabilistic Couplings

Given a countable set A , a probability (sub-)distribution μ over A is a function $\mu: A \rightarrow [0, 1]$ such that $\sum_a \mu(a) \leq 1$. We use the notation $\mathcal{D}(A)$ to denote the set of distributions over A . Distributions have a well-known monadic structure, see e.g. [21]. We use $\text{ret } a$ and $\mu \gg f$ to denote the return and bind operations. Given a random variable $f: A \rightarrow [0, 1]$, its expected value is given by $\mathbb{E}_\mu[f] \triangleq \sum_a \mu(a)f(a)$.

We recall the notion of (ε, δ) -approximate coupling, due to [7, 37]:

DEFINITION A.1. *Let A, B be countable types, and $\Phi \subseteq A \times B$ a relation. Given two real-valued parameters $0 \leq \varepsilon, \delta$, we say that there is an (ε, δ) -approximate Φ -coupling between distributions $\mu_1: \mathcal{D}(A), \mu_2: \mathcal{D}(B)$ if, for any real-valued random variables $f: A \rightarrow [0, 1], g: B \rightarrow [0, 1]$ such that $\forall (a, b) \in \Phi, fa \leq gb$, the following holds: $\mathbb{E}_{\mu_1}[f] \leq \exp(\varepsilon) \cdot \mathbb{E}_{\mu_2}[g] + \delta$. We denote the existence of such a coupling by $\mu_1 \Phi^{(\varepsilon, \delta)} \mu_2$.*

The reason we are interested in approximate couplings is that there is a tight connection between them and DP. Indeed, when instantiating Φ to be the equality relation we have the following result:

THEOREM A.2. *Let $f: \text{DB} \rightarrow \mathcal{D}(X)$ be a function over databases. Then, f is (ε, δ) -dp iff for any inputs $b, b': \text{DB}$ such that $d_{\text{DB}}(b, b) \leq 1$, we have $f \ b \ (=)^{(\varepsilon, \delta)} \ f \ b'$.*

Approximate couplings are the key component of our relational program logic. One can think of μ_1, μ_2 as two randomized computations, Φ as a relational postcondition we wish to establish and (ε, δ) as a privacy budget that we can spend in order to prove the postcondition. We will see later in the section how to interpret ε and δ as separation logic resources, but for now notice that they have an affine flavor to them, since they satisfy the following monotonicity lemma:

LEMMA A.3. *Let $\mu_1: \mathcal{D}(A), \mu_2: \mathcal{D}(B), \Phi \subseteq A \times B$. Assume that $\mu_1 \Phi^{(\varepsilon, \delta)} \mu_2$. Then, for all $\varepsilon', \delta', \Phi'$ such that $\varepsilon \leq \varepsilon', \delta \leq \delta', \Phi \subseteq \Phi'$, we also have $\mu_1 \Phi'^{(\varepsilon', \delta')} \mu_2$.*

In order to build a logic around couplings, we need them to be able to compose them along the operations of the underlying distribution monad. The following results are well-known:

THEOREM A.4. *The distribution monad operations lift to couplings, in the sense that:*

- Let $a: A, b: B$ and $\Phi \subseteq A \times B$ such that $(a, b) \in \Phi$. Then $\text{ret } a \Phi^{(\varepsilon, \delta)} \text{ret } b$.
- Let $\mu_1: \mathcal{D}(A), \mu_2: \mathcal{D}(B), f: A \rightarrow \mathcal{D}(A'), g: B \rightarrow \mathcal{D}(B')$ and $\Phi \subseteq A \times B, \Psi \subseteq A' \times B'$ such that $\mu_1 \Phi^{(\varepsilon, \delta)} \mu_2$ and, for all $(a, b) \in \Phi, fa \Psi^{(\varepsilon', \delta')} gb$. Then, $(\mu_1 \gg f) \Psi^{(\varepsilon+\varepsilon', \delta+\delta')} (\mu_2 \gg g)$.

In order to support a more flexible form of composition, such as in our Above Threshold example (§4.1.1), we introduce a more general version of the composition lemma above, where the choice of coupling in the continuation can depend on the result of the first step. This is heavily inspired by choice couplings [7], but it really shines in our setting, since we have a resourceful treatment of the privacy budget.

THEOREM A.5. *Let $\mu_1: \mathcal{D}(A), \mu_2: \mathcal{D}(B), f: A \rightarrow \mathcal{D}(A'), g: B \rightarrow \mathcal{D}(B')$. Assume we have a predicate $\Xi \subseteq A$, and $\Phi_1, \Phi_2 \subseteq A \times B, \Psi \subseteq A' \times B'$, with Φ_1, Φ_2 disjoint in the sense that, $\forall a, a', b. a \in \Xi \wedge a' \notin \Xi \rightarrow ((a, b) \notin \Phi_1 \vee (a', b) \notin \Phi_2)$. Assume further that:*

- $\mu_1 \Phi_1^{(\varepsilon_1, \delta_1)} \mu_2$
- $\mu_1 \Phi_2^{(\varepsilon_2, \delta_2)} \mu_2$

- For all a, b such that $a \in \Xi$ and $(a, b) \in \Phi_1$, $fa \Psi^{(\varepsilon'_1, \delta'_1)} gb$.
- For all a, b such that $a \notin \Xi$ and $(a, b) \in \Phi_2$, $fa \Psi^{(\varepsilon'_2, \delta'_2)} gb$.

Then, $(\mu_1 \gg f) \Psi^{(\varepsilon, \delta)} (\mu_2 \gg g)$, where $\varepsilon \triangleq \max(\varepsilon_1 + \varepsilon'_1, \varepsilon_2 + \varepsilon'_2)$ and $\delta \triangleq \delta_1 + \delta_2 + \max(\delta'_1, \delta'_2)$

The idea behind the statement above is that it allows us to use two different couplings for the first step: one over Φ_1 and another over Φ_2 . Then get a sample a from μ_1 , and we choose which of the couplings to use depending on whether a lands in Ξ , and we continue the rest of the execution. In particular, this allows us to optimize the use of the ε component of the privacy budget, since we can set the amount to use in each of the variants and only spend what we need. In the case of the δ component, we can also set the optimal amount for the continuation, but not for the first step, both δ_1 and δ_2 need to be spent no matter the result.

A.2 Hoare Quadruples and the Weakest Precondition

Hoare Quadruples are not a primitive notion in Clutch-DP. Instead, they are defined in terms of a weakest precondition predicate (WP). This is a unary predicate about the left-hand side program. The right-hand program e_R is tracked by a resource $\text{spec}(e_R)$.

$$\{P\} e \preceq e' \{v v'. Q\} \triangleq \square(P \multimap \exists K. \text{spec}(K[e']) \multimap \text{wp } e \{v. \exists v'. \text{spec}(K[v']) * Q(v, v')\})$$

to state that e refines e' and that if e terminates with v then e' terminates with v' and the postcondition P is satisfied.

The WP, whose definition is shown below, couples the execution of the implementation program together with the execution of the specification program and ensures that the postcondition holds at the end. The construction is similar to the one presented in Approxix [21], but adapted to our more general notion of coupling. Crucially, we also employ an updated program coupling modality to support choice coupling composition at the level of program steps. Notice that this definition tracks the two components of the error, both ε and δ . As usual, the weakest precondition is a predicate defined as a guarded fixpoint, whose existence is ensured by the presence of the later modality (\triangleright) in front of the recursive occurrence of the weakest precondition:

$$\begin{aligned} \text{wp}_{\mathcal{E}} e_1 \{\Phi\} &\triangleq \forall \sigma_1, \rho'_1, \varepsilon_1, \delta_1. S(\sigma_1, \rho'_1, \varepsilon_1, \delta_1) \multimap \\ &\quad \mathcal{E} \Rightarrow_{\emptyset} \text{scpl } \sigma_1 \precsim_{(\varepsilon_1, \delta_1)} \rho'_1 \{\sigma_2, \rho'_2, \varepsilon_2, \delta_2. \\ &\quad (e_1 \in \text{Val} * \mathcal{E} \Rightarrow \mathcal{E} S(\sigma_2, \rho'_2, \varepsilon_2, \delta_2) * \Phi(e_1)) \vee \\ &\quad (e_1 \notin \text{Val} * \text{pcpl } (e_1, \sigma_2) \precsim_{(\varepsilon_2, \delta_2)} \rho'_2 \{\sigma_3, \rho'_3, \varepsilon_3, \delta_3. \\ &\quad \triangleright \text{scpl } \sigma_3 \precsim_{(\varepsilon_3, \delta_3)} \rho'_3 \{\sigma_4, \rho'_4, \varepsilon_4, \delta_4. \mathcal{E} \Rightarrow \mathcal{E} S(\sigma_4, \rho'_4, \varepsilon_4, \delta_4) * \text{wp}_{\mathcal{E}} e_2 \{\Phi\}\})\}) \end{aligned}$$

The weakest precondition requires us to own the *state interpretation* $S(\sigma_1, \rho'_1, \varepsilon_1, \delta_1)$, which is a predicate that maps the different resources in Clutch-DP (the physical state of the implementation program, the configuration of the specification program, and the privacy budget) to their corresponding resource algebras. We then have to prove two different cases, depending on whether the current program e is or is not a value. In the former case, we can just give back the state interpretation and prove Φ . Otherwise, we have to prove a property about the distribution that results after taking one program step. This is phrased in terms of the composition of two modalities, the *specification coupling modality* and the *program coupling modality*.

The specification coupling modality, initially introduced in Approxix, accounts for logical updates to the current state of the execution. In particular, it allows for the specification program to take

physical steps. It is defined as and inductive predicate given by the following rules:

$$\begin{array}{c}
 \text{SPEC-COUP-RET} \\
 \frac{\Phi(\sigma, \rho', \varepsilon, \delta)}{\text{scpl } \sigma \precsim_{(\varepsilon, \delta)} \rho' \{ \Phi \}}
 \end{array}$$

$$\begin{array}{c}
 \text{SPEC-COUP-BIND} \\
 \frac{\begin{array}{c} \mu_1 R^{(\varepsilon_1, \delta_1)} (\mu'_1 \gg \lambda \sigma'_2. \text{step}_n(e'_1, \sigma'_2)) \quad \varepsilon_1 + \varepsilon_2 \leq \varepsilon \quad \delta_1 + \delta_2 \leq \delta \\ \text{erasable}(\mu_1, \sigma_1) \quad \text{erasable}(\mu'_1, \sigma'_1) \quad \forall (\sigma_2, \rho'_2) \in R. \models_{\emptyset} \text{scpl } \sigma_2 \precsim_{(\varepsilon_2, \delta_2)} \rho'_2 \{ \Phi \} \end{array}}{\text{scpl } \sigma_1 \precsim_{(\varepsilon, \delta)} (e'_1, \sigma'_1) \{ \Phi \}}
 \end{array}$$

For the purpose of this work, the notion of erasability can be ignored, this is used to support presampling tapes [20]. By the **SPEC-COUP-RET** rule, if the current implementation and specification programs, and error budget satisfy Φ , we can immediately conclude $\text{scpl } \sigma \precsim_{(\varepsilon, \delta)} \rho' \{ \Phi \}$. The **SPEC-COUP-BIND** considers the recursive case. We can split our current credit budget ε, δ into $\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2$ and use ε_1, δ_1 to take a number n of steps on the specification program under an approximate coupling that satisfies an intermediate relation R . Then from $(\sigma_2, \rho'_2) \in R$, we continue to recursively establish $\text{scpl } \sigma_2 \precsim_{(\varepsilon_2, \delta_2)} \rho'_2 \{ \Phi \}$, with our remaining credit budget ε_2, δ_2 .

The program coupling modality couples exactly one physical step of the left-hand side program with an arbitrary number of physical steps n of right-hand side program. In addition, the right-hand side program is also allowed to execute an erasable update before its physical steps. Choice couplings are baked into the definition of the program coupling modality, thus the intermediate pair of states that the program reach can satisfy one of either R_1 or R_2 before executing the continuation. The amount of error credits that is available for the continuation may also depend on this choice.

$$\begin{array}{c}
 \text{PROG-COUP-BIND} \\
 \frac{\begin{array}{c} \text{step}(e_1, \sigma_1) R_1^{(\varepsilon_1, \delta_1)} (\mu'_1 \gg \lambda \sigma'_2. \text{step}_n(e'_1, \sigma'_2)) \quad \text{step}(e_1, \sigma_1) R_2^{(\varepsilon_2, \delta_2)} (\mu'_1 \gg \lambda \sigma'_2. \text{step}_n(e'_1, \sigma'_2)) \\ \text{red}(e_1, \sigma_1) \quad \varepsilon_1 + \varepsilon_2 \leq \varepsilon \quad \varepsilon'_1 + \varepsilon'_2 \leq \varepsilon \quad \delta_1 + \delta_2 + \max(\delta'_1 + \delta'_2) \leq \delta \quad \text{erasable}(\mu'_1, \sigma'_1) \\ \forall ((e_2, \sigma_2), (e'_2, \sigma'_2)). ((P(e_2, \sigma_2) \wedge R_1((e_2, \sigma_2), (e'_2, \sigma'_2))) \models_{\emptyset} \Phi(e_2, \sigma_2, e'_2, \sigma'_2, \varepsilon'_1, \delta'_1)) * \\ ((\neg P(e_2, \sigma_2) \wedge R_2((e_2, \sigma_2), (e'_2, \sigma'_2))) \models_{\emptyset} \Phi(e_2, \sigma_2, e'_2, \sigma'_2, \varepsilon'_2, \delta'_2)) \end{array}}{\text{pcpl } (e_1, \sigma_1) \precsim_{(\varepsilon, \delta)} (e'_1, \sigma'_1) \{ \Phi \}}
 \end{array}$$

A.3 An Adequacy Theorem

With these definitions in hand, we can state the adequacy theorem of Clutch-DP, now in full generality:

THEOREM A.6. *Let e, e' be two RandML functions and $\Psi: \text{Val} \times \text{Val} \rightarrow \text{Prop}$ a proposition over (final) values. If $\text{spec}(e') * \mathbb{F}^X(\varepsilon) * \mathbb{F}^+(\delta) \rightarrow* \text{wp } e \{ v, v'.\text{spec}(v') * \Psi \}$ holds in Clutch-DP, then, for any initial states σ, σ' , we have $\text{exec}(e, \sigma) \Psi^{(\varepsilon, \delta)} \text{exec}(e', \sigma')$.*

The proof is done by induction on the execution of the left-hand side program, by proving the lemma below, and finally taking a limit:

LEMMA A.7. *Let e, e' be two RandML functions and $\Psi: \text{Val} \times \text{Val} \rightarrow \text{Prop}$ a proposition over (final) values. If $\text{spec}(e') * \mathbb{F}^X(\varepsilon) * \mathbb{F}^+(\delta) \rightarrow* \text{wp } e \{ v, v'.\text{spec}(e') * \Psi \}$ holds in Clutch-DP, then, for any $n \in \mathbb{N}$ and any initial states σ, σ' , we have $\text{exec}_n(e, \sigma) \Psi^{(\varepsilon, \delta)} \text{exec}(e', \sigma')$.*

The key part of the proof happens when the program on the left is not a value, and has to take a step. This requires two applications of STATE-COUP-BIND and one application of PROG-COUP-BIND

to establish the specification coupling in the non-value case of the WP. In particular, **PROG-COUPLE-BIND** allows us to advance the program on the left exactly one step, which lets us use our inductive hypothesis about $\text{exec}_n(e, -)$ to reason about $\text{exec}_{n+1}(e, \sigma)$.

B Case Study: Report Noisy Max

The Report Noisy Max mechanism (RNM) is a DP mechanism that receives a list of N 1-sensitive queries, runs them privately by adding noise sampled from $\text{Laplace}(\varepsilon/2)$, and returns the index of the query with the highest value. While a naïve analysis of the algorithm using composition theorems would show that the algorithm is $(N \cdot \varepsilon/2)$ -dp, the only information that is released is the index of the highest query. All other (noisy) results are discarded. This suggests that a better analysis should be possible. In fact, one can prove that RNM is ε -dp, in other words, its privacy cost is completely independent of the number of queries.

We now show how this is proven in Clutch-DP. The implementation of RNM in RandML can be seen in [Figure 13](#), where qs is a list of queries, and $qs[i]$ returns its i -th element. The algorithm simply iterates over the list of queries, while keeping track of the highest result so far (maxA) as well as its index (maxI), and returns maxI at the end.

```
let RNM  $\varepsilon$  qs N db =
  let maxI = ref (-1) in
  let maxA = ref (0) in
  let rec rnm i =
    if i = N then !maxI
    else let a = Laplace ( $\varepsilon/2$ ) (qs[i] db) in
      if i = 0 || (!maxA) < a then
        maxA  $\leftarrow$  a;
        maxI  $\leftarrow$  i;
      else ();
      rnm (i + 1)
  in rnm 0
```

Fig. 13. The Report Noisy Max mechanism

We can prove that RNM satisfies the following specification:

$$\begin{aligned} & \{ \mathcal{F}^X(\varepsilon) * \text{Adj}(db, db') * \forall i. \text{1-sens}(qs[i]) \} \\ \forall qs N db db'. j. \quad & RNM \varepsilon qs N db \precsim RNM \varepsilon qs N db' \\ & \{ v v'. v = j \Rightarrow v' = j \} \end{aligned} \tag{14}$$

The precondition states that we own a privacy budget of $\mathcal{F}^X(\varepsilon)$, that the databases we are considering are adjacent, and that all queries in qs are 1-sensitive. This quadruple is stated in a pointwise manner: we show that for every possible outcome j if it is the case that the left-hand program returns j , then so does the right-hand program. From this specification, by applying the general version of the adequacy theorem ([Theorem A.6](#)), we can recover that RNM is externally ε -dp.

References

- [1] Chiké Abuah, David Darais, and Joseph P. Near. 2022. Solo: A Lightweight Static Analysis for Differential Privacy. *Proc. ACM Program. Lang.* 6, OOPSLA2, Article 150 (Oct. 2022). <https://doi.org/10.1145/3563313>
- [2] Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, Shin-ya Katsumata, and Tetsuya Sato. 2021. Higher-order probabilistic adversarial computations: categorical semantics and program logics. *Proc. ACM Program. Lang.* 5, ICFP, Article 93 (Aug 2021), 30 pages. <https://doi.org/10.1145/3473598>
- [3] Alejandro Aguirre, Philipp G. Haselwarter, Markus de Medeiros, Kwing Hei Li, Simon Oddershede Gregersen, Joseph Tassarotti, and Lars Birkedal. 2024. Error Credits: Resourceful Reasoning about Error Bounds for Higher-Order Probabilistic Programs. *Proc. ACM Program. Lang.* 8, ICFP, Article 246 (Aug 2024), 33 pages. <https://doi.org/10.1145/3674635>
- [4] Aws Albarghouthi and Justin Hsu. 2017. Synthesizing Coupling Proofs of Differential Privacy. *Proceedings of the ACM on Programming Languages* 2, POPL (Dec. 2017), 58:1–58:30. <https://doi.org/10.1145/3158146>
- [5] Jialu Bao, Emanuele D’Osualdo, and Azadeh Farzan. 2025. Bluebell: An Alliance of Relational Lifting and Independence for Probabilistic Reasoning. *Proc. ACM Program. Lang.* 9, POPL, Article 58 (Jan. 2025), 31 pages. <https://doi.org/10.1145/3704894>
- [6] Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, Léo Stefanesco, and Pierre-Yves Strub. 2015. Relational Reasoning via Probabilistic Coupling. In *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*. https://doi.org/10.1007/978-3-662-48899-7_27
- [7] Gilles Barthe, Noémie Fong, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2016. Advanced Probabilistic Couplings for Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016 (CCS ’16)*. 55–67. <https://doi.org/10.1145/2976749.2978391>
- [8] Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Pierre-Yves Strub. 2015. Higher-Order Approximate Relational Refinement Types for Mechanism Design and Differential Privacy. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL ’15)*. Association for Computing Machinery, New York, NY, USA, 55–68. <https://doi.org/10.1145/2676726.2677000>
- [9] Gilles Barthe, Marco Gaboardi, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2016. Proving Differential Privacy via Probabilistic Couplings. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science* (New York, NY, USA) (LICS ’16). Association for Computing Machinery, New York, NY, USA, 749–758. <https://doi.org/10.1145/2933575.2934554>
- [10] Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. 2012. Probabilistic relational reasoning for differential privacy. *SIGPLAN Not.* 47, 1 (Jan 2012), 97–110. <https://doi.org/10.1145/2103621.2103670>
- [11] Clément L. Canonne, Gautam Kamath, and Thomas Steinke. 2020. The Discrete Gaussian for Differential Privacy. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS ’20)*. Curran Associates Inc., Red Hook, NY, USA, 15676–15688.
- [12] Ethan Cowan, Michael Shoemate, and Mayana Pereira. 2024. *Hands-on Differential Privacy: Introduction to the Theory and Practice Using OpenDP* (first edition ed.). O’Reilly Media, Inc, Sebastopol, CA.
- [13] Arthur Azevedo de Amorim, Marco Gaboardi, Justin Hsu, and Shin-ya Katsumata. 2021. Probabilistic Relational Reasoning via Metrics. In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS ’19)*. IEEE Press, Vancouver, Canada, 1–19.
- [14] Markus de Medeiros, Muhammad Naveed, Tancrède Lepoint, Temesghen Kahsai, Tristan Ravitch, Stefan Zetzsche, Anjali Joshi, Joseph Tassarotti, Aws Albarghouthi, and Jean-Baptiste Tristan. 2025. Verified Foundations for Differential Privacy. *Artifact for Verified Foundations for Differential Privacy* 9, PLDI (June 2025), 191:1094–191:1118. <https://doi.org/10.1145/3729294>
- [15] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/11681878_14
- [16] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (Aug. 2014), 211–407. <https://doi.org/10.1561/0400000042>
- [17] Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2021. ReLoC Reloaded: A Mechanized Relational Logic for Fine-Grained Concurrency and Logical Atomicity. *Log. Methods Comput. Sci.* 17, 3 (2021). [https://doi.org/10.46298/lmcs-17\(3:9\)2021](https://doi.org/10.46298/lmcs-17(3:9)2021)
- [18] Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. 2013. Linear Dependent Types for Differential Privacy. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL ’13)*. Association for Computing Machinery, New York, NY, USA, 357–370. <https://doi.org/10.1145/2429069.2429113>
- [19] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2012. Universally Utility-maximizing Privacy Mechanisms. *SIAM J. Comput.* 41, 6 (Jan. 2012), 1673–1693. <https://doi.org/10.1137/09076828X>

[20] Simon Oddershede Gregersen, Alejandro Aguirre, Philipp G. Haselwarter, Joseph Tassarotti, and Lars Birkedal. 2024. Asynchronous Probabilistic Couplings in Higher-Order Separation Logic. *Proc. ACM Program. Lang.* 8, POPL (2024), 753–784. <https://doi.org/10.1145/3632868>

[21] Philipp G. Haselwarter, Kwing Hei Li, Alejandro Aguirre, Simon Oddershede Gregersen, Joseph Tassarotti, and Lars Birkedal. 2025. Approximate Relational Reasoning for Higher-Order Probabilistic Programs. *Proc. ACM Program. Lang.* 9, POPL, Article 41 (Jan. 2025), 31 pages. <https://doi.org/10.1145/3704877>

[22] Justin Hsu. 2017. *Probabilistic Couplings for Probabilistic Reasoning*. Ph. D. Dissertation. University of Pennsylvania. arXiv:1710.09951

[23] Seidu Inusah and Tomasz J. Kozubowski. 2006. A Discrete Analogue of the Laplace Distribution. *Journal of Statistical Planning and Inference* 136, 3 (March 2006), 1090–1102. <https://doi.org/10.1016/j.jspi.2004.08.014>

[24] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>

[25] Kelly Kostopoulou, Pierre Tholoniat, Asaf Cidon, Roxana Geambasu, and Mathias Lécuyer. 2023. Turbo: Effective Caching in Differentially-Private Databases. In *Proceedings of the 29th Symposium on Operating Systems Principles (SOSP '23)*. Association for Computing Machinery, New York, NY, USA, 579–594. <https://doi.org/10.1145/3600006.3613174>

[26] Ios Kotsogiannis, Yuchao Tao, Xi He, Maryam Fanaeepour, Ashwin Machanavajjhala, Michael Hay, and Gerome Miklau. 2019. PrivateSQL: A Differentially Private SQL Query Engine. *Proc. VLDB Endow.* 12, 11 (July 2019), 1371–1384. <https://doi.org/10.14778/3342263.3342274>

[27] Elisabet Lobo-Vesga, Alejandro Russo, and Marco Gaboardi. 2020. A Programming Framework for Differential Privacy with Accuracy Concentration Bounds. In *2020 IEEE Symposium on Security and Privacy (SP)*. 411–428. <https://doi.org/10.1109/SP40000.2020.00086>

[28] Elisabet Lobo-Vesga, Alejandro Russo, Marco Gaboardi, and Carlos Tomé Cortiñas. 2024. Sensitivity by Parametricity. *Paper Artifact: Sensitivity by Parametricity* 8, OOPSLA2 (Oct. 2024), 286:415–286:441. <https://doi.org/10.1145/3689726>

[29] Min Lyu, Dong Su, and Ninghui Li. 2017. Understanding the Sparse Vector Technique for Differential Privacy. *Proc. VLDB Endow.* 10, 6 (Feb. 2017), 637–648. <https://doi.org/10.14778/3055330.3055331>

[30] Miti Mazmudar, Thomas Humphries, Jiaxiang Liu, Matthew Rafuse, and Xi He. 2022. Cache Me If You Can: Accuracy-Aware Inference Engine for Differentially Private Data Exploration. *Proc. VLDB Endow.* 16, 4 (Dec. 2022), 574–586. <https://doi.org/10.14778/3574245.3574246>

[31] Ilya Mironov. 2012. On Significance of the Least Significant Bits for Differential Privacy. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS '12)*. Association for Computing Machinery, New York, NY, USA, 650–661. <https://doi.org/10.1145/2382196.2382264>

[32] Joseph P. Near and Chiké Abuah. 2021. *Programming Differential Privacy*. Vol. 1. <https://programming-dp.com/>

[33] Joseph P. Near, David Darais, Chike Abuah, Tim Stevens, Pranav Gaddamadugu, Lun Wang, Neel Soman, Mu Zhang, Nikhil Sharma, Alex Shan, and Dawn Song. 2019. Duet: An Expressive Higher-Order Language and Linear Type System for Statically Enforcing Differential Privacy. *Proc. ACM Program. Lang.* 3, OOPSLA (Oct. 2019). <https://doi.org/10.1145/3360598>

[34] Shangfu Peng, Yin Yang, Zhenjie Zhang, Marianne Winslett, and Yong Yu. 2013. Query Optimization for Differentially Private Data Management Systems. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. 1093–1104. <https://doi.org/10.1109/ICDE.2013.6544900>

[35] Jason Reed and Benjamin C. Pierce. 2010. Distance Makes the Types Grow Stronger: A Calculus for Differential Privacy. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP '10)*. Association for Computing Machinery, New York, NY, USA, 157–168. <https://doi.org/10.1145/1863543.1863568>

[36] Ryan M Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. 2016. Privacy Odometers and Filters: Pay-as-you-Go Composition. In *Advances in Neural Information Processing Systems*, Vol. 29. Curran Associates, Inc.

[37] Tetsuya Sato. 2016. Approximate Relational Hoare Logic for Continuous Random Samplings. In *The Thirty-second Conference on the Mathematical Foundations of Programming Semantics, MFPS 2016*. <https://doi.org/10.1016/JENTCS.2016.09.043>

[38] Tetsuya Sato, Gilles Barthe, Marco Gaboardi, Justin Hsu, and Shin-ya Katsumata. 2019. Approximate Span Liftings: Compositional Semantics for Relaxations of Differential Privacy. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. 1–14. <https://doi.org/10.1109/LICS.2019.8785668>

[39] Tetsuya Sato and Yasuhiko Minamide. 2025. Formalization of Differential Privacy in Isabelle/HOL. In *Proceedings of the 14th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP '25)*. Association for Computing Machinery, New York, NY, USA, 67–82. <https://doi.org/10.1145/3703595.3705875>

[40] Matías Toro, David Darais, Chike Abuah, Joseph P. Near, Damián Árquez, Federico Olmedo, and Éric Tanter. 2023. Contextual Linear Types for Differential Privacy. *ACM Trans. Program. Lang. Syst.* 45, 2 (May 2023), 8:1–8:69. <https://doi.org/10.1145/3589207>

- [41] Daniel Winograd-Cort, Andreas Haeberlen, Aaron Roth, and Benjamin C. Pierce. 2017. A Framework for Adaptive Differential Privacy. *Proc. ACM Program. Lang.* 1, ICFP (Aug. 2017), 10:1–10:29. <https://doi.org/10.1145/3110254>
- [42] Danfeng Zhang and Daniel Kifer. 2017. LightDP: Towards Automating Differential Privacy Proofs. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL '17)*. Association for Computing Machinery, New York, NY, USA, 888–901. <https://doi.org/10.1145/3009837.3009884>
- [43] Hengchu Zhang, Edo Roth, Andreas Haeberlen, Benjamin C. Pierce, and Aaron Roth. 2019. Fuzzi: A Three-Level Logic for Differential Privacy. *Prototype Implementation of Fuzzi: A Three-Level Logic for Differential Privacy* 3, ICFP (July 2019), 93:1–93:28. <https://doi.org/10.1145/3341697>