
GEOGRAPHIC SPINES IN THE 2020 CENSUS DISCLOSURE AVOIDANCE SYSTEM

Ryan Cumings-Menon,[†] Robert Ashmead,^{†, ‡} Daniel Kifer,^{†, ◊}
Philip Leclerc,[†] Jeffrey Ocker,[△] Michael Ratcliffe,[†] Pavel Zhuravlev,[†] John M. Abowd[▽]

[†] U.S. Census Bureau

[‡] The Ohio Colleges of Medicine Government Resource Center

[◊] Penn State University

[△] Federal Communications Commission, formerly U.S. Census Bureau

[▽] Cornell University, U.S. Census Bureau (retired)

ABSTRACT. The 2020 Census Disclosure Avoidance System (DAS) is a formally private mechanism that first adds independent noise to cross tabulations for a set of pre-specified hierarchical geographic units, which is known as the geographic spine. After post-processing these noisy measurements, the DAS outputs a formally private database with fields indicating location in the standard census geographic spine, which is defined by the United States as a whole, states, counties, census tracts, block groups, and census blocks. This paper describes how the geographic spine used internally within the DAS to define the initial noisy measurements impacts accuracy of the output of the DAS. Specifically, tabulations for geographic areas tend to be most accurate for geographic areas that both 1) can be derived by aggregating together the geographic units of the internal spine other than block geographic units, and 2) are closer to the geographic units of the internal spine. After describing the methods supported by the DAS for defining the internal DAS geographic spine, we provide the settings used to define the 2020 Census production DAS executions. To demonstrate the accuracy impact of the choice of internal spine, we also provide accuracy metrics for three DAS executions using settings similar to the ones used for the 2020 Census Redistricting Data (P.L. 94-171) persons production DAS execution, but using three different choices of internal geographic spine.

Key words and phrases: Differential Privacy, 2020 Census, TopDown Algorithm.

*We thank Randall Akee and Norman DeWeaver for providing summary metrics that motivated the exploration of alternative geographic spines within the DAS [2]. We also thank Simson Garfinkel and Nicholas Nagle for early contributions to the concept of using alternative geographic spines within the DAS. The authors are grateful for helpful comments from Justin Doty, Brian Finley, Mark Fleischer, Sallie Keller, and the participants of the DIMACS and Rutgers University Department of Statistics' Workshop on the Analysis of Census Noisy Measurement Files and Differential Privacy on April 27 – 29, 2022. The views expressed in this technical paper are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data (Project No. 7502798, Disclosure Review Board (DRB) approval number: CBDRB-FY24-CED005-0001).

1. INTRODUCTION

The U.S. Census Bureau’s 2020 Disclosure Avoidance System (DAS) was used to produce a formally private file in support of the 2020 Census Redistricting Data (P.L. 94-171) Summary File (hereafter redistricting data) and the 2020 Demographic and Housing Characteristics File (DHC). As described by [1, 5] in more detail, the DAS takes the confidential Census Edited File (CEF) as input, which is a set of microdata files with components that include data on individuals, housing units, group quarters (GQs), and households; the output of the DAS is called the Microdata Detailed File (MDF). We call the DAS formally private because it provides privacy guarantees based on either differential privacy (DP) [8] or zero-concentrated differential privacy (zCDP) [3], as we describe in Section 2 in more detail.

Each DAS execution begins by converting the microdata in the CEF to a histogram, *i.e.*, a flattened and fully saturated contingency table, for each census block. Afterward, histograms for less granular *geographic units*, or *geounits*, are created by adding these block level histograms to one another. For example, in early versions of the DAS, these block histograms were added to one another to define the histograms for each census block group. This recursion was repeated to define the histograms for each census tract, county, and state geographic units, and also, in US DAS executions, a final geounit corresponding to the U.S. as a whole.¹ The collection of all of these hierarchical geographic units is known as a *geographic spine*, or a *spine*, and this particular spine is known as the conventional geographic spine, as defined by the Census Bureau Geography Division. The geounits in the geographic spine are split into several levels of granularity, which are known as *geographic levels*, or *geolevels*; for example, each geounit in the conventional geographic spine is a member of either the U.S., state, county, census tract, block group, or the census block geolevels. In other words, the spine can be viewed as a rooted tree, with a root vertex given by the U.S. geounit. This connection with rooted trees also motivates the terms we use to describe adjacency relationships on the spine, such as parent and child geounits. For example, for the conventional spine, the child geounits of the U.S. geounit are the geounits in the state geographic level.

To see the importance of the internal spine used within the DAS, a general overview of what takes place during a DAS execution will be helpful. After the CEF-based histograms of each geounit on the internal spine are computed, the DAS uses a formally private implementing mechanism to observe noisy measurements for each geounit, which are defined as tabulations, *i.e.*, counts within various population groups within the geounit (such as population counts within each race group), added to realizations of mean-zero random variables. Afterward, the DAS uses the noisy measurements of the U.S. geounit, to compute a formally private data histogram for this geounit. Next, the DAS uses the noisy measurements of progressively more granular (lower) geolevels on the spine to estimate formally private data histograms for these geounits, subject to the constraint that these histograms are consistent with the estimates of the next-higher (parent) geolevels. After this recursion is run down to the block geolevel, the output MDF is defined by converting the block histogram estimates to microdata. For more detail on the DAS, including the optimization problems used to define the histogram estimates of each geounit, see [1, 5].

¹Aside from US DAS executions discussed in this paragraph, there are also Puerto Rico (PR) DAS executions, which have a least granular geounit at the state geographic level that corresponds to the Commonwealth of Puerto Rico as a whole.

Note that the main topic of this paper is the geographic spine that is used internally in the DAS. Even when this geographic spine does not correspond to the conventional spine, the published statistical tabulations based on the DAS output are still provided for geounits on the conventional spine, defined using the Census Bureau Geography Division’s standard definitions.

1.1. Contributions. The conventional spine has two primary shortcomings when used for the internal spine in the DAS. First, legally defined American Indian/Alaska Native/Native Hawaiian (AIAN) tribal areas are usually far from this spine, meaning that many on-spine geounits must be added to or subtracted from one another to compose these geographic areas. Being distant from the spine often results in query count estimates with greater mean squared error than on-spine geounits. Second, the conventional spine also places many other important legal, political, or census-designated off-spine entities (OSEs) far from the spine, *e.g.*, minor civil divisions (MCDs) and incorporated places.

In order to ameliorate these two issues, the Census Bureau experimented with using alternative spines within the DAS. The purpose of this paper is to describe these alternative spines, the methods we used to define them, and the spines used within the DAS for the production 2020 redistricting and DHC DAS implementations.

The starting point for these alternative spines is a spine called an *AIAN spine*. The goal for these spines is to ameliorate the issue of AIAN tribal areas being far from the spine. To define an AIAN spine, first let the *Aggregated AIAN Areas OSE* be defined as the collection of all blocks that are within the AIAN areas. An AIAN spine is defined by replacing geounits in the state geolevel and below with up to two geounits. Specifically, each such original geounit is replaced by geounits defined as the portion of this geounit that is inside, and the portion that is outside, of the Aggregated AIAN Areas OSE.² For example, the geounit corresponding to the state of Arizona would be replaced with two state geolevel geounits in an AIAN spine: 1) the portion of Arizona that is inside of the AIAN areas, and 2) the portion of Arizona that is outside of the AIAN areas.

This process of splitting a spine into an AIAN and a non-AIAN branch at the state geolevel and below can be applied to a variety of different spines. In the case of the 2020 redistricting DAS implementation, we apply this process to the conventional spine; a graphical summary of the 2020 redistricting AIAN spine is shown in Figure 1. The 2020 DHC AIAN spine is defined by the output of this same process applied to an alternative input spine, as described in Section 6 in more detail.

This paper describes several optimization heuristics that perform updates to an AIAN spine to enhance specific features of the spine. We call geographic spines that are output from these spine optimization methods *optimized spines*. The spine optimization routines update an input spine in two stages. In the first stage, certain geolevels are redefined to bring a given set of OSEs closer to the spine. For example, all of the spine optimization methods the Census Bureau has experimented with involve redefining the block groups in an AIAN spine as *optimized block groups*. We also describe methods that add additional geolevels, *e.g.*, a geolevel between tracts and counties known as the tract group geolevel, to provide a further improvement in the accuracy of tabulations for the set of target OSEs. After the

²In cases in which one of these two portions of the original geounit do not contain any blocks, the original geounit is replaced with a single geounit. For example, since there are not any AIAN areas in Kentucky, an AIAN spine would not include a geounit corresponding to the AIAN portion of Kentucky, but there would be a geounit that corresponds to the non-AIAN portion of Kentucky.

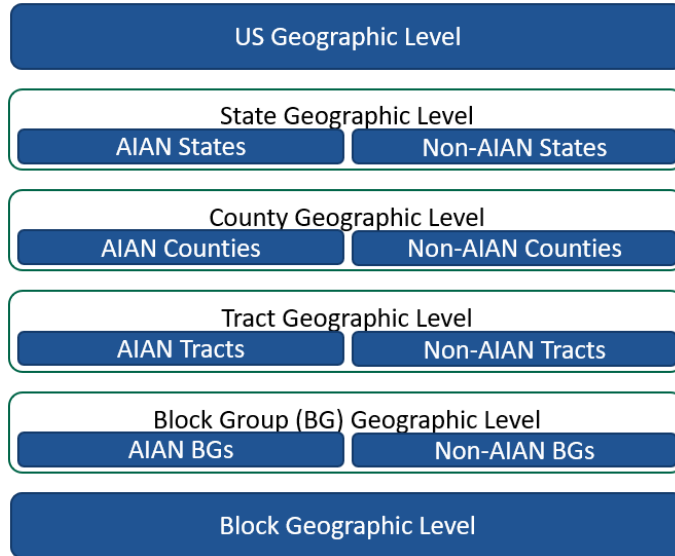


Figure 1: The redistricting AIAN spine, shown without the optional tract group geographic level between the tract and county geographic levels. This spine uses the standard census spine as a starting point and includes an AIAN branch. The optimized spine uses an AIAN spine as a starting point and redefines block groups as optimized block groups.

initial definitions of the geounits have been found in the first stage of the spine optimization routines, the variance of the final histogram estimate is reduced in the second stage of spine optimization by removing certain redundant geounits, which we call the bypassing step. For example, in cases in which a parent geounit has only one child geounit, this decision rule results in bypassing the parent geounit, *i.e.*, removing the parent geounit from the spine, which, as we will see below in more detail, also allows us to decrease the variance of the noisy measurements of the child geounit. Since the consistency-with-parent constraints used in the DAS require that the count estimates of the child are equal to those of the parent in cases in which a parent geounit has only one child, bypassing the parent geounit in this way increases accuracy of the child’s histogram estimate by reducing the variance of the noisy measurements that are used to derive this estimate.

In addition to describing the spine optimization routines that are supported by the DAS, we also provide the spine optimization settings of the 2020 production redistricting and DHC DAS executions. To demonstrate the differences in the accuracy provided by the conventional spine, an AIAN spine, and an optimized spine, we compare accuracy metrics for DAS executions that use each of these three spines.

1.2. Paper Outline. The heuristics used in the spine optimization routines are based on existing results from the DP literature, which are provided in Section 2. Section 3 describes how geographic spines are represented in this paper using matrices and provides a result on the privacy guarantees of the DAS when the internal spine is defined as either the conventional spine or an AIAN spine. The methods used in the first stage of spine optimization, *i.e.*, the methods that redefine certain geolevels to bring a given set of OSEs

closer to the spine, are described in Section 4. Section 5 describes the bypassing method used in the second stage and provides a result on the privacy guarantees of the DAS when an optimized spine is used. Section 6 describes the geographic spines that were used in the 2020 production DAS executions. Section 7 concludes the paper with a comparison of accuracy metrics for redistricting persons DAS executions with an internal spine defined by either the conventional spine, an AIAN spine, or an optimized spine. The next subsection outlines the notation used throughout this paper.

1.3. Notation. Throughout the paper we denote matrices using capitalized font, and vectors using lower case bold font. We denote the Kronecker product of real matrices A, B by $A \otimes B$, and elementwise division, when these matrices are conformable, by $A \oslash B$ respectively. We denote the elementwise absolute value of the real matrix A by $|A|$. A length N column vector with each element equal to $c \in \mathbb{R}$ is denoted by \mathbf{c}_N , and, when there is little risk of confusion, we omit the subscript N . We use I_n to denote the $n \times n$ identity matrix. We denote the i^{th} row of $A \in \mathbb{R}^{m \times n}$ by $A[i, \cdot]$ and the i^{th} column by $A[\cdot, i]$. Given two datasets (possibly with unequal numbers of rows) x, x' , let $d_{\mathcal{H}}(x, x')$ denote the minimum number of set addition and subtraction operations with singleton sets required to transform the dataset x to be equal to the dataset x' . We also denote the set of nonnegative real numbers by \mathbb{R}_+ , the positive reals by \mathbb{R}_{++} , and the nonnegative integers by \mathbb{Z}_+ . The number of elements in the set A will be denoted by $\text{Card}(A)$.

We make use of several probability distributions, but two are worth pointing out explicitly, as their definitions are less standard. First, let the discrete Gaussian random variable be defined so that its probability mass function, $f : \mathbb{Z} \rightarrow \mathbb{R}$, is given by $f(x) \propto \exp(-(x - \mu)^2 / (2\sigma^2))$, and denote the distribution of this random variable by $N_{\mathbb{Z}}(\mu, \sigma^2)$; see for example, [4]. In contrast, to denote the (continuous) Gaussian distribution we use the standard notation, $N(\mu, \sigma^2)$. We also define the discrete Laplace random variable so that its probability mass function, $f : \mathbb{Z} \rightarrow \mathbb{R}$, is given by $f(x) \propto \exp(-|x - \mu|/b)$, and denote the distribution of this random variable by $\text{Laplace}_{\mathbb{Z}}(\mu, \sigma^2)$; see for example, [9]. We also denote the distributions of length n column vectors with each element $i \in \{1, \dots, n\}$ independently distributed as $N_{\mathbb{Z}}(\boldsymbol{\mu}[i], \mathbf{c}[i])$, $\text{Laplace}(\boldsymbol{\mu}[i], \mathbf{c}[i])$, and $\text{Laplace}_{\mathbb{Z}}(\boldsymbol{\mu}[i], \mathbf{c}[i])$, where $\boldsymbol{\mu} \in \mathbb{R}^n$ and $\mathbf{c} \in \mathbb{R}_+^n$, by $N_{\mathbb{Z}}(\boldsymbol{\mu}, \text{diag}(\mathbf{c}))$, $\text{Laplace}(\boldsymbol{\mu}, \text{diag}(\mathbf{c}))$, and $\text{Laplace}_{\mathbb{Z}}(\boldsymbol{\mu}, \text{diag}(\mathbf{c}))$ respectively.

2. DIFFERENTIAL PRIVACY

2.1. Privacy Definitions. The definitions of differential privacy and ρ -zCDP are provided by [7, 8, 3], and these definitions are also restated below. Both definitions can be viewed as a bound on the difference between the distribution of a mechanism's output when the input is a given database and when the input is a *neighboring database*. The specific notion of difference used in these definitions, along with their required upper bounds, is generally self-evident, but more care is required for the definition of a neighboring database because there are multiple definitions for this term that are commonly used. For example, two databases are said to be *unbounded neighbors* if one database can be derived by adding *or* removing a record from the other database; likewise, two databases are said to be *bounded neighbors* if one database can be derived from the other by adding one record *and* removing one record [12]. In other words, the set of bounded neighbors can be written as $\{x, x' \in \mathcal{X}^d \mid d_{\mathcal{H}}(x, x') = 2\}$,

where \mathcal{X}^d is the set of databases with d records, and the set of unbounded neighbors as $\{x, x' \mid d_{\mathcal{H}}(x, x') = 1\}$.

The definition of neighbor has important implications because the upper bounds provided in the privacy guarantees below can be translated into an upper bound on the power of any statistical test with a given significance level of the null hypothesis that the input database is x , versus the alternative that it is a neighbor of x [6, 15]. In other words, using bounded neighbors does not provide a theoretical bound on the accuracy of inferences on the number of records in the input database. However, in the case of tests with a null hypothesis that the database is x and an alternative hypothesis that it is a given bounded neighbor of x , formal privacy definitions that use bounded neighbors provide stronger limitations on the accuracy of privacy attackers' inferences. This can be shown for the privacy definitions below using the fact that a bounded neighbor of a given database is also an "unbounded neighbor of an unbounded neighbor" of the database.

The privacy guarantees provided by the DAS use a unique definition of neighbor that does not protect against certain inferences. Specifically, the DAS does not protect each state's total population or the locations of each group quarters type. In other words, the DAS outputs a database with state total populations that are identical to those of the input database, and, since the frame of a U.S. decennial census excludes addresses that are unoccupied group quarters, at least one person must reside in each group quarter that appears in the census.³ The population residing in each group quarter type in a given geounit is constrained to be at least the number of group quarters of that type within the geounit. In addition to these data-dependent constraints, the DAS also imposes data-independent constraints called imputation and edit constraints. For example, respondents that reside in college/university student housing are required to be between the ages of 16 and 65 and respondents in nursing facilities are required to be at least 20 years of age. The imputation and edit constraints imposed on the CEF ensure that each row of the CEF satisfies these constraints, so we also impose these data-independent imputation and edit constraints on the output MDF produced by the DAS. For more detail on the data-dependent and data-independent constraints imposed by the DAS, see [1, 5].

For this reason, in the context of the DAS, two databases are said to be neighbors if one database can be derived from the other by adding one record to a given state and then removing one record from that same state, such that both databases also satisfy the group quarters invariant constraints and the imputation and edit constraints of the input database of the DAS. To emphasize this definition of neighbors in the presence of invariant constraints in our privacy definitions below, we define the set of pairs of neighboring databases as $\{x, x' \in \mathcal{U} \mid d_{\mathcal{H}}(x, x') = 2\}$, where \mathcal{U} is a given universal set of databases. For example, this notation is a generalization of the definition of bounded neighbors, which used $\mathcal{U} = \mathcal{X}^d$. In the results throughout this paper that use these privacy definitions, we denote the set of databases that satisfy these group quarter invariant constraints as well as the imputation and edit constraints of the confidential input data as \mathcal{G} and the set of databases with $d[i]$ records in each state i as \mathcal{X}^d . Using this notation, the definition of neighbors used by the DAS can be written as, $\{x, x' \in \mathcal{X}^d \cap \mathcal{G} \mid d_{\mathcal{H}}(x, x') = 2\}$.

³That is, the definition of a living quarter eligible for inclusion in the census of population and housing does not include vacant group quarters. By contrast, housing units may be either occupied or vacant.

Definition 1. (Differential Privacy [8]) A randomized algorithm $\mathcal{M} : \mathcal{U} \rightarrow \mathcal{Y}$ satisfies ϵ -differential privacy if, for all $(x, x') \in \{x, x' \in \mathcal{U} \mid d_{\mathcal{H}}(x, x') = 2\}$ and all $E \subset \mathcal{Y}$, we have $P(\mathcal{M}(x) \in E) \leq \exp(\epsilon)P(\mathcal{M}(x') \in E)$.

The definition of ρ -zCDP [3] is provided below; this privacy framework is used for privacy-loss budget (PLB) accounting in both the 2020 redistricting data file and the DHC production DAS executions. Kifer et al. [11] provide more information on privacy semantic statements that can be made for the DAS.

Definition 2. (Zero-Concentrated Differential Privacy (zCDP) [3]) A randomized mechanism $\mathcal{M} : \mathcal{U} \rightarrow \mathcal{Y}$ is ρ -zero-concentrated differentially private (ρ -zCDP) if, for all $x, x' \in \{x, x' \in \mathcal{U} \mid d_{\mathcal{H}}(x, x') = 2\}$ and all $\alpha \in (1, \infty)$,

$$D_{\alpha}(\mathcal{M}(x) \parallel \mathcal{M}(x')) \leq \rho\alpha,$$

where $D_{\alpha}(P \parallel Q) = \log(\sum_{E \in \mathcal{Y}} P(E)^{\alpha} Q(E)^{1-\alpha}) / (\alpha - 1)$ is the Rényi divergence of order α between the distributions P and Q .

2.2. Privacy-Loss Accounting Results. A general pattern in our results showing that the DAS is formally private for various spine settings is to first establish that the DP implementing mechanism for each query i is either ϵ_i -DP or ρ_i -zCDP, and then use parallel and sequential composition to show that the combination of all of the DP implementing query answers used in the DAS are $\sum_i \epsilon_i$ -DP or $\sum_i \rho_i$ -zCDP. This section contains these intermediate results that are used in these proofs. The first two results are used in the next subsection to show that each implementing mechanism is either ϵ -DP or ρ -zCDP.

Lemma 1. (Theorem 3.2 [9]) Let $\Delta, \epsilon > 0$. Let $q : \mathcal{U} \rightarrow \mathbb{Z}$ satisfy $|q(x) - q(x')| \leq \Delta$ for all $x, x' \in \{x, x' \in \mathcal{U} \mid d_{\mathcal{H}}(x, x') = 2\}$. Define a randomized algorithm by $\mathcal{M}(x) = q(x) + Y$ where either $Y \sim \text{Laplace}_{\mathbb{Z}}(0, \Delta/\epsilon)$ or $Y \sim \text{Laplace}(0, \Delta/\epsilon)$. Then \mathcal{M} satisfies ϵ -DP.

Lemma 2. (Theorem 4 [4]) Let $\Delta, \epsilon > 0$. Let $q : \mathcal{U} \rightarrow \mathbb{Z}$ satisfy $|q(x) - q(x')| \leq \Delta$ for all $x, x' \in \{x, x' \in \mathcal{U} \mid d_{\mathcal{H}}(x, x') = 2\}$. Define a randomized algorithm by $\mathcal{M}(x) = q(x) + Y$ where either $Y \sim N_{\mathbb{Z}}(0, \Delta^2/(2\rho))$ or $Y \sim N(0, \Delta^2/(2\rho))$. Then \mathcal{M} satisfies ρ -zCDP.

While providing a comprehensive compilation of the properties that are implied by a mechanism satisfying either ϵ -DP or ρ -zCDP is beyond the scope of this paper, a few of these properties are worth explicitly pointing out. First, these privacy guarantees share the property that they are invariant to post-processing. In other words, if $\mathcal{M} : \mathcal{U} \rightarrow \mathcal{Y}$ satisfies either of these privacy guarantees, then so does the mechanism $f \circ \mathcal{M}$, where $f : \mathcal{Y} \rightarrow \mathcal{Z}$. Second, if $\mathcal{M}_1(x)$ is ϵ_1 -DP (respectively, ρ_1 -zCDP) and $\mathcal{M}_2(x)$ is ϵ_2 -DP (ρ_2 -zCDP) then releasing the output of $\mathcal{M}_1(x)$ and $\mathcal{M}_2(x)$ simultaneously is itself $(\epsilon_1 + \epsilon_2)$ -DP ($(\rho_1 + \rho_2)$ -zCDP), which is called sequential composition. Also, when $\mathcal{M}_1(x), \mathcal{M}_2(x), \dots, \mathcal{M}_k(x)$ only depend on disjoint subsets of the input dataset x , releasing the outputs of these mechanisms simultaneously is $(\max\{\epsilon_1, \epsilon_2\})$ -DP (respectively, $(\max\{\rho_1, \rho_2\})$ -zCDP) when using unbounded neighbors and $\max_{i \neq j} \epsilon_i + \epsilon_j$ (respectively, $(\max_{i \neq j} \rho_i + \rho_j)$ -zCDP) when using bounded neighbors, which is known as parallel composition. The following lemma uses sequential and parallel composition to provide privacy guarantees for a mechanism that is defined by multiple univariate mechanisms described in Lemmas 1 and 2.

Lemma 3. (Theorem 14 [4], Proposition 1 [7], Theorem 3.2 [9]) Suppose $x, x' \in \mathcal{X}^d$ differ on a single entry, and let a randomized algorithm $\mathcal{M} : \mathcal{X}^d \rightarrow \mathcal{Y}^m$ be defined by $\mathcal{M}(x) = q(x) + \mathbf{y}$ where $q : \mathcal{X}^d \rightarrow \mathcal{Y}^m$ and \mathbf{y} is an m dimensional column vector of independent random variables. Then we have the following.

- (1) Let $\mathbf{b} \in \mathbb{R}_{++}^m$, $\epsilon > 0$, and, for all neighbors $x, x' \in \mathcal{X}^d$, suppose $\sum_{j=1}^m |q(x)[j] - q(x')[j]|/\mathbf{b}[j] \leq \epsilon$. If, for all $j \in \{1, \dots, m\}$, either $\mathbf{y}_j \sim \text{Laplace}_{\mathbb{Z}}(\mathbf{0}, \mathbf{b}[j])$ or $\mathbf{y}_j \sim \text{Laplace}(\mathbf{0}, \mathbf{b}[j])$, then \mathcal{M} satisfies ϵ -DP.
- (2) Let $\boldsymbol{\sigma}^2 \in \mathbb{R}_{++}^m$, $\rho > 0$, and, for all neighbors $x, x' \in \mathcal{X}^d$, suppose $\sum_{j=1}^m \frac{(q(x)[j] - q(x')[j])^2}{\boldsymbol{\sigma}^2[j]} \leq 2\rho$. If, for all $j \in \{1, \dots, m\}$, either $\mathbf{y}_j \sim \text{N}_{\mathbb{Z}}(\mathbf{0}, \boldsymbol{\sigma}^2[j])$ or $\mathbf{y}_j \sim \text{N}(\mathbf{0}, \boldsymbol{\sigma}^2[j])$, then \mathcal{M} satisfies ρ -zCDP.

2.3. Linear Queries and Marginal Query Groups. The notation for each query $q(\cdot)$ in the preceding subsections is more general than required to describe the DAS. Specifically, the DAS uses only linear queries. We define a *linear query* as a linear map represented by a length n vector, and a linear query answer as the inner product of a linear query and the histogram cell counts vector. Since we only consider linear queries from this point forward, we sometimes refer to a linear query simply as a *query* below. We also define a *query matrix* as the matrix representation of a collection of linear queries, *i.e.*, a matrix defined by vertically stacking the length n row vector representations of the collection of linear queries together.

All elements of the linear queries used within the DAS have elements that are either zero or one, and they are defined so that each linear query answer provides an individual count for a marginal of the full histogram. It is convenient to group the queries providing counts for the same marginal together, so we will define a *query group* as the linear queries that provide counts for the same marginal, vertically stacked on top of one another. For example, suppose the schema of the original database is CENRACE \times HISPANIC \times VOTINGAGE, where CENRACE, HISPANIC, and VOTINGAGE indicate one of 63 census race combination categories, one of two ethnicity categories, and one of two age categories, respectively.⁴ In this case, the CENRACE \times VOTINGAGE query group matrix is defined as the matrix

$$Q = I_{63} \otimes \mathbf{1}_2^\top \otimes I_2,$$

and right multiplying Q by the vector of histogram cell counts of a database provides the CENRACE \times VOTINGAGE query group answers. More generally, the query matrix of a given query group is defined as a matrix that can be written as a series of Kronecker products of identity matrices and row vectors of ones. Note that this definition encompasses cases in which each of these Kronecker factors are all either row vectors of ones, *i.e.*, the total population query group matrix, or are all identity matrices, *i.e.*, the detailed cell query group matrix.

One of the inputs of the DAS is the proportion of PLB allocated to each query i of a given geolevel, which we denote by $\alpha_i \in \mathbb{R}_{++}$. The following lemma uses this notational convention, while only considering a single geounit, to provide the ϵ -DP and ρ -zCDP privacy

⁴While we are using attributes that are part of the census, note that this example, and the queries in this section, are simply described in the context of a hypothetical survey that does not use a geographic spine. The linear query notation introduced here is extended in the next section to account for the presence of the geographic spine.

guarantees of releasing the output of the implementing mechanisms of this geounit.⁵ This is used in the next section to provide the privacy guarantees of the DAS when either the conventional spine or an AIAN spine is used. Note that this lemma also makes use of a notational convention that we use throughout the rest of the paper; the histogram of detailed cell counts computed on the confidential dataset $x \in \mathcal{X}^d$ is denoted by \mathbf{x} .

Lemma 4. If each $Q[i] \in \{Q[i]\}_{i=1}^q$ is a query group matrix of dimension $m[i] \times n$ and the mechanism $\mathcal{M}(\mathbf{x})$, outputs $\{Q[i]\mathbf{x} + \mathbf{y}_i\}_{i=1}^q$, where \mathbf{x} is a vector of the histogram counts for $x \in \mathcal{X}^d$, and \mathbf{y}_i is a length $m[i]$ column vector of independent random variables, then we have the following.

(1) If, for each $i \in \{1, \dots, q\}$, either

$$\mathbf{y}_i \sim \text{Laplace}(\mathbf{0}, \text{diag}(\mathbf{2}_{m[i]}/(\epsilon\alpha_i))) \text{ or } \mathbf{y}_i \sim \text{Laplace}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{2}_{m[i]}/(\epsilon\alpha_i))),$$

where $\sum_i \alpha_i = 1$, then $\mathcal{M}(\cdot)$ satisfies ϵ -DP.

(2) If, for each $i \in \{1, \dots, q\}$, either

$$\mathbf{y}_i \sim \text{N}(\mathbf{0}, \text{diag}(\mathbf{1}_{m[i]}/(\rho\alpha_i))) \text{ or } \mathbf{y}_i \sim \text{N}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{1}_{m[i]}/(\rho\alpha_i))),$$

where $\sum_i \alpha_i = 1$, then $\mathcal{M}(\cdot)$ satisfies ρ -zCDP.

Proof. Since, for each $i \in \{1, \dots, q\}$, only one element of each column of $Q[i]$ is nonzero, we use Lemma 3 to leverage parallel composition to prove both cases. This reduces the problem to carrying out sequential composition of the set of mechanisms, $\{\mathcal{M}[i](\mathbf{x})\}_i^q$, where $\mathcal{M}[i](\mathbf{x}) := Q[i]\mathbf{x} + \mathbf{y}_i$. In both cases, we use the fact that $|Q[i]\mathbf{x} - Q[i]\mathbf{x}'|$, where \mathbf{x}, \mathbf{x}' are vectors of histogram cell counts of databases differing on a single entry, is a vector with at most two elements that are equal to one, with the remaining elements equal to zero.

In the first case, since $\sum_j |Q[i][j, \cdot]\mathbf{x} - Q[i][j, \cdot]\mathbf{x}'| \leq 2$, the first case in Lemma 3 implies $\mathcal{M}[i](\mathbf{x})$ is $\alpha_i\epsilon$ -DP. Thus, using sequential composition, releasing the output of $\{\mathcal{M}[i](\mathbf{x})\}_i$ is $\sum_i \alpha_i\epsilon = \epsilon$ -DP.

In the second case, since $\sum_j |Q[i][j, \cdot]\mathbf{x} - Q[i][j, \cdot]\mathbf{x}'|^2 \leq 2$, the second case in Lemma 3 implies $\mathcal{M}[i](\mathbf{x})$ is $\alpha_i\rho$ -zCDP. Thus, using sequential composition, releasing the output of $\{\mathcal{M}[i](\mathbf{x})\}_i$ is $\sum_i \alpha_i\rho = \rho$ -zCDP. \square

2.4. The Matrix Mechanism. This section describes a class of random mechanisms known as matrix mechanisms in the context of linear queries composed of query groups [13]. Using the notation introduced in Lemma 4, we assume that the errors of the i^{th} query group are random variables distributed as either $\mathbf{y}_i \sim \text{Laplace}(\mathbf{0}, \mathbf{2}_{m[i]}/(\epsilon\alpha_i))$, in the case of ϵ -DP, or $\mathbf{y}_i \sim \text{N}(\mathbf{0}, \mathbf{1}_{m[i]}/(\rho\alpha_i))$, in the case of ρ -zCDP, where $\alpha_i \in \mathbb{R}_{++}$ satisfies $\sum_i \alpha_i = 1$. Also, let the *workload* matrix $W \in \{0, 1\}^{m \times n}$ be defined as $W := \text{stack}(\{Q[i]\}_{i=1}^q)$, the errors vector as $\mathbf{y} := \text{stack}(\{\mathbf{y}_i\}_i^q)$, the PLB proportions vector as $\boldsymbol{\alpha} := \text{stack}(\{\mathbf{1}_{m[i]}\alpha_i\}_{i=1}^q)$, and the response variable as $\mathbf{z} := W\mathbf{x} + \mathbf{y}$, which is also the vertically stacked output of the

⁵Alternatively, for datasets that contain an attribute identifying each respondent's geounit for a given geolevel, this theorem can also be used to provide the privacy guarantees of releasing the output of all of the implementing DP mechanisms for the geolevel. However, depending on how the operation of bypassing geounits is defined, this can be less straightforward in the case of the optimized spine because there may be respondents in the dataset that are not included in any geounit in the geolevel. This alternative interpretation of the Lemma is discussed in more detail in Sections 3, and the operation of bypassing geounits is defined in 5.

mechanism described in Lemma 4. We assume that W has full column rank, which can be ensured by including the detailed cell query group in W , *i.e.*, the identity matrix.

For both of the distributions that we consider in this section, we define an alternative response variable that is observationally equivalent to \mathbf{z} , but with homoscedastic errors. For example, in the case of ϵ -DP, we have $\mathbf{y} \sim \text{Laplace}(\mathbf{0}, \text{diag}(\mathbf{2} \odot (\epsilon\boldsymbol{\alpha})))$, so we can define the rescaled errors as $\tilde{\mathbf{y}} := \text{diag}(\epsilon\boldsymbol{\alpha}/2)\mathbf{y}$, which are distributed as $\tilde{\mathbf{y}} \sim \text{Laplace}(\mathbf{0}, I)$, the rescaled workload as $\tilde{W} := \text{diag}(\epsilon\boldsymbol{\alpha}/2)W$, and the rescaled response variable as $\tilde{\mathbf{z}} := \text{diag}(\epsilon\boldsymbol{\alpha}/2)\mathbf{z} = \tilde{W}\mathbf{x} + \tilde{\mathbf{y}}$. In the case of ρ -zCDP, we have $\mathbf{y} \sim \text{N}(\mathbf{0}, \text{diag}(\mathbf{1} \odot (\rho\boldsymbol{\alpha})))$, so we can define the rescaled errors as $\tilde{\mathbf{y}} := \text{diag}(\sqrt{\rho\boldsymbol{\alpha}})\mathbf{y}$, which are distributed as $\tilde{\mathbf{y}} \sim \text{N}(\mathbf{0}, I)$, the rescaled workload as $\tilde{W} := \text{diag}(\sqrt{\rho\boldsymbol{\alpha}})W$, and the rescaled response variable as $\tilde{\mathbf{z}} := \text{diag}(\sqrt{\rho\boldsymbol{\alpha}})\mathbf{z} = \tilde{W}\mathbf{x} + \tilde{\mathbf{y}}$.

One simple example of a matrix mechanism is a mechanism that releases the weighted least squares estimates of $W\mathbf{x}$. Specifically, this can be done by first estimating \mathbf{x} as $\hat{\mathbf{x}} := \arg \min_{\mathbf{x}} \|\tilde{W}\mathbf{x} - \tilde{\mathbf{z}}\|_2^2 = (\tilde{W}^\top \tilde{W})^{-1} \tilde{W}^\top \tilde{\mathbf{z}}$, and then defining the output of the mechanism as $W\hat{\mathbf{x}}$. Note that, in either the case of ϵ -DP or ρ -zCDP, the privacy guarantee of the final mechanism follows from the fact that a mechanism that released $\tilde{\mathbf{z}}$ would satisfy the same privacy guarantee, along with the invariance to post-processing property.

This simple mechanism can be generalized by making a distinction between the linear query answers provided as output and the linear queries used to estimate $\hat{\mathbf{x}}$. Specifically, let the *strategy* matrix $A \in \{0, 1\}^{p \times n}$ be defined as $A := \text{stack}(\{R[i]\}_i^r)$, where each $R[i] \in \{R[i]\}_i^r$ is a query group matrix, and the query group PLB proportions as $\{\gamma_i\}_i^r$. We also assume that the strategy matrix has full column rank. In the same manner as described above for W , we define the rescaled strategy matrix as $\tilde{A} := \text{diag}(\epsilon\boldsymbol{\gamma}/2)A$, when deriving an ϵ -DP mechanism, or $\tilde{A} := \text{diag}(\sqrt{\rho\boldsymbol{\gamma}})A$, when deriving a ρ -zCDP mechanism. Likewise, the rescaled error vector and response variable are defined as above so that $\tilde{\mathbf{z}} := \tilde{A}\mathbf{x} + \tilde{\mathbf{y}}$. The final output of this matrix mechanism is then given by this alternative estimate of the linear queries in the workload W , which are $W\hat{\mathbf{x}} = W(\tilde{A}^\top \tilde{A})^{-1} \tilde{A}^\top \tilde{\mathbf{z}}$.⁶

Note that the variance matrix of this output vector is given by,

$$\text{Var}(W\hat{\mathbf{x}}) = W(\tilde{A}^\top \tilde{A})^{-1} W^\top.$$

Past work focuses on using this variance matrix to find strategies that provide a low expected sum of squared errors, which is given by $\text{Trace}(\text{Var}(W\hat{\mathbf{x}})) = \text{Trace}(W^\top W(\tilde{A}^\top \tilde{A})^{-1})$; see for example, [13, 14]. We use an alternative approach to motivate the heuristic used to bypass geounits in Section 5, after introducing how we represent the spine using matrices in the next section.

3. REPRESENTING LINEAR QUERIES ON THE SPINE

In this section we describe how we represent the workload and the strategy matrices that include the linear queries for each geounit on the spine. To do so, we first consider the case in which the same query groups are used in all such geounits, and generalize this notation afterward to the case where the query matrix is dependent on the geolevel.

Let the strategy matrix of the linear queries for each geounit be denoted by $B \in \{0, 1\}^{m \times n}$, where $B = \text{stack}(\{Q[i]\}_i^q)$ and each $Q[i]$ is a query group matrix. We assign an

⁶This discussion is also straightforward to extend to the more general case in which the workloads and strategy matrices are not required to be vertically stacked query group matrices. We focus on these specific workload and strategy matrices because they are similar to the ones we consider in the rest of the paper.

integer to each block geounit for the purpose of ordering the blocks. Specifically, suppose that the blocks are ordered lexicographically so that blocks in the same state are adjacent to one another, within each state, blocks in the same county are adjacent, etc. For example, in the case of the conventional geographic spine, this can be achieved by sorting the blocks by their 15 digit census GEOID, as the format of the GEOID is [2 digit state FIPS code][3 digit county FIPS code][6 digit census tract code][4 digit census block code].⁷ We refer to geounit $u \in \{1, \dots, U[l]\}$ in geolevel $l \in \{1, \dots, L\}$ as geounit (l, u) .

Now let $b[l, u]$ denote the number of block level descendants of geounit (l, u) . For example, since all census blocks are descendants of the root geounit, there is a total of $b[1, 1]$ blocks.⁸ Let $\mathbf{x} \in \mathbb{Z}_+^{nb[1,1]}$ be a vector of histogram cell counts over all blocks, ordered in the manner described in the preceding paragraph. Using this notation, we can express the query answers to the query matrix B for the root geounit in terms of these block-level cell counts as $(\mathbf{1}_{b[1,1]}^\top \otimes B)\mathbf{x}$. Likewise, the query answers for all block geounits is $(I_{b[1,1]} \otimes B)\mathbf{x}$. More generally, the query answers of geolevel l can be expressed in terms of the block-level cell counts by $(A[l] \otimes B)\mathbf{x}$, where $A[l] := \text{block_diag}(\{\mathbf{1}_{b[l,u]}^\top\}_{u=1}^{U[l]})$ and $\text{block_diag}(\{E[1], E[2], \dots\})$ denotes a block diagonal matrix with the (possibly non-square) matrices $E[1], E[2], \dots$ along the block diagonal. We will refer to $A := \text{stack}(\{A[l]\}_{l=1}^L)$ as the matrix representation of the spine throughout the paper because this matrix encodes the adjacency relationships between geounits on the spine, as described in the following example.

Example 1. In this example we will describe a matrix representation of a simple spine. Specifically, suppose there are three geolevels. For geolevel $l = 1$, suppose the root (or US) geounit has two children in geolevel $l = 2$. Suppose the first geounit of geolevel $l = 2$ has one child geounit in geolevel $l = L = 3$, and the second geounit in geolevel $l = 2$ has two children. In this case, $A[1], A[2], A[3]$, and the matrix representation of the spine, *i.e.*, the matrix A , are given by

$$\begin{aligned} A[1] &= \text{block_diag}(\{\mathbf{1}_{b[1,1]}^\top\}) = \mathbf{1}_3^\top = [1 \quad 1 \quad 1], \\ A[2] &= \text{block_diag}(\{\mathbf{1}_{b[2,1]}^\top, \mathbf{1}_{b[2,2]}^\top\}) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix}, \\ A[3] &= \text{block_diag}(\{\mathbf{1}_{b[3,1]}^\top, \mathbf{1}_{b[3,2]}^\top, \mathbf{1}_{b[3,3]}^\top\}) = I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \end{aligned}$$

and

$$A = \text{stack}(\{A[1], A[2], A[3]\}) = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

⁷Federal Information Processing Standard (FIPS) codes are fixed width codes assigned to states and counties. Each two digit state FIPS code is unique, and within each state, each three digit county FIPS code is also unique.

⁸For the Commonwealth of Puerto Rico, all blocks are descendant from the PR geounit.

□

This notation allows us to express the strategy matrix containing all linear queries in all geolevels for the case in which the per-geounit strategy matrix is the same in each geolevel; specifically, the overall strategy matrix in this case is given by

$$S := A \otimes B \in \{0, 1\}^{m \sum_l U[l] \times nb[1,1]},$$

where, as described above, the per-geounit strategy matrix B has dimension $m \times n$. As a summary of this notation, an example of $A \otimes B$ is shown in Figure 2.

The notation above is not quite general enough to capture all possible strategy matrices that a user may specify for a DAS execution in the most general case. This is because the DAS also allows users to specify distinct strategy matrices in each geolevel, but the notation above assumes that the per-geounit strategy matrix in each geolevel is B . For this reason, we use an alternative definition for the strategy matrix that is general enough to encompass all possible strategies used by the DAS. To do so, let $B[l] \in \{B[l]\}_{l=1}^L$ denote the per-geounit strategy matrix for geounits in geolevel l . Using this notation, the full strategy matrix is,

$$S := \begin{bmatrix} A[1] \otimes B[1] \\ A[2] \otimes B[2] \\ \vdots \\ A[L-1] \otimes B[L-1] \\ A[L] \otimes B[L] \end{bmatrix}. \quad (3.1)$$

In contrast to the strategy matrix, S , we do not require the workload to consist solely of linear queries for a hierarchical set of geographic entities; instead we only place very limited restrictions on this matrix. Specifically, we use W to denote the workload and only require that $W \in \mathbb{R}^{f \times nb[1,1]}$.

When either the conventional spine or an AIAN spine is used within the DAS, the spine optimization routines are not used and the strategy matrix used by the DAS is simply the unaltered initial strategy matrix. In these cases, we use $\beta_l \in \{\beta_l\}_{l=1}^L$, where $\beta_l \in \mathbb{R}_{++}$ and $\sum_l \beta_l = 1$, to denote the proportion of the global PLB that is used for the query groups of geolevel l . For each geolevel l , we use $\alpha_{i,l} \in \{\alpha_{i,l}\}_{i=1}^{q[l]}$, where $\alpha_{i,l} \in \mathbb{R}_{++}$ and $\sum_i \alpha_{i,l} = 1$, to denote the proportion of the geolevel PLB that is used for query group i . In summary, the proportion of the global PLB assigned to query group i for geounit (l, u) is $\beta_l \alpha_{i,l}$ in these cases.

The following theorem uses the notation above to show that the DAS is either ϵ -DP or ρ -zCDP, depending on whether a pure DP or zCDP framework and implementing mechanism is used, when either the conventional spine or an AIAN spine is used. The fact that geounits can be bypassed in the case of the optimized spine requires some modifications to this proof, so this result is provided in Section 5.

Theorem 1. Suppose the implementing mechanism $\mathcal{M}(\mathbf{x})$ outputs $\{(A[l] \otimes B[l])\mathbf{x} + \mathbf{y}_l\}_{l=1}^L$, where each $B[l]$ is defined by vertically stacking the query group matrices $\{Q[i, l]\}_{i=1}^{q[l]}$ with $Q[i, l]$ of dimension $m[i, l] \times n$, and \mathbf{y}_l is a vector of independent random variables. Also, let $\boldsymbol{\alpha}[l] := \text{stack}(\{\mathbf{1}_{m[i, l]} \alpha_{i, l}\}_i)$. Then we have the following.

(1) If either

$$\mathbf{y}_l \sim \text{Laplace}(\mathbf{0}, \text{diag}(\mathbf{2} \odot (\epsilon \beta_l \boldsymbol{\alpha}[l]))) \text{ or } \mathbf{y}_l \sim \text{Laplace}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{2} \odot (\epsilon \beta_l \boldsymbol{\alpha}[l])),$$

1	1	1	1	1	1	1	1	1
1	1	0	0	1	1	0	0	0
0	0	1	1	0	0	1	1	1
1	0	1	0	1	0	1	0	0
0	1	0	1	0	1	0	1	0
1	0	0	0	1	0	0	0	0
0	1	0	0	0	1	0	0	0
0	0	1	0	0	0	0	1	0
0	0	0	1	0	0	0	0	1
1	1	1	1	0	0	0	0	0
1	1	0	0	0	0	0	0	0
0	0	1	1	0	0	0	0	0
1	0	1	0	0	0	0	0	0
0	1	0	1	0	0	0	0	0
1	0	0	0	0	0	0	0	0
0	1	0	0	0	0	0	0	0
0	0	1	0	0	0	0	0	0
0	0	0	1	0	0	0	0	0
0	0	0	0	1	1	1	1	1
0	0	0	0	1	1	0	0	0
0	0	0	0	0	0	1	1	1
0	0	0	0	1	0	1	0	0
0	0	0	0	0	1	0	1	1
0	0	0	0	1	0	0	0	0
0	0	0	0	0	1	0	0	0
0	0	0	0	0	0	1	0	0
0	0	0	0	0	0	0	1	0
0	0	0	0	0	0	0	0	1

Figure 2: An example of $A \otimes B$ is given above. In this example, A is given by $\text{stack}(\mathbf{1}_2^\top, I_2)$, which corresponds to the case in which there is one US geounit with two block geolevel child geounits. Each geounit contains a 2×2 histogram, and the linear queries in B are given by a total sum query ($\mathbf{1}_2^\top \otimes \mathbf{1}_2^\top$), the marginal query groups for each attribute ($I_2 \otimes \mathbf{1}_2^\top$ and $\mathbf{1}_2^\top \otimes I_2$), and the detailed cell query group ($I_2 \otimes I_2$). For each of the geounits, the total sum query is highlighted in green, the marginal query groups are highlighted in light and dark red, and the detailed cell query group is unhighlighted.

then both $\mathcal{M}(\cdot)$ and the DAS are ϵ -DP with respect to the neighbor definition $\{x, x' \in \mathcal{X}^d \cap \mathcal{G} \mid d_{\mathcal{H}}(x, x') = 2\}$.

(2) If either

$$\mathbf{y}_l \sim \mathcal{N}(\mathbf{0}, \text{diag}(\mathbf{1} \otimes (\rho\beta_l\alpha[l]))) \text{ or } \mathbf{y}_l \sim \mathcal{N}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{1} \otimes (\rho\beta_l\alpha[l])),$$

then both $\mathcal{M}(\cdot)$ and the DAS are ρ -zCDP with respect to the neighbor definition $\{x, x' \in \mathcal{X}^d \cap \mathcal{G} \mid d_{\mathcal{H}}(x, x') = 2\}$.

Proof. Note that for each fixed l , we have

$$A[l] \otimes B[l] = A[l] \otimes \text{stack}(\{Q[i, l]\}_i) = \text{stack}(\{A[l] \otimes Q[i, l]\}_i).$$

Since each $A[l] \otimes Q[i, l]$ satisfies the definition of a query group matrix, Lemma 4 implies that releasing the output of \mathcal{M}_l , with output defined as $(A[l] \otimes B[l])\mathbf{x} + \mathbf{y}_l$ is $(\beta_l\epsilon)$ -DP in the first case and $(\beta_l\rho)$ -zCDP in the second case. Thus, sequential composition implies that the DP implementing mechanism $\mathcal{M}(\cdot)$ is $\sum_l \beta_l\epsilon = \epsilon$ -DP in the first case and $\sum_l \beta_l\rho = \rho$ -zCDP in the second case. Since both ρ -zCDP and ϵ -DP guarantees are invariant to post-processing, the DAS satisfies the same privacy guarantee as the implementing mechanisms. \square

In contrast to cases in which either the conventional spine or an AIAN spine are used, the spine optimization routines alter the initial matrix representation of the spine, A , and

the PLB proportions allocated to each geounit. We use $\gamma_{l,u}$ to denote the proportion of the global PLB allocated to geounit (l, u) after spine optimization, and γ to denote a vector composed of these values, *i.e.*, $\gamma := \text{stack}(\{\gamma_{l,u}\}_{l,u})$. In other words, after spine optimization is complete, the proportion of the global PLB that is allocated to query group i for geounit (l, u) is given by $\gamma_{l,u}\alpha_{i,l}$. Note that the spine optimization routines do not alter the per-geounit workloads, $\{B[l]\}_l$, or the per-geounit query group proportions, $\{\alpha_{i,l}\}_{i,l}$. Table 1 provides a summary of the notation introduced in this section that we use in the rest of the paper.

$U[l] \in \mathbb{N}$	The number of geounits in geolevel $l \in \{1, \dots, L\}$
$b[l, u] \in \{1, \dots, U[l]\}$	The number of block descendants of geounit (l, u)
$\alpha_{i,l} \in \mathbb{R}_{++}$	For each geounit in geolevel l , the proportion of the geounit's PLB allocated to query i , <i>i.e.</i> , $\sum_i \alpha_{i,l} = 1$ for every l
$\beta_l \in \mathbb{R}_{++}$	The proportion of the global PLB that is allocated to each geounit in geolevel l before spine optimization
$\gamma_{l,u} \in \mathbb{R}_+$	The global PLB proportion of geounit (l, u) after spine optimization
$B[l] \in \{0, 1\}^{m \times n}$	The per-geounit strategy matrix for geolevel l , defined by stacking the query group matrices for geolevel l
$A[l] \in \{0, 1\}^{U[l] \times b[1,1]}$	The matrix representation of geolevel l of the spine, <i>i.e.</i> , $A[l] = \text{block_diag}(\{\mathbf{1}_{b[l,u]}^\top\}_u)$
$A \in \{0, 1\}^{\sum_l U[l] \times b[1,1]}$	The matrix representation of the spine, <i>i.e.</i> , $A = \text{stack}(\{A[l]\}_{l=1}^L)$
$S \in \{0, 1\}^{m \times \sum_l U[l] \times nb[1,1]}$	The strategy matrix, as defined in (3.1)

Table 1: Definitions used to construct the strategy matrices used by the DAS

4. BRINGING OFF-SPINE ENTITIES CLOSER TO THE SPINE

This section describes two methods supported by the spine optimization routines to enhance accuracy in a set of target OSEs by bringing these geographic units closer to the spine in the first stage of the spine optimization routines. Both of these methods sacrifice accuracy in one or more geographic levels of the input spine to achieve this goal because they replace one or more geographic levels of the input geographic spine with alternative geographic levels that are closer to the OSEs. For example, the 2020 production redistricting and DHC spine settings resulted in the block group geographic level in an AIAN input spine being replaced with optimized block group geolevels, which degraded accuracy for census block groups.

First, the approach used by the 2020 production redistricting and DHC DAS implementations to bring target OSEs closer to the geographic spine is conceptually the simplest approach. These DAS implementations define this first stage of the spine optimization settings by directly defining geographic levels as the intersections of geographic entities. For example, each optimized block group in the optimized spine used for the 2020 redistricting data file production execution is defined as an intersection of a combination of OSEs and geounits, *e.g.*, a tract, an AIAN area (or the region outside of all AIAN areas), *etc.*; see Section 6 for a full description of this approach.

The next subsection describes the second alternative approach that is supported by the DAS but that is not used in our 2020 production DAS execution settings. Specifically, rather than directly defining the new geographic levels in the optimized spine as intersections of

OSEs, this alternative method automates the geonit definitions in the new geographic levels to bring OSEs closer to the optimized spine.

4.1. Automating the Choice of Alternative Geolevel Definitions. This section describes one approach to redefine geolevels to bring OSEs closer to the spine. This is done by using a heuristic to reduce an objective function based on the *off-spine entity distance* (OSED), which is the minimum number of geonits that must be added or subtracted from one another to derive an OSE.⁹ This heuristic generally results in a reduction in the variance of estimates for these geographic regions, particularly when a sufficiently high proportion of the global PLB is allocated to each geolevel.

Before defining a systematic algorithm that outputs the OSEDs for each OSE, we define some additional notation. Let the set of OSEs be denoted by \mathcal{K} , and let $\{C_k(u)\}_{k \in \mathcal{K}}$, where $C_k : \mathbb{N} \rightarrow \{0, 1\}$, denote a set of functions such that $C_k(u)$ is equal to one when the OSE $k \in \mathcal{K}$ contains the block geonit u and is equal to zero otherwise. We assume the input and output spines of this stage of the spine optimization routines are encoded using the matrix representation of the spine, as defined in the previous section and denoted by A . The algorithm we describe for computing OSED next requires computing intermediate OSEDs of intersections of a given geographic entity and geonits on the spine; we denote the OSED of the intersection of the geographic entity $k \in \mathcal{K}$ and the geonit (u, l) as $c[k, l, u]$.

We begin by fixing some $k \in \mathcal{K}$ and finding the OSEDs for both k and its complement, which we denote by k' , under the (temporary) assumption that the only geolevel is the block geolevel. In this case, each block would contribute $c[k, L, u] := C_k(u)$ to the OSED of OSE k . Likewise, each block would contribute $c[k', L, u] := 1 - C_k(u)$ to the OSED of the complement of k .

If the block-group geolevel were to be added to the spine at this point, we could compute the intersection of a single block group and k in one of two ways. First, we could add all the blocks together that are both inside of entity k and inside of the block group. This would result in the block-group u contributing $\sum_{v \in \text{Children}(u)} c[k, L, v]$ to the OSED of k . On the other hand, we could also take all the geographic extent that the block group occupies and then subtract off the geographic region of the blocks in the complement of k . This would result in this block group contributing $1 + \sum_{v \in \text{Children}(u)} c[k', L, v]$ to the OSED of k . Note that an additional one is added in this case because of the additional step of subtracting the complement of k in the block group from the block group itself. Since OSED is defined as the *minimum* number of geonits that must be added or subtracted to one another to define an entity, we choose the option that results in a smaller value. Since similar derivations can be carried out for the complement of k , by symmetry, we have,

⁹Our use of the term “off-spine entities” refers to geographic entities that may be off of the DAS geographic spine, but we do not assume that each OSE has a geographic extent that differs from each (on-spine) geonit. An implication of this definition of OSED is that, when the geographic extent of an OSE is identical to that of a geonit on the spine, its OSED is equal to one.

$$c[k, L - 1, u] := \min \left\{ \sum_{v \in \text{Children}(u)} c[k, L, v], 1 + \sum_{v \in \text{Children}(u)} c[k', L, v] \right\} \quad (4.1)$$

$$c[k', L - 1, u] := \min \left\{ \sum_{v \in \text{Children}(u)} c[k', L, v], 1 + \sum_{v \in \text{Children}(u)} c[k, L, v] \right\}. \quad (4.2)$$

Similar logic can be repeated to derive the following recursive system of equations for the OSED of the intersection of entity k and an arbitrary geounit $(l - 1, u)$

$$c[k, l - 1, u] := \min \left\{ \sum_{v \in \text{Children}(u)} c[k, l, v], 1 + \sum_{v \in \text{Children}(u)} c[k', l, v] \right\} \quad (4.3)$$

$$c[k', l - 1, u] := \min \left\{ \sum_{v \in \text{Children}(u)} c[k', l, v], 1 + \sum_{v \in \text{Children}(u)} c[k, l, v] \right\} \quad (4.4)$$

Since all entities are assumed to be contained within the US, the final OSED for entity k can be found by applying these recursions up to the root geounit and defining this OSED as $c[k, 1, 1]$. Afterward, our final objective function is defined by applying the reduce operation $h : \mathbb{N}^{\text{Card}(\mathcal{K})} \rightarrow \mathbb{R}$ to the OSEDs, which, for example, can be defined as the arithmetic mean or the max function. This computation is summarized in Algorithm 1.

Algorithm 1: OSEDs.Reduced($A, \mathcal{K}, \{C_k(\cdot)\}_k, h(\cdot)$)

```

for  $k \in \mathcal{K}$  do
  for geounit  $u$  in geolevel  $L$  do
     $c[k, L, u] \leftarrow C_k(u)$ 
     $c[k', L, u] \leftarrow 1 - C_k(u)$ 
  for geolevel  $l$  in  $\{L - 1, L - 2, \dots, 1\}$  do
    for geounit  $u$  in geolevel  $l$  do
       $x \leftarrow \sum_{v \in \text{Children}(u)} c[k, l + 1, v]$ 
       $y \leftarrow \sum_{v \in \text{Children}(u)} c[k', l + 1, v]$ 
       $c[k, l, u] \leftarrow \min \{x, y + 1\}$ 
       $c[k', l, u] \leftarrow \min \{y, x + 1\}$ 
  return  $h(\{c[k, 1, 1]\}_{k \in \mathcal{K}})$ 

```

In the most general setting, in which the OSEs are not necessarily disjoint, formulating an algorithm that redefines certain geolevels in order to minimize the OSEDs with a polynomial time complexity appears to be a difficult problem because of the similarity of this optimization problem to a set covering problem. For this reason, Algorithm 2 describes an example of a greedy approach to approximate the redefinition of tract groups that minimizes the OSEDs of a set of off-spine entities in a computationally tractable manner. This example first redefines block groups as optimized block groups by combining blocks within a given tract geounit, and within the same intersection of the OSEs. As described in the pseudocode below, each optimized block group is composed of up to $\sqrt{n} + \text{fanout_cutoff}$ blocks, where n is the number of blocks in the optimized block group's parent tract and fanout_cutoff is a user choice parameter. The goal of this choice is to ensure the fanout value, *i.e.*, the

number of child geounits of a given geounit, of the tract and the fanout values of its child optimized block groups are reasonably low. This is beneficial because high fanout values can increase the runtime of DAS executions (by increasing the size of the optimization problems used within the DAS), and we have found that high fanout values also appear to degrade accuracy. To see why using this upper bound for the number of blocks within each block group results in reasonably low fanout values, note that we can minimize the highest fanout value among a given tract and its children by defining the optimized block groups so that they each contain at most $\lceil \sqrt{n} \rceil$, where $\lceil \cdot \rceil$ denotes the ceiling function. For example, for a tract with 100 block geounit descendants, redefining its children by 10 optimized block group child geounits, each with 10 block child geounits, would result in the lowest possible maximum fanout of 10 among these 11 geounits. After the optimized block group geolevel has been defined in Algorithm 2, the tract group geolevel is redefined to reduce the OSEDs of the entities in \mathcal{K} . Note that this stage of the algorithm also ensures that the maximum number of tracts within each tract group is less than $\sqrt{n} + \text{fanout_cutoff}$.

Algorithm 2 uses $\mathbf{x} \preceq_{\text{lexicographic}} \mathbf{y}$, where $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, to denote the lexicographic less than or equal to partial ordering, which is defined as \perp when $\mathbf{x} \neq \mathbf{y}$ and the first index i for which $\mathbf{x}[i]$ and $\mathbf{y}[i]$ differ satisfies $\mathbf{x}[i] > \mathbf{y}[i]$, and \top otherwise.

Note that Algorithm 2 does not alter the PLB proportions of the input spine. The next section describes the algorithm used to update these proportions in the second spine optimization stage.

5. A PARETO FRONTIER OF GEOUNIT DEFINITIONS

This section uses a matrix mechanism to derive decision rules for whether or not to bypass a geounit. We will consider the case in which a pure-DP implementing mechanism is used within the DAS first. This decision rule is motivated by a setting that is less general than that of the DAS in several ways. First, for every $l, l' \in \{1, \dots, L\}$, we suppose $B[l] = B[l']$, and let $B := B[l]$. Second, for each query group i , we suppose $\alpha_{i,l} = \alpha_{i,l'}$, and let $\alpha_i := \alpha_{i,l}$. Third, we also constrain our attention to DP implementing mechanisms that use noise drawn from the continuous distributions $\text{Laplace}(\cdot)$ rather than its discrete counterpart $\text{Laplace}_{\mathbb{Z}}(0, b)$. The decision rule is based on the variance matrix of the WLS estimator of the detailed cell histogram counts for all block geounits in this simplified setting. It will be helpful to derive this variance prior to describing the decision rule. To do so, recall that the number of rows of query group i is denoted by $m[i, 1]$, and let $\boldsymbol{\alpha} := \text{stack}(\{\mathbf{1}_{m[i,1]} \alpha_i\}_i)$. In the terminology introduced in Section 2.4, let the rescaling vector $r \in \mathbb{R}_+^{m \sum_i U[i]}$ in this case be defined as $r := \boldsymbol{\gamma} \otimes \boldsymbol{\alpha} \epsilon$ and let the rescaled strategy matrix be defined as

$$\tilde{S} := \text{diag}(r)S = (\text{diag}(\boldsymbol{\gamma}) \otimes \text{diag}(\boldsymbol{\alpha} \epsilon))(A \otimes B) = (\text{diag}(\boldsymbol{\gamma})A) \otimes (\text{diag}(\boldsymbol{\alpha} \epsilon)B).$$

Also, let the stacked DP implementing mechanism answers be defined as $\tilde{\mathbf{z}} := \tilde{S}\mathbf{x} + \tilde{\mathbf{y}}$, where $\tilde{\mathbf{y}} = \text{diag}(r)\mathbf{y}$, $\mathbf{y} \sim \text{Laplace}(\mathbf{0}, \text{diag}(\mathbf{2} \circ (\epsilon \boldsymbol{\gamma} \otimes \boldsymbol{\alpha})))$.

Then the WLS estimate for this strategy matrix can be expressed as

$$\hat{\mathbf{x}} = (\tilde{S}^{\top} \tilde{S})^{-1} \tilde{S} \tilde{\mathbf{z}}, \quad (5.1)$$

and the variance matrix of the output of the matrix mechanism, $\text{Var}(W\hat{\mathbf{x}})$, is proportional to

Algorithm 2: `Redefine_Block_Groups_and_Tract_Groups($A, \mathcal{K}, \{C_k(\cdot)\}_k$)`

```

// Within each tract, redefine block groups by combining groups of
//  $\sqrt{n} + \text{fanout\_cutoff}$  blocks in the intersections of the same OSEs,
// where  $n$  is the number of blocks in the tract.
// Initialize tract groups so that they each have one child Tract.
current_OSEDs  $\leftarrow$  OSEDs_Reduced( $A, \mathcal{K}, \{C_k(\cdot)\}_k, \text{Sort\_Descending}(\cdot)$ )
for  $i \in \{1, \dots\}$  do
  altered_spine  $\leftarrow \perp$ 
  for county  $\in$  Counties do
     $n \leftarrow$  Number_of_tracts_in_county(county)
    for  $u, v \in$  Children(county) do
      if  $\text{Card}(\text{Children}(u) \cup \text{Children}(v)) > \sqrt{n} + \text{fanout\_cutoff}$  then
         $\perp$  continue
      // Combine_Siblings( $u, v, A$ ) returns a new spine defined as the
      // spine  $A$  after replacing the sibling geounits  $u$  and  $v$ 
      // with a single geounit with children given by
      // Children( $u$ )  $\cup$  Children( $v$ ):
       $A' \leftarrow$  Combine_Siblings( $u, v, A$ )
      test_OSEDs  $\leftarrow$  OSEDs_Reduced( $A', \mathcal{K}, \{C_k(\cdot)\}_k, \text{Sort\_Descending}(\cdot)$ )
      if test_OSEDs  $\preceq_{\text{lexicographic}}$  current_OSEDs then
         $A \leftarrow A'$ 
        current_OSEDs  $\leftarrow$  test_OSEDs
        altered_spine  $\leftarrow \top$ 
    if not altered_spine then
       $\perp$  break
return  $A$ 

```

$$W(\tilde{S}^\top \tilde{S})^{-1} W^\top. \quad (5.2)$$

It is also worth explicitly stating how we define the operation of bypassing a parent geounit. We define the operation of bypassing a parent geounit with c children, each with equal PLB proportions, by 1) creating c geounits in the geolevel of the parent, each with a geolevel proportion given by the sum of the proportions allocated to the parent and one of the children, 2) defining the single child of each of these c new geounits by one of the children, 3) removing the old parent geounit, 4) redefining the geolevel proportion of each child to be zero. Thus, even though we call this operation “bypassing a parent,” this operation actually moves the geolevel PLB proportion to a higher geolevel. This ensures that this operation does not change the total number of geolevels, and since the DAS fixes the final estimates in a top-down manner, so that the consistency with parent constraints are satisfied, this also ensures that the entire share of the geolevel PLB allocations are used in cases in which a parent geounit has only one child geounit. Although this definition describes the operation used within the DAS, the decision rules developed in this section only depend on properties of the matrix $\tilde{S}^\top \tilde{S}$, and using this definition of the bypass operation impacts this matrix in the same way as simply reallocating the PLB of the parent geounit to the child geounits. For

this reason, we motivate the decision rules described in this section on this simpler definition of the bypass operation. In other words, unlike what is done in the DAS codebase, here we define the operation of bypassing a parent geounit as redefining the PLB of each of the child geounits as the sum of the child's PLB and the PLB of the parent and then redefining the PLB of the parent geounit to be zero.

The next result describes a case in which each of the expected squared errors of the matrix mechanism, *i.e.*, the diagonal of $\text{Var}(W\hat{\mathbf{x}})$, can be decreased, or remain unchanged, by bypassing a parent. In cases in which the initial PLB proportion of a parent and its children are equal, this result implies that accuracy can be improved by bypassing the parent geounit when it has less than or equal to three child geounits. Note that the decision rule in the following theorem is only used in the spine optimization routines of the DAS when discrete or continuous ϵ -DP Laplace mechanisms are used for the implementing mechanisms. The decision rule used for cases in which either discrete or continuous ρ -zCDP Gaussian mechanisms are used for the implementing mechanisms within the DAS, *i.e.*, the decision rule used for the redistricting data file and DHC production executions, is described after the proof.

Theorem 2. Suppose that ϵ -DP is implemented using Laplace mechanisms and that the per-geolevel query strategies and query PLB proportions are the same in each geolevel. Also, suppose that the geolevel PLB allocated to each of the c children of geounit $(l, 1)$ are equal (*i.e.*: $\gamma_{l+1,u} = \gamma_{l+1,v}$ for all $u, v \in \{1, \dots, c\}$). If $\gamma_{l+1,1} \geq (c-1)\gamma_{l,1}/2$, then, for any B and $\{\alpha_i\}_i$, reallocating the PLB assigned to geounit $(l, 1)$ to its children will either decrease or leave unchanged each of the diagonal elements of $\text{Var}(W\hat{\mathbf{x}})$.

Proof. Let $\{\gamma_0[l, u]\}_{l,u}$ and $\{\gamma_1[l, u]\}_{l,u}$ denote the geounit PLB proportions before and after reallocation, respectively. Likewise, let $\hat{\mathbf{x}}_0$ and $\hat{\mathbf{x}}_1$ denote the WLS estimate before and after reallocation, respectively. Also, for the symmetric matrices $C, D \in \mathbb{R}^{n \times n}$ we will use $C \leq D$ to denote the condition that $C - D$ is negative semidefinite and $C \leq 0$ to denote the condition that C is negative semidefinite. The variance matrix of the WLS estimate before (respectively, after) bypassing geounit $(l, 1)$ is proportional to

$$\begin{aligned} \text{Var}(\hat{\mathbf{x}}_i) &= ((\text{diag}(\gamma_i)A \otimes \text{diag}(\boldsymbol{\alpha})B)^\top (\text{diag}(\gamma_i)A \otimes \text{diag}(\boldsymbol{\alpha})B))^{-1} \\ &= ((A^\top \text{diag}(\gamma_i) \otimes B^\top \text{diag}(\boldsymbol{\alpha}))(\text{diag}(\gamma_i)A \otimes \text{diag}(\boldsymbol{\alpha})B))^{-1} \\ &= ((A^\top \text{diag}(\gamma_i^2)A) \otimes (B^\top \text{diag}(\boldsymbol{\alpha}^2)B))^{-1} \\ &= (A^\top \text{diag}(\gamma_i^2)A)^{-1} \otimes (B^\top \text{diag}(\boldsymbol{\alpha}^2)B)^{-1}, \end{aligned}$$

where $i = 0$ (respectively, $i = 1$). We prove the sufficient condition that $\text{Var}(\hat{\mathbf{x}}_1) \leq \text{Var}(\hat{\mathbf{x}}_0)$. Given the variance matrix of the WLS estimator above, this condition is equivalent to

$$\begin{aligned} \text{Var}(\hat{\mathbf{x}}_1) \leq \text{Var}(\hat{\mathbf{x}}_0) &\iff \\ (A^\top \text{diag}(\gamma_1^2)A)^{-1} \otimes (B^\top \text{diag}(\boldsymbol{\alpha}^2)B)^{-1} &\leq (A^\top \text{diag}(\gamma_0^2)A)^{-1} \otimes (B^\top \text{diag}(\boldsymbol{\alpha}^2)B)^{-1} \iff \\ ((A^\top \text{diag}(\gamma_1^2)A)^{-1} - (A^\top \text{diag}(\gamma_0^2)A)^{-1}) &\otimes (B^\top \text{diag}(\boldsymbol{\alpha}^2)B)^{-1} \leq 0. \end{aligned}$$

Since we assume that $B^\top \text{diag}(\boldsymbol{\alpha}^2)B$ is positive definite, this condition holds if and only if,

$$\begin{aligned}
((A^\top \text{diag}(\gamma_1^2)A)^{-1} - (A^\top \text{diag}(\gamma_0^2)A)^{-1}) &\leq 0 \iff \\
(A^\top \text{diag}(\gamma_1^2)A)^{-1} &\leq (A^\top \text{diag}(\gamma_0^2)A)^{-1} \iff \\
A^\top \text{diag}(\gamma_1^2)A &\geq A^\top \text{diag}(\gamma_0^2)A \iff \\
A^\top \text{diag}(\gamma_1^2 - \gamma_0^2)A &\geq 0
\end{aligned}$$

Note that the only elements of γ_1 that are not equal to γ_0 can be defined in terms of γ_0 by $\gamma_1[l+1, u] = \gamma_0[l+1, u] + \gamma_0[l, 1]$ for all $u \in \{1, \dots, c\}$ and $\gamma_1[l, 1] = 0$. Thus, $A^\top \text{diag}(\gamma_0^2 - \gamma_1^2)A$ is a block matrix with a single block that is nonzero. Let $n[u] := b[l+1, u]$ for $u \in \{1, \dots, c\}$. This block is of dimension $b[l, 1] \times b[l, 1]$, and is given by

$$D := \gamma_0[l, 1] \begin{bmatrix} 2\gamma_0[l+1, 1]\mathbf{1}_{n[1]}\mathbf{1}_{n[1]}^\top & -\gamma_0[l, 1]\mathbf{1}_{n[1]}\mathbf{1}_{n[2]}^\top & -\gamma_0[l, 1]\mathbf{1}_{n[1]}\mathbf{1}_{n[3]}^\top & \cdots \\ -\gamma_0[l, 1]\mathbf{1}_{n[2]}\mathbf{1}_{n[1]}^\top & 2\gamma_0[l+1, 2]\mathbf{1}_{n[2]}\mathbf{1}_{n[2]}^\top & -\gamma_0[l, 1]\mathbf{1}_{n[2]}\mathbf{1}_{n[3]}^\top & \\ -\gamma_0[l, 1]\mathbf{1}_{n[3]}\mathbf{1}_{n[1]}^\top & -\gamma_0[l, 1]\mathbf{1}_{n[3]}\mathbf{1}_{n[2]}^\top & 2\gamma_0[l+1, 3]\mathbf{1}_{n[3]}\mathbf{1}_{n[3]}^\top & \\ \vdots & & & \ddots \end{bmatrix}.$$

For each $i \in \{1, \dots, c\}$, let $a[i] := \sum_{j=1}^{i-1} n[j]$, $b[i] := \sum_{j=i+1}^c n[j]$, and $\mathbf{t}_i := (\mathbf{0}_{a[i]}^\top, \mathbf{1}_{n[i]}^\top/n[i], \mathbf{0}_{b[i]}^\top)^\top$. Note that the matrix D has only c unique columns, and the vectors $\{\mathbf{t}_i\}_{i=1}^c$ provide an orthogonal basis for the span of these columns. Let $T := \text{horizontal_stack}(\{\mathbf{t}_i\}_i)$. Thus,

$$\begin{aligned}
D \geq 0 &\iff TDT^\top \geq 0 \iff \\
\begin{bmatrix} 2\gamma_0[l+1, 1]\gamma_0[l, 1] & -\gamma_0[l, 1]^2 & -\gamma_0[l, 1]^2 & \cdots \\ -\gamma_0[l, 1]^2 & 2\gamma_0[l+1, 2]\gamma_0[l, 1] & -\gamma_0[l, 1]^2 & \\ -\gamma_0[l, 1]^2 & -\gamma_0[l, 1]^2 & 2\gamma_0[l+1, 3]\gamma_0[l, 1] & \\ \vdots & & & \ddots \end{bmatrix} &\geq 0 \iff \\
D' := \text{diag}(\mathbf{1}_r(2\gamma_0[l+1, 1]\gamma_0[l, 1] + \gamma_0[l, 1]^2)) - \mathbf{1}_r\mathbf{1}_r^\top \gamma_0[l, 1]^2 &\geq 0.
\end{aligned}$$

Since $\gamma_0[l, u] \geq 0$ for all l and u , the eigenvector of D' corresponding to the smallest eigenvalue is $\mathbf{1}_r$, so this matrix is positive semidefinite when

$$2\gamma_0[l+1, 1]\gamma_0[l, 1] \geq (c-1)\gamma_0[l, 1]^2.$$

Since $(l, 1)$ was not bypassed in our initial PLB allocation, $\gamma_0[l, 1] \neq 0$, so we have

$$\gamma_0[l+1, 1] \geq (c-1)\gamma_0[l, 1]/2.$$

□

Remark 1. Consider two cases in which the only query in B is a total population query. First, if the parent is not bypassed, we could construct unbiased estimates of the total population of the parent by either observing the DP answer for the parent directly, which has a variance of $2/\gamma_{i,1}^2$, or by summing the DP answers of the children together, which has a variance of $2c/\gamma_{i+1,1}^2$. The mean with inverse-variance weighting provides the linear combination of these estimates with the lowest possible variance of $2c/(c\gamma_{i,1}^2 + \gamma_{i+1,1}^2)$ in this case. Second, if we did bypass the parent, we could estimate the total population query for the parent by summing the total population of the children, which would have a variance of $2c/(\gamma_{i+1,1} + \gamma_{i,1})^2$ in this case. The theorem above can be viewed as a statement that limiting our attention to a single total population query, and to only the parent and its children, in this way is without loss of generality, at least for the purpose of finding the cases in which bypassing the parent does not increase the expected error of the WLS estimator for any query in any geounit. This is because $2c/(c\gamma_{i,1}^2 + \gamma_{i+1,1}^2) \geq 2c/(\gamma_{i+1,1} + \gamma_{i,1})^2$ if and only if $\gamma_{i+1,1} \geq (c-1)\gamma_{i,1}/2$, which is the same requirement given in the statement of the theorem. \square

Remark 2. This theorem may be of independent interest in the DP literature because other authors have already considered strategy matrices that have a hierarchical structure that is analogous to the matrix representation of the spine A above, and this result can be used to narrow the search space of the set of hierarchical strategy matrices considered when choosing a strategy matrix with this property [10, 13]. For example, when using the WLS approach described by [10] with PLB proportions that are the same for each level of the hierarchy and with the children of each non-leaf node defined by either two or three sub-intervals of the range of the parent node, this result implies that the expected squared error of an arbitrary linear query can be reduced by increasing the number of sub-intervals used to define these child nodes. \square

One can use a similar technique to show that, in the case in which ρ -zCDP implementing Gaussian mechanisms are used to define \mathbf{y} , bypassing a parent will increase at least one diagonal element of $\text{Var}(W\hat{\mathbf{x}})$ whenever the parent has two or more children, and the diagonal elements of $\text{Var}(W\hat{\mathbf{x}})$ will remain unchanged whenever the parent has only one child. In part for this reason, our decision rule for the case in which ρ -zCDP discrete or continuous Gaussian mechanisms are used within the DAS consists of only bypassing parent geounits with only one child. As described in the introduction, this decision rule can also be motivated more directly than what was done in Theorem 2. Specifically, given the top-down manner in which the DAS fixes estimates to ensure consistency with the estimates of the parent geounits, the estimates of a child geounit are fixed by those of the parent geounit whenever the parent only has one child geounit. Thus, if we did not bypass parent geounits with only one child, the PLB allocation of these child geounits would simply not be used, in the sense that the noisy measurements of these child geounits would not impact the histogram estimates of the child geounits, as is also described in the following observation. Note that the decision rule provided in Theorem 2 results in at least the same number of geounits being bypassed, since, whenever $c = 1$, the decision rule provided in Theorem 2 results in bypassing being chosen for any PLB allocations of the parent and child geounits.

Observation 1. The noisy answers for each child geounit that does not have a sibling geounit do not impact the output of the DAS. Thus, the variance of the noisy answers that

are used to construct the histogram estimates of these child geounits can be decreased by bypassing all parent geounits with only one child.

Algorithm 3 summarizes how both decision rules are used within the spine optimization routines of the DAS. Note that this algorithm starts at the block group geolevel and iterates up the spine to the US geolevel, rather than starting at the US geolevel and moving downward, which ensures that no remaining geounits can be bypassed after only one pass through the spine.

Algorithm 3: Move_Spine_to_Pareto_Frontier($A, \{\beta_l\}_l, \text{Pure_DP} \in \{\top, \perp\}$)

```

for  $l \in \{1, \dots, L\}$  do
  for  $u$  in geolevel  $l$  do
     $\gamma_{l,u} \leftarrow \beta_l$ 
for  $l \in \{L-1, L-2, \dots, 1\}$  do
  for  $u$  in geolevel  $l$  do
    Min_Child_PLB  $\leftarrow \min_{c \in \text{Children}(u)} \gamma_{l+1,c}$ 
     $c \leftarrow \text{Card}(\text{Children}(u))$ 
    if Pure_DP and Min_Child_PLB  $\geq (c-1)\gamma_{l,u}/2$  then
       $A, \{\gamma_{l,u}\}_{l,u} \leftarrow \text{Bypass\_Parent\_Geounit}(u, A, \{\gamma_{l,u}\}_{l,u})$ 
    else if  $c = 1$  then
       $A, \{\gamma_{l,u}\}_{l,u} \leftarrow \text{Bypass\_Parent\_Geounit}(u, A, \{\gamma_{l,u}\}_{l,u})$ 
return  $A, \{\gamma_{l,u}\}_{l,u}$ 

```

5.1. The impact of Spine Optimization on the Privacy Guarantees of the DAS.

Using the DAS with the spine that is output from Algorithm 2 does not alter the privacy guarantees of the DAS because the same arguments used in the proof of Theorem 1 apply to this case as well. However, the same cannot be said for the spine that is output from Algorithm 3 when at least one geounit is bypassed. The following theorem generalizes the argument used in Theorem 1 to show that using the spine output from Algorithm 3 within the DAS does not impact the privacy guarantees.

Theorem 3. Suppose the implementing mechanism $\mathcal{M}(x)$ outputs $\{(A[l][u, \cdot] \otimes B[l])x + \mathbf{y}_{l,u}\}_{l,u}$, where each $A[l]$ is the matrix representation of geolevel $l \in \{1, \dots, L\}$ of the spine output from the spine optimization routines described in Sections 4 and 5, each $B[l]$ is defined by vertically stacking the query group matrices $\{Q[i, l]\}_{i=1}^{q[l]}$ with $Q[i, l]$ of dimension $m[i, l] \times n$, x is the confidential histogram cell counts for dataset $x \in \mathcal{X}^d \cap \mathcal{G}$, and $\mathbf{y}_{l,u}$ is a vector of either independent random variables or has all elements equal to ∞ . Also, let $\alpha[l] := \text{stack}(\{\mathbf{1}_{m[i,l]} \alpha_{i,l}\}_i)$. Then we have the following.

- (1) If $\mathbf{y}_{l,u} = \infty$ when $\gamma_{l,u} = 0$, and is distributed as either $\mathbf{y}_{l,u} \sim \text{Laplace}(\mathbf{0}, \text{diag}(\mathbf{2} \otimes (\epsilon\gamma_{l,u}\alpha[l])))$ or $\mathbf{y}_{l,u} \sim \text{Laplace}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{2} \otimes (\epsilon\gamma_{l,u}\alpha[l])))$ when $\gamma_{l,u} > 0$, then both $\mathcal{M}(\cdot)$ and the DAS are ϵ -DP with respect to the neighbor definition $\{x, x' \in \mathcal{X}^d \cap \mathcal{G} \mid d_{\mathcal{H}}(x, x') = 2\}$.
- (2) If $\mathbf{y}_{l,u} = \infty$ when $\gamma_{l,u} = 0$, and is distributed as either $\mathbf{y}_{l,u} \sim \text{N}(\mathbf{0}, \text{diag}(\mathbf{1} \otimes (\rho\gamma_{l,u}\alpha[l])))$ or $\mathbf{y}_{l,u} \sim \text{N}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\mathbf{1} \otimes (\rho\gamma_{l,u}\alpha[l])))$ when $\gamma_{l,u} > 0$, then both $\mathcal{M}(\cdot)$ and the DAS are ρ -zCDP with respect to the neighbor definition $\{x, x' \in \mathcal{X}^d \cap \mathcal{G} \mid d_{\mathcal{H}}(x, x') = 2\}$.

Proof. For $x \in \mathcal{X}^d \cap \mathcal{G}$, let the set of datasets $x' \in \mathcal{X}^d \cap \mathcal{G}$ that differ from x on a single entry be denoted by $N(x)$. Also, let $T[l] := \{u \in \mathbb{N} \mid \gamma_{l,u} > 0\}$. The output of all the DP implementing mechanisms is observationally equivalent to the set of finite output elements, or $\{(A[l][u, \cdot] \otimes B[l])\mathbf{x} + \mathbf{y}_{l,u}\}_{l,u \in T[l]}$, because, independently of the data, the remaining DP mechanism answers are infinite with probability one. Thus, without loss of generality, we can restrict our attention to an alternative mechanism that outputs $\{(A[l][u, \cdot] \otimes B[l])\mathbf{x} + \mathbf{y}_{l,u}\}_{l,u \in T[l]}$. Let $\hat{A}[l] := \text{stack}(\{A[l][u, \cdot]\}_{u \in T[l]})$, so the strategy matrix for this alternative mechanism can be written as,

$$D := \text{stack}(\{\hat{A}[l] \otimes B[l]\}_l).$$

To prove the first result of the theorem, let $\hat{\gamma}[l] := \text{stack}(\{\gamma_{l,u}\}_{l,u \in T[l]})$, and $\hat{\gamma} := \text{stack}(\{\hat{\gamma}[l]\}_l)$. In this case, the noise for this mechanism is distributed as either $\hat{\mathbf{y}} \sim \text{Laplace}(\mathbf{0}, \text{diag}(\hat{\mathbf{b}}))$ or $\hat{\mathbf{y}} \sim \text{Laplace}_{\mathbb{Z}}(\mathbf{0}, \text{diag}(\hat{\mathbf{b}}))$, where $\hat{\mathbf{b}} := \text{stack}(\{\hat{\mathbf{b}}[l]\}_l)$ and $\hat{\mathbf{b}}[l] := \mathbf{2} \otimes (\epsilon(\hat{\gamma}[l] \otimes \boldsymbol{\alpha}[l]))$. Thus, we will show the following condition, which, by the first result in Lemma 3, will imply the first result of the theorem,

$$\max_{x \in \mathcal{X}^d \cap \mathcal{G}, x' \in N(x)} \sum_i \left| D[i, \cdot](\mathbf{x} - \mathbf{x}') / \hat{\mathbf{b}}[i] \right| \leq \epsilon. \quad (5.3)$$

Starting from the left hand side of this inequality, we have,

$$\begin{aligned} \max_{x \in \mathcal{X}^d \cap \mathcal{G}, x' \in N(x)} \sum_i \left| \left(\text{diag}(\mathbf{1} \otimes \hat{\mathbf{b}}) D \right) [i, \cdot](\mathbf{x} - \mathbf{x}') \right| &\leq 2 \max_j \sum_i \left(\text{diag}(\mathbf{1} \otimes \hat{\mathbf{b}}) D \right) [i, j] = \\ &2 \max_j \sum_i \text{stack} \left(\left\{ \text{diag}(\epsilon(\hat{\gamma}[l] \otimes \boldsymbol{\alpha}[l])/2) (\hat{A}[l] \otimes B[l]) \right\}_l \right) [i, j] = \\ &\epsilon \max_j \sum_l \sum_i \left(\text{diag}(\hat{\gamma}[l]) \hat{A}[l] \otimes \text{diag}(\boldsymbol{\alpha}[l]) B[l] \right) [i, j]. \end{aligned} \quad (5.4)$$

Note that each of the row sums of $\text{diag}(\boldsymbol{\alpha}[l]) B[l]$, *i.e.*, the vector $B[l]^\top \text{diag}(\boldsymbol{\alpha}[l]) \mathbf{1} = B[l]^\top \boldsymbol{\alpha}[l]$, is equal to $\mathbf{1}$ because $\sum_k Q[i, l][k, \cdot] = \mathbf{1}^\top$ and $\sum_k \alpha_{k,l} = 1$. Thus, since each column j of $\text{diag}(\hat{\gamma}[l]) \hat{A}[l]$ contains at most one nonzero value, we have,

$$\begin{aligned} \epsilon \max_j \sum_l \sum_i \left(\left(\text{diag}(\hat{\gamma}[l]) \hat{A}[l] \right) \otimes \left(\text{diag}(\boldsymbol{\alpha}[l]) B[l] \right) \right) [i, j] &= \\ \epsilon \max_j \sum_l \sum_i \left(\text{diag}(\hat{\gamma}[l]) \hat{A}[l] \right) [i, j] &= \epsilon \max_j \sum_i \left(\text{diag}(\hat{\gamma}) \hat{A} \right) [i, j], \end{aligned}$$

Note that $\sum_i \left(\text{diag}(\hat{\gamma}) \hat{A} \right) [i, j]$ is equal to the sum of the geolevel PLB proportions along a path from the US geounit to the block geounit j . Since the bypass operation used within Algorithm 3 does not change this sum, and its initial value is $\sum_l \beta_l = 1$, we have

$$\epsilon \max_j \sum_i \left(\text{diag}(\hat{\gamma}) \hat{A} \right) [i, j] = \epsilon \sum_l \beta_l = \epsilon,$$

which implies inequality (5.3).

Similar logic also implies the second result of the theorem. Specifically, in this case we can use the second result of Lemma 3, and the fact that for each $x \in \mathcal{X}^d \cap \mathcal{G}$ and $x' \in \mathcal{N}(x)$ we have $(D[i, \cdot](\mathbf{x} - \mathbf{x}'))^2 = |D[i, \cdot](\mathbf{x} - \mathbf{x}')|$, to derive the sufficient condition,

$$\max_{x \in \mathcal{X}^d, x' \in \mathcal{N}(x)} \sum_i |D[i, \cdot](\mathbf{x} - \mathbf{x}')| / \hat{\sigma}_i^2 \leq 2\rho, \quad (5.5)$$

where $\hat{\sigma}^2 := \text{stack}(\{\hat{\sigma}^2[\mathbf{l}]\}_l)$ and $\hat{\sigma}^2[\mathbf{l}] := \mathbf{1} \odot (\rho(\hat{\gamma}[\mathbf{l}] \otimes \alpha[\mathbf{l}]))$. This inequality follows from,

$$\begin{aligned} \max_{x \in \mathcal{X}^d \cap \mathcal{G}, x' \in \mathcal{N}(x)} \sum_i |D[i, \cdot](\mathbf{x} - \mathbf{x}')| / \hat{\sigma}_i^2 &\leq 2 \max_j \sum_i (\text{diag}(\mathbf{1} \odot \hat{\sigma}^2) D) [i, j] = \\ &2\rho \max_j \sum_i (\text{diag}(\hat{\gamma}[\mathbf{l}] \otimes \alpha[\mathbf{l}]) D) [i, j]. \end{aligned} \quad (5.6)$$

Since this last value is equal to the value of (5.4) multiplied by $2\rho/\epsilon$, and the logic above implies the value of (5.4) is equal to ϵ , we have,

$$2\rho \max_j \sum_i (\text{diag}(\hat{\gamma}[\mathbf{l}] \otimes \alpha[\mathbf{l}]) D) [i, j] = \epsilon \cdot (2\rho/\epsilon) = 2\rho,$$

which implies the sufficient condition (5.5). \square

6. SPINE SETTINGS USED FOR 2020 CENSUS PRODUCTION EXECUTIONS

In this section we will describe the geographic spines used in the 2020 Census data product production executions, which include files for the persons and units universes for both the Redistricting Data (P.L. 94-171) Summary File and the Demographics and Housing Characteristics File (DHC) data products.

For the redistricting data product persons production execution, the geolevels included in the spine were US, state, county, tract, optimized block group, and block.¹⁰ The same approach described above was used to include both an AIAN and a non-AIAN branch to this spine at the state geolevel and below. Likewise, optimized block groups were defined using the approach described above using the following four categories of OSEs. First, each AIAN OSE is composed of an individual AIAN area, along with one additional OSE defined by the region outside of all AIAN areas. Second, each GQ OSE is defined as the union of all blocks that contain the same combination of major GQ types. For example, one GQ OSE is the union of all blocks that only contain GQs that are college/university student housing. This OSE category was included to decrease the impact that blocks with GQs have on their neighbors. Third, each minor civil division (MCD) OSE is composed of an MCD in the twelve strong-MCD states, *i.e.*, Connecticut, Maine, Massachusetts, Michigan, Minnesota, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, Vermont, and Wisconsin, along with one additional OSE defined by the region outside of all of these MCDs. Fourth, each place OSE is defined by an incorporated or census-designated place that are outside of the twelve strong MCD states, along with one additional OSE that is

¹⁰While we focus our attention on the settings for US executions, similar settings were used for the Puerto Rico executions, with the exception that the PR root is at the same level as state in the US hierarchy. The PLB allocated to each geolevel in the Puerto Rico executions was also normalized to ensure the global PLB of each Puerto Rico execution was the same as the global PLB of each corresponding US execution.

defined by the region outside of these places. As described above in more detail, optimized block groups were defined by grouping together blocks that are within the same AIAN, GQ, MCD, and place OSEs as well as the same census tract. After all geounits on the spine are defined in this way, the bypassing approach described in Algorithm 3 was used to define the final PLB values for each geounit.

The spine used in the redistricting data product housing unit production execution was nearly identical to the spine used in the redistricting data product persons production execution with the exception that the GQ OSE category was not considered when defining optimized block groups because occupied GQs are excluded from the universe of the housing units.

The 2020 production DHCP and DHCH executions both use identical spines. The geolevels included on this spine are US, state, county, prim, tract subset group, tract subset, optimized block group, and block. The US, state, county, and block geolevels are defined in the same way as the redistricting data product executions, so next we will define the remaining geolevels. First, prim geounits are defined as the most granular geography unit used internally by the Population Estimates and Projections (PEP) Area of the Census Bureau. All tabulations published by PEP can be derived by aggregation using prims as the aggregation atom.¹¹ For example, PEP produces estimates for four places in Autauga County, AL, *i.e.*, Autaugaville town, Billingsley town, Millbrook city, and Prattville city, so the intersection of each of these places and Autauga County defines one prim geounit that is included in this geolevel. In addition, since estimates for each county in the US are also published by PEP, the final prim within Autauga County is defined as the area within Autauga County that is outside of these first four prims. Second, each tract subset geounit was defined as the intersection of a census tract, a prim geounit, and an AIAN OSE, as defined above. Third, tract subset groups were defined by grouping together tract subsets using the same approach described above for grouping blocks to define optimized block groups.

Fourth, optimized block groups were defined in a similar way as in the redistricting spine, but with a different choice of categories of OSEs, which we will define next. Specifically, each school district (SD) OSE is composed of an individual SD, along with one additional OSE defined by the region outside of all SDs. Also, each conventional block group (CBG) OSE is composed of a block group on the standard census geographic spine, along with one additional OSE defined by the region outside of all CBGs. Afterward, optimized block groups were defined by grouping together blocks within the same GQ OSE, tract subset, SD OSE, and CBG OSE. Like the case of the redistricting spines, after all geounits were defined, the bypassing approach described in Algorithm 3 was used to define the final PLB values for each geounit.

7. IMPACT OF SPINE OPTIMIZATION: SUMMARY METRICS

To demonstrate the impact of spine optimization, this section provides summary metrics of the total population query for three redistricting person-level DAS executions. The 2010 CEF is used as the input data file for all three DAS executions. The execution labeled “Optimized” uses settings that are identical to those of the 2020 production redistricting

¹¹These geounits are most commonly called *primitives* in PEP, rather than *prim*s. Since the spine used within the DAS also includes more granular(/primitive) geounits than prims, we use prims instead in this paper to avoid confusion.

person-level DAS execution. The execution labeled “Conventional” uses the production settings with the exception that the internal spine is defined by the official 2010 tabulation geography. The execution labeled “AIAN” uses the production settings with the exception that the internal spine is defined by the version used to produce the official 2020 redistricting data.

The tables report the mean absolute error (MAE) of the total population query for various geographic levels and for each of these three DAS executions. Each such MAE value is defined by first finding the total population error of each geounit in the geolevel; afterward, the MAE is defined as the arithmetic mean of the absolute value of these errors. While we focus on the accuracy of the total population query in this section, we have found that the relative comparisons of MAEs between spine settings typically do not strongly depend on the query being considered, since the PLB allocated to the total population query in a given geounit is proportional to the PLB allocated to any alternative query group.

Table 2 provides total population MAE values for geolevels below the county in the conventional spine. Since we include a state geolevel total population invariant, the total population MAE for the state and US geolevel would be zero, so we only include the MAE for geolevels below the state geolevel. Regarding the column containing the count of the geounits in each geolevel, note that the universe of geounits that we consider only includes the geounits from the conventional spine containing at least one housing unit and/or at least one occupied GQ, which lowers the values of these counts in some cases. Table 3 provides total population MAE values for OSEs, including the OSEs that were targeted in the spine optimization routines, as described in Section 6 in more detail. The row “Aggregated AIAN Areas in States” in this table provides the total population MAE values, averaged over the geographic entities defined as the union of all AIAN blocks within each state. Note that the number of geographic entities is given by 36 because there are 36 states that contain AIAN blocks. The next two rows, “Aggregated AIAN Areas in Counties” and “Aggregated AIAN Areas in Tracts” are defined similarly.

Table 2: Total Population query MAE values of US redistricting persons DAS execution for geographic units on the standard census spine.

Geolevel	Geounit Count	Total Population MAE		
		Spine Type		
		Conventional	AIAN	Optimized
County	3,143	1.829	1.943	1.867
Tract	72,544	1.986	1.989	1.949
Block Group	216,886	1.407	1.409	15.99
Block	6,398,202	5.006	5.007	4.849

The accuracy metrics in these tables demonstrate several impacts of the DAS internal spine on accuracy that are worth highlighting. For example, Table 2 provides evidence that, relative to the execution that used the conventional spine for the DAS internal spine, the DAS execution that used the AIAN spine exhibited very slightly lower accuracy for the geounits on the conventional spine. In addition to random variability between DAS executions, this change is likely caused by the OSEDs of the geounits in the conventional spine increasing when moving from a DAS internal spine defined as the conventional spine to the AIAN spine. However, this impact is not large because most of the OSEDs of the geounits in the conventional spine are not impacted, and, since each geounit on the conventional spine can

Table 3: Total Population query MAE values of a US redistricting persons DAS execution for off-spine entities.

Geolevel	Geounit Count	Total Population MAE		
		Spine Type		
		Conventional	AIAN	Optimized
AIAN Areas	621	29.19	9.741	1.876
All Aggregated AIAN Areas	1	6428	15.00	3.000
Aggregated AIAN Areas in States	36	259.4	0.806	0.639
Aggregated AIAN Areas in Counties	412	29.09	1.629	1.391
Aggregated AIAN Areas in Tracts	1,532	11.08	1.793	1.656
AIAN Block	165,647	4.491	4.465	4.256
Place [†]	29,250	35.36	35.26	2.679
MCD [‡]	11,914	14.74	14.78	8.254

The universe of AIAN considered includes all AIAN areas other than state and tribal designated statistical areas.

[†]Place refers to both census designated and incorporated places outside of the 12 strong-MCD states.

[‡]MCD refers to all MCDs in each of the 12 strong-MCD states.

be defined by combining at most two geounits on the AIAN spine, the OSEDs that are impacted only increase from one to two. In contrast to these very small changes, Table 3 shows that moving from the conventional spine to the AIAN spine resulted in a large improvement in the accuracy within AIAN areas, since these areas are much closer to the AIAN spine than the conventional spine.

Several MAE values in Table 2 for the optimized spine are also noteworthy. First, since using the optimized spine moves conventional (tabulation) block groups further from the spine, the block group geounits in the conventional spine become much less accurate. Second, in contrast, the remaining geolevels improve in accuracy relative to the AIAN spine DAS execution, and, in most cases, also relative to the conventional spine DAS execution. These improvements are a result of using the bypassing step of the spine optimization routines described in Section 5 and also of reducing the largest fanout values, which, as described above, is defined as the number of child geounits of a parent geounit, at the tract geolevel and below. Likewise, all of the MAEs for the optimized spine DAS execution in Table 3 are lower than the corresponding values of the AIAN spine and the conventional spine DAS executions, which is also a result of using the bypassing step, of decreasing the largest fanout values at the tract geolevel and below, and of bringing these target OSEs closer to the spine. Since block group geounits in the conventional spine are not legally or politically defined, do not correspond to functioning governmental units, and are typically used in cases in which a non-specific geolevel is required with a granularity between that of census tracts and census blocks, the Census Bureau’s redistricting tuning experiments pointed to the accuracy improvements provided by spine optimization outweighing the cost of decreased accuracy in conventional (tabulation) block groups.

REFERENCES

- [1] John Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajhala, Brett Moran, William Sexton,

- Matthew Spence, and Pavel Zhuravlev. The 2020 Census Disclosure Avoidance System TopDown Algorithm. *Harvard Data Science Review*, (Special Issue 2), 2022. <https://doi.org/10.1162/99608f92.529e3cb9>.
- [2] Randall Akee and Norman DeWeaver. *2020 Census Data Products: Data Needs and Privacy Considerations*, chapter Identification of Rural and Special Populations: American Indians and Alaska Natives. The National Academies Press, Washington, DC, 2020. <https://doi.org/10.17226/25978>. URL: <https://nap.nationalacademies.org/catalog/25978/2020-census-data-products-data-needs-and-privacy-considerations-proceedings>, <https://doi.org/10.17226/25978>.
 - [3] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016. https://doi.org/10.1007/978-3-662-53641-4_24.
 - [4] Clement Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. *Journal of Privacy and Confidentiality*, 12(1), 2022. <https://doi.org/10.29012/jpc.784>.
 - [5] Ryan Cumings-Menon, Robert Ashmead, Daniel Kifer, Philip Leclerc, Matthew Spence, Pavel Zhuravlev, and John M Abowd. Disclosure Avoidance for the 2020 Census Demographic and Housing Characteristics File. *arXiv preprint arXiv:2312.10863*, 2023. <https://arxiv.org/abs/2312.10863>.
 - [6] Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B*, 84(1):3–37, 2022. <https://doi.org/10.1111/rssb.12454>.
 - [7] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006. https://doi.org/10.1007/11761679_29.
 - [8] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006. https://doi.org/10.1007/11681878_14.
 - [9] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. *SIAM Journal on Computing*, 41(6):1673–1693, 2012. <https://doi.org/10.1137/09076828x>.
 - [10] Michael Hay, Vibhor Rastogi, Gerome Miklau, and Dan Suciu. Boosting the accuracy of differentially private histograms through consistency. *Proceedings of the VLDB Endowment*, 3(1), 2010. <https://doi.org/10.14778/1920841.1920970>.
 - [11] Daniel Kifer, John M Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 census. *arXiv preprint arXiv:2209.03310*, 2022. <https://arxiv.org/abs/2209.03310>.
 - [12] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 193–204, 2011. <https://doi.org/10.1145/1989323.1989345>.
 - [13] Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau, and Andrew McGregor. Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 123–134, 2010. <https://doi.org/10.1145/1807085.1807104>.
 - [14] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. Optimizing error of high-dimensional statistical queries under differential privacy. *Proceedings of the VLDB Endowment*, 11(10):1206–1219, 2018. <https://doi.org/10.14778/3231751.3231769>.
 - [15] Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010. <https://doi.org/10.1198/jasa.2009.tm08651>.