The Saddle-Point Accountant for Differential Privacy

Shahab Asoodeh

Joint work with Wael Alghamdi, Felipe Gomez, Flavio Calmon (Harvard University),
Oliver Kosut, Lalitha Sankar (ASU)

October 19, 2022

McMaster University

زن_زندگی_آزادی

#Women-Life-Freedom

Joint work with

Wael Alghamdi
(Harvard)

Flavio Calmon
(Harvard)

Felipe Gomez
(Harvard)

Oliver Kosut
(ASU)

Lalitha Sankar
(ASU)

State-of-the-art composition

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

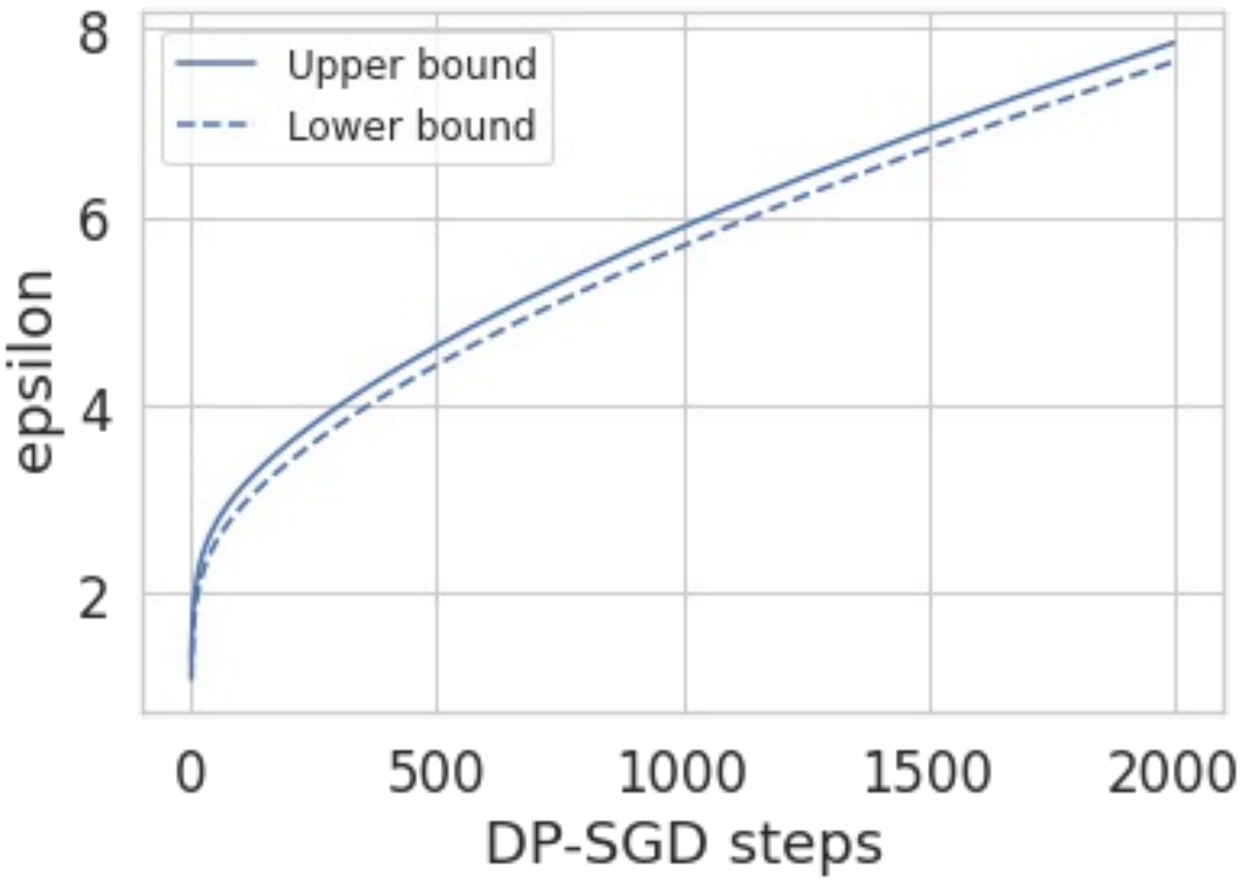

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

State-of-the-art composition

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

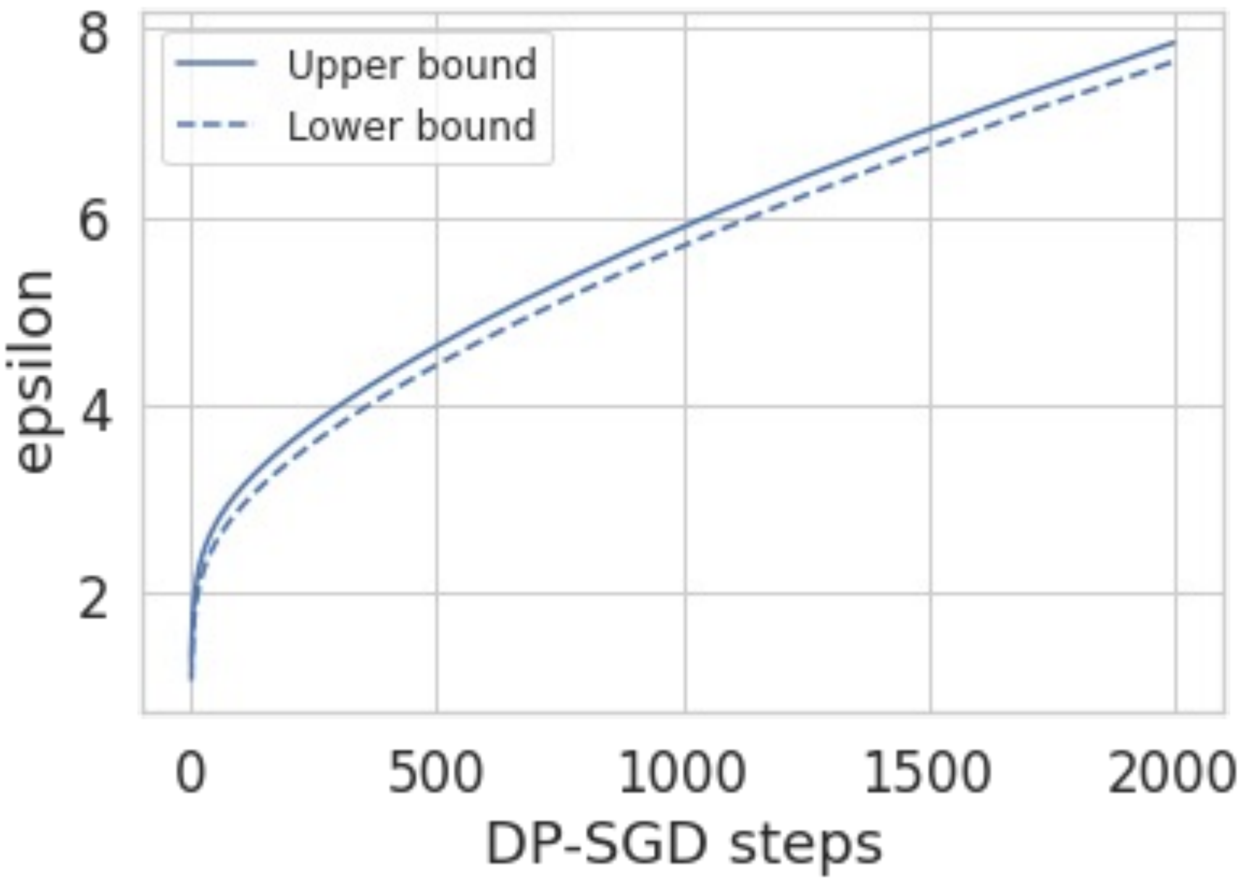

runtime complexity

$O(\sqrt{n} \log n)$

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

State-of-the-art composition

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

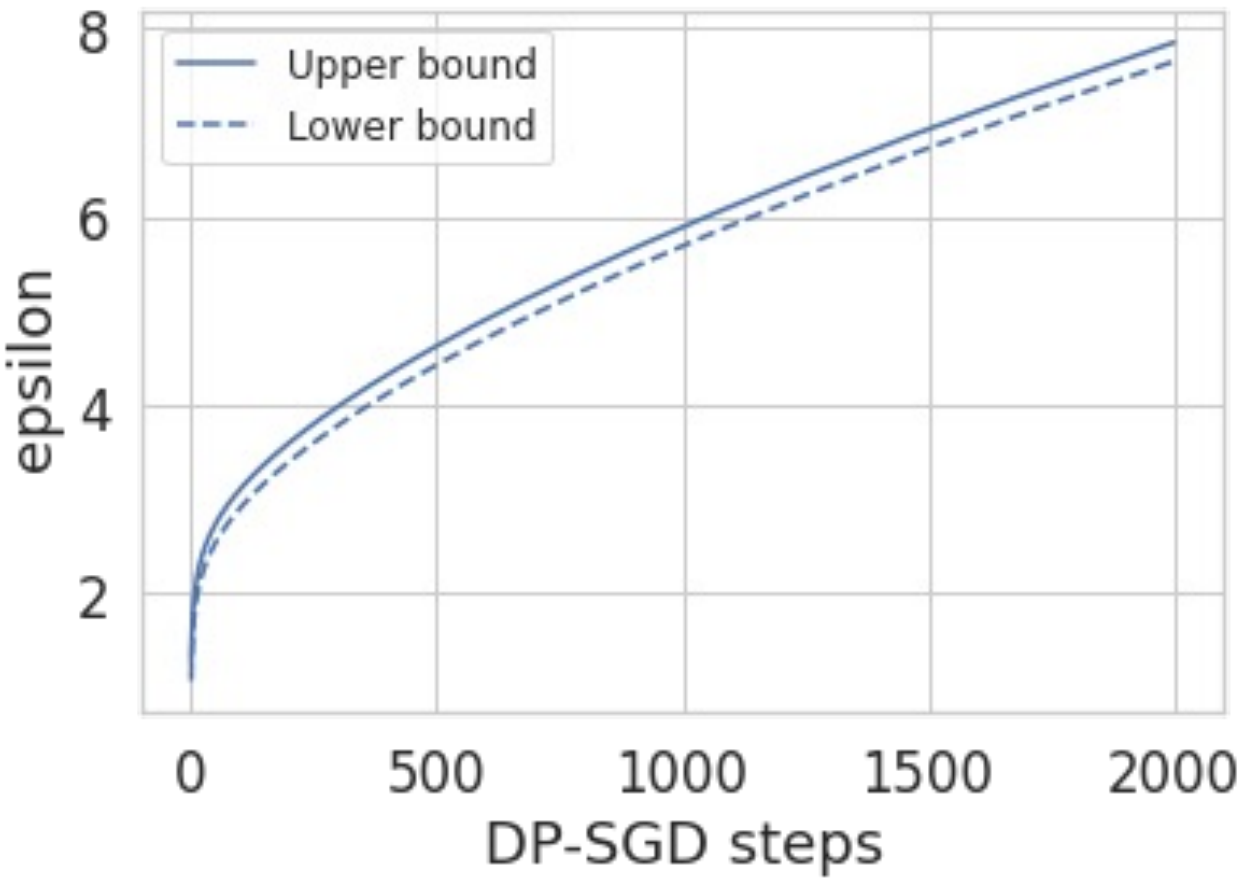

runtime complexity

$O(\sqrt{n} \log n)$

↓

$O(\text{polylog}(n))$

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

Ghazi, Kamath, Kumar, and Manurangsi, Faster Privacy Accounting via Evolving Discretization, ICML 2022

State-of-the-art composition

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$n = 2000$

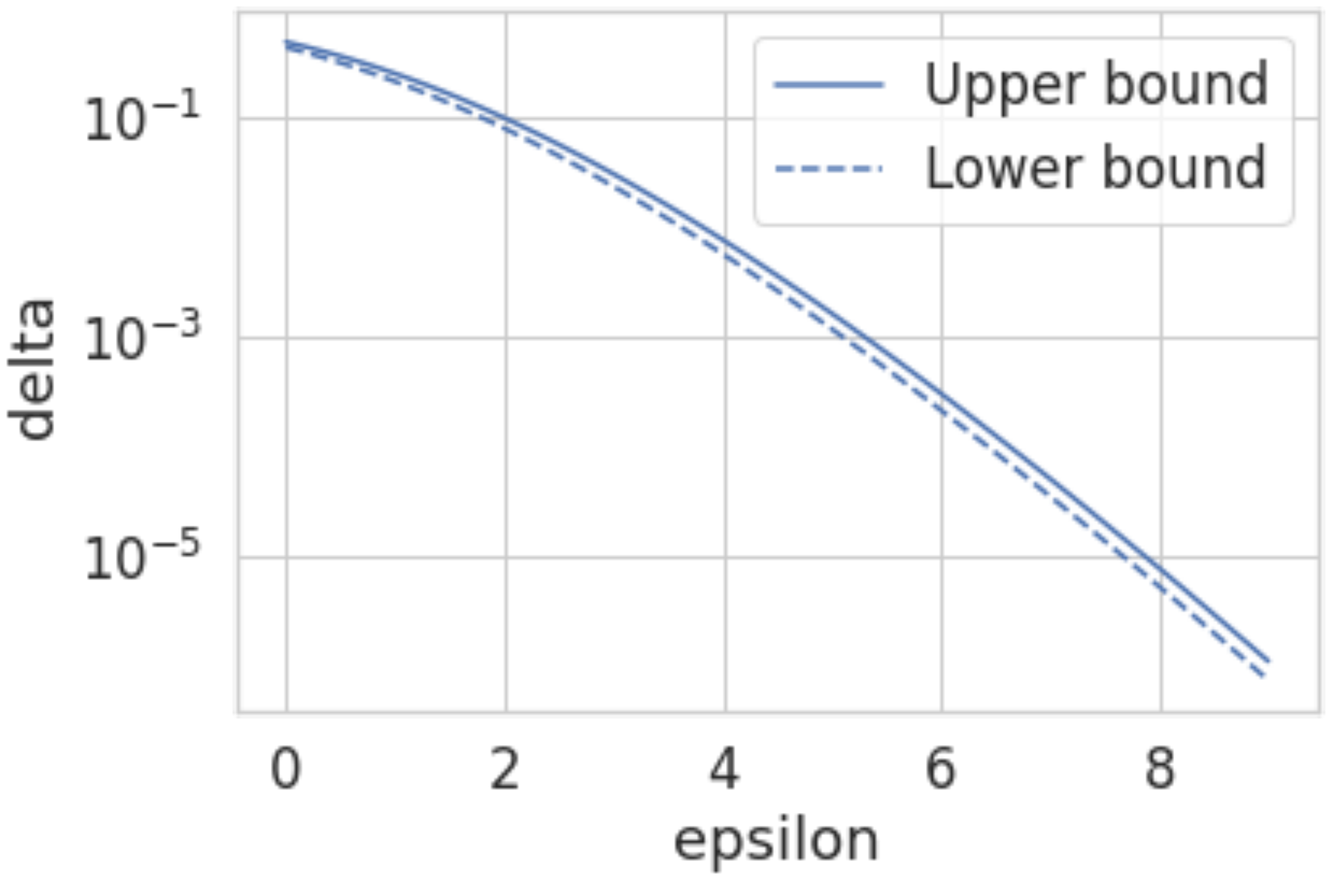

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

State-of-the-art composition

DP-SGD:

$\sigma = 1$

subsampling $= 10^{-2}$

$n = 2000$

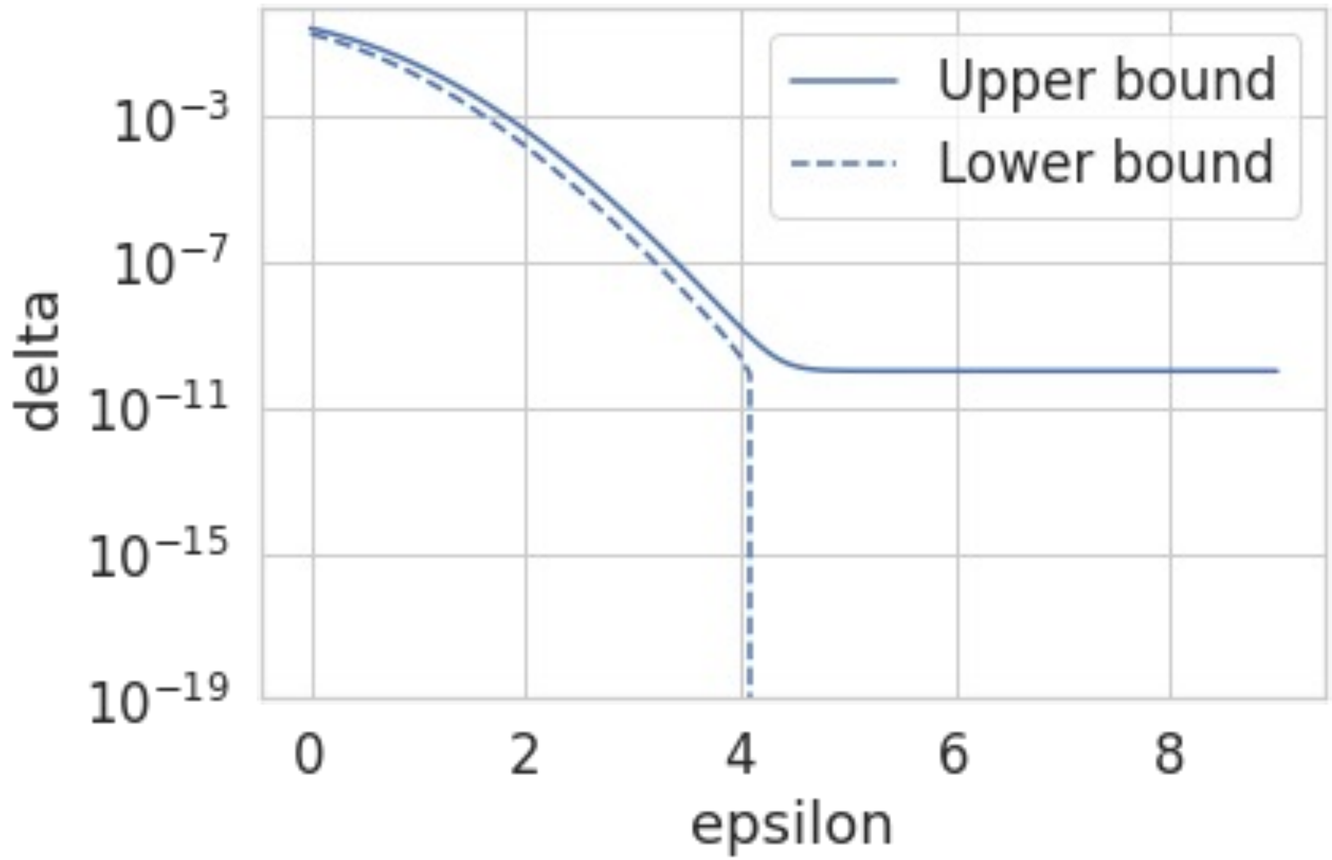

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

State-of-the-art composition

DP-SGD:

$\sigma = 1$

subsampling $= 10^{-2}$

$n = 2000$

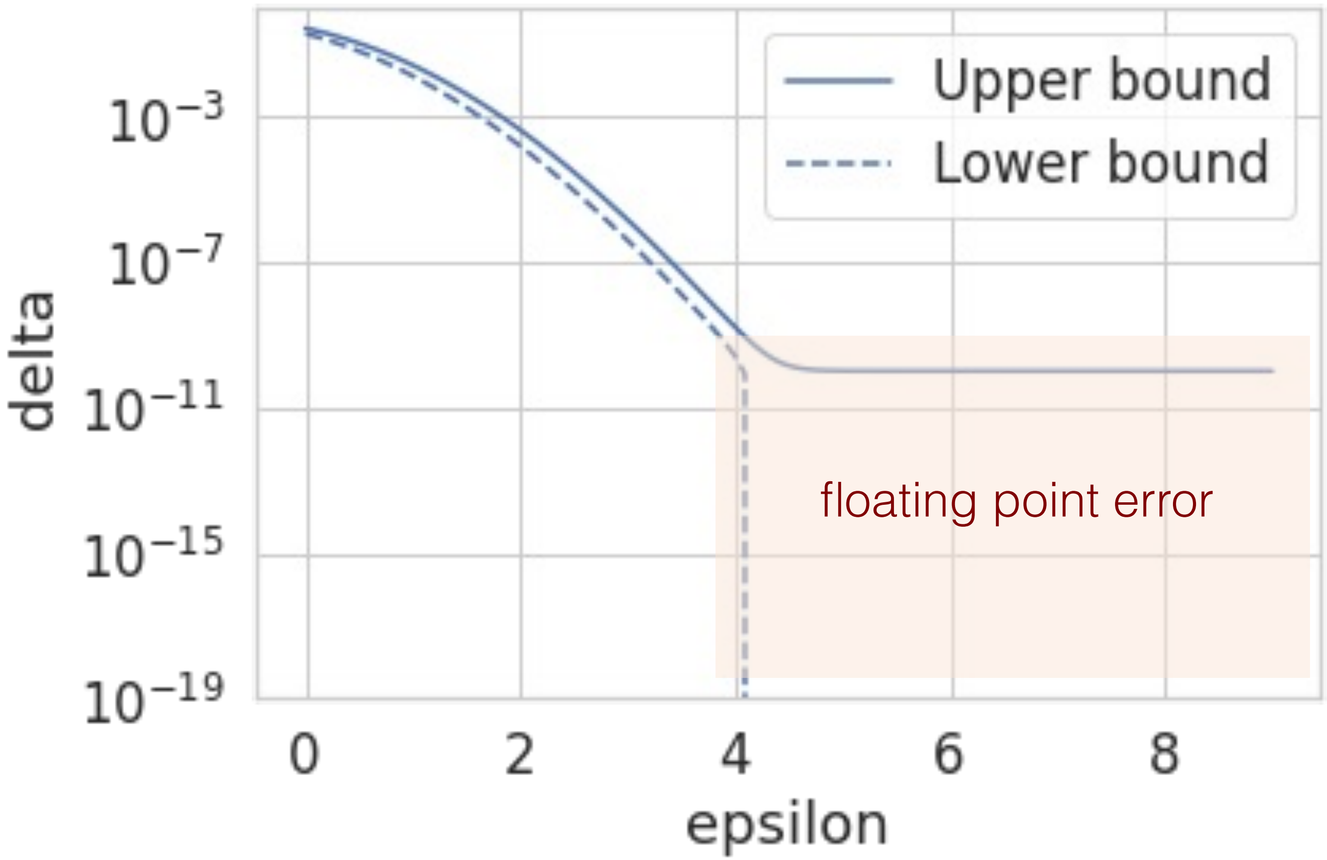

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

Today's talk: Develop DP numerical composition using saddle-point approximation:

- Runtime complexity independent of # composition
- Works for all epsilon and delta

Saddle-point vs. state-of-the-art

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

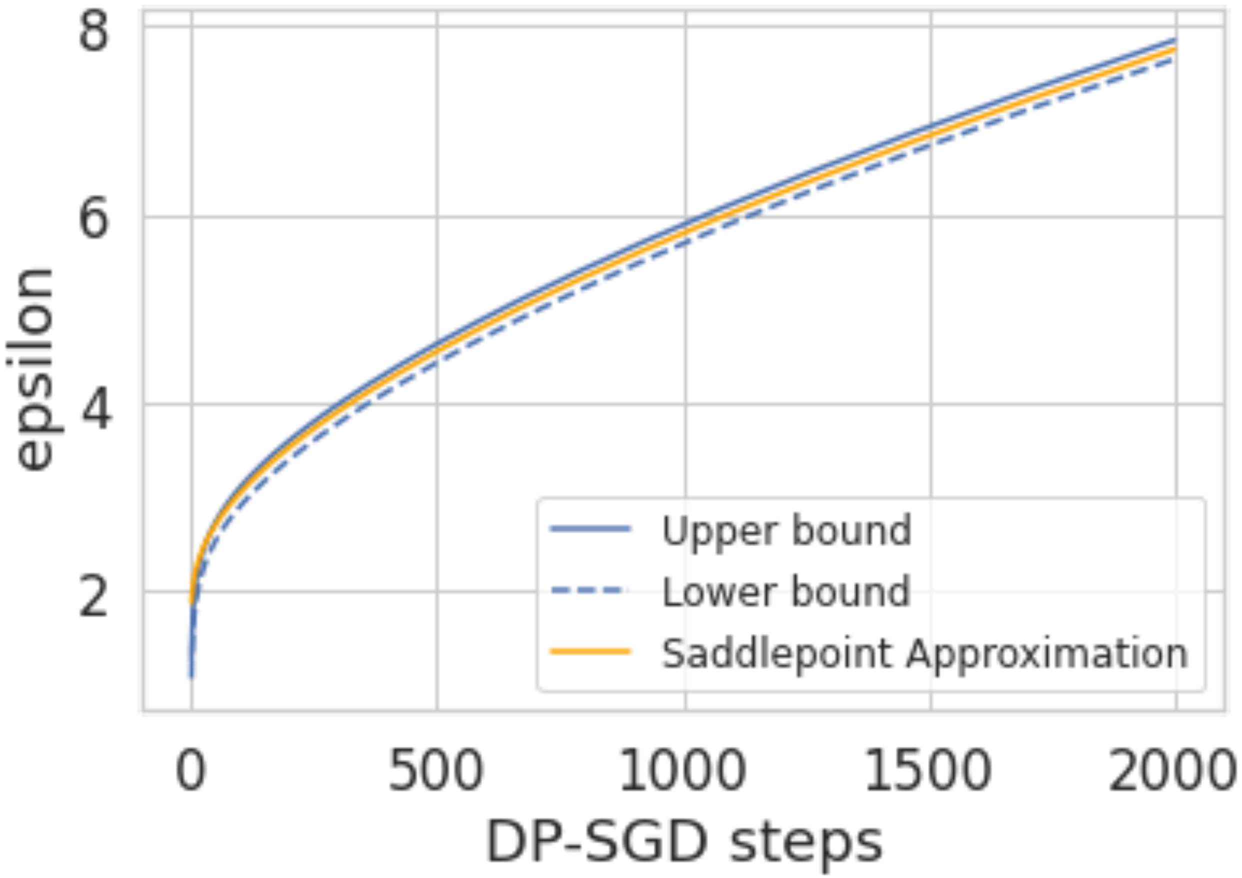

runtime complexity
$O(1)$

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

Saddle-point vs. state-of-the-art

DP-SGD:

$\sigma = 1$

subsampling $= 10^{-2}$

$n = 2000$

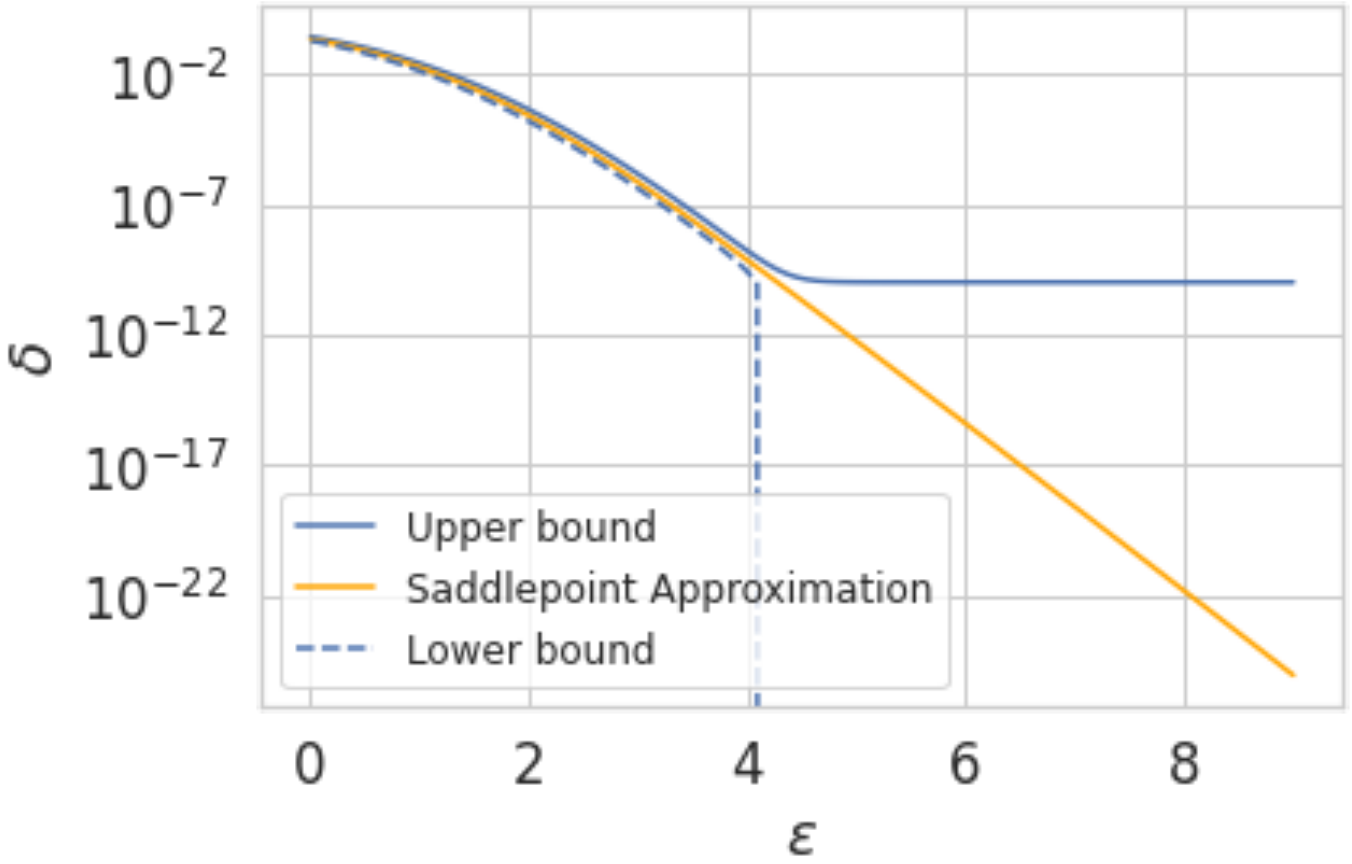

Gopi, Lee, and Wutschitz, Numerical Composition of Differential Privacy, NeurIPS 2021

Differential privacy

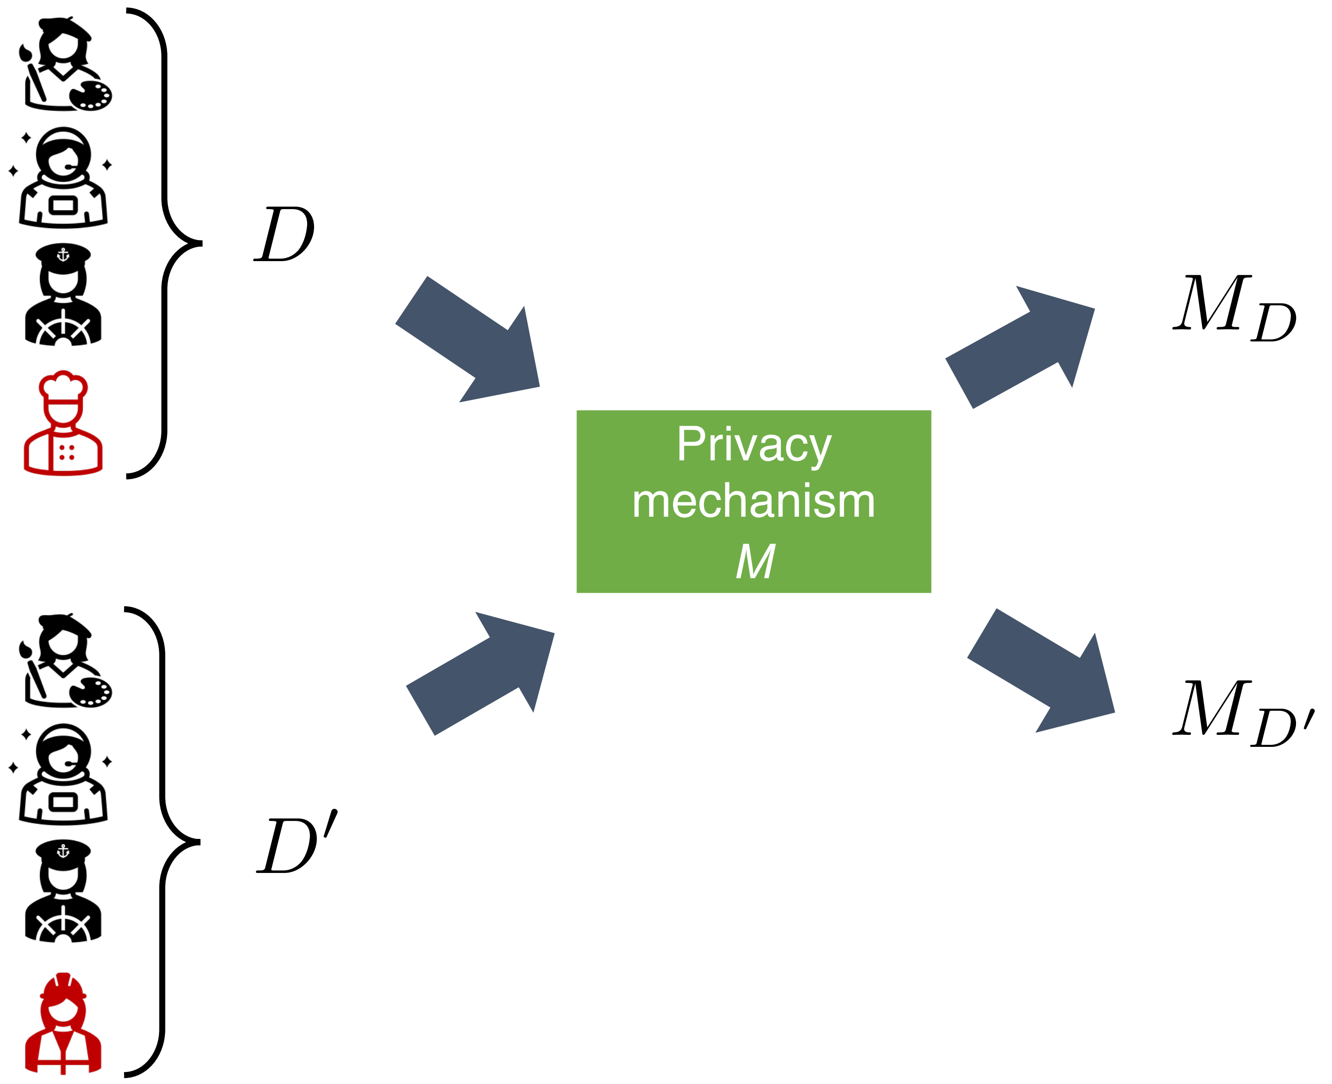

M is (ε, δ)-DP if $\forall\ D \sim D'$

$$\sup_{\text{subset } A} [M_D(A) - e^{\varepsilon} M_{D'}(A)] \leq \delta$$

Hockey-stick divergence
$\mathsf{E}_{\varepsilon}(M_D \| M_{D'})$

M is (ε, δ)-DP if $\forall\ D \sim D'$

$$\mathsf{E}_{\varepsilon}(M_D \| M_{D'}) \leq \delta$$

Dominating distribution, PLRV

A pair (P, Q) is said to dominate M if

$$\sup_{D \sim D'} \mathsf{E}_\varepsilon(M_D \| M_{D'}) \leq \mathsf{E}_\varepsilon(P \| Q)$$

or tightly dominate M if equality is achieved for all ε.

$\delta(\varepsilon) \triangleq$ smallest δ such that M is (ε, δ)-DP

privacy curve

If (P, Q) tightly dominates M, then

$$\delta(\varepsilon) = \mathsf{E}_\varepsilon(P \| Q) = \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)_+\right]$$

where

$$L = \log \frac{\mathrm{d}P}{\mathrm{d}Q}(X) \qquad \text{with } X \sim P$$

privacy loss random variable

Dominating distribution, PLRV

A pair (P, Q) is said to dominate M if

$$\sup_{D \sim D'} \mathsf{E}_\varepsilon(M_D \| M_{D'}) \leq \mathsf{E}_\varepsilon(P \| Q)$$

or tightly dominate M if equality is achieved for all ε.

$\delta(\varepsilon) \triangleq$ smallest δ such that M is (ε, δ)-DP

privacy curve

If (P, Q) ~~tightly~~ dominates M, then

$$\delta(\varepsilon) \overset{\leq}{\cancel{=}} \mathsf{E}_\varepsilon(P \| Q) = \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)_+\right]$$

where

$$L = \log \frac{\mathrm{d}P}{\mathrm{d}Q}(X) \qquad \text{with } X \sim P$$

privacy loss random variable

DP-SGD

Algorithm 1 Differentially private SGD (Outline)

Input: Examples $\{x_1, \ldots, x_N\}$, loss function $\mathcal{L}(\theta) = \frac{1}{N}\sum_i \mathcal{L}(\theta, x_i)$. Parameters: learning rate η_t, noise scale σ, group size L, gradient norm bound C.

Initialize θ_0 randomly

for $t \in [T]$ **do**

Take a random sample L_t with sampling probability L/N

Compute gradient

For each $i \in L_t$, compute $\mathbf{g}_t(x_i) \leftarrow \nabla_{\theta_t}\mathcal{L}(\theta_t, x_i)$

Clip gradient

$\bar{\mathbf{g}}_t(x_i) \leftarrow \mathbf{g}_t(x_i) / \max\left(1, \frac{\|\mathbf{g}_t(x_i)\|_2}{C}\right)$

Add noise

$\tilde{\mathbf{g}}_t \leftarrow \frac{1}{L}\left(\sum_i \bar{\mathbf{g}}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})\right)$

Descent

$\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{\mathbf{g}}_t$

Output θ_T and compute the overall privacy cost (ε, δ) using a privacy accounting method.

DP-SGD

Algorithm 1 Differentially private SGD (Outline)

Input: Examples $\{x_1, \ldots, x_N\}$, loss function $\mathcal{L}(\theta) = \frac{1}{N}\sum_i \mathcal{L}(\theta, x_i)$. Parameters: learning rate η_t, noise scale σ, group size L, gradient norm bound C.

Initialize θ_0 randomly

for $t \in [T]$ **do**

Take a random sample L_t with sampling probability L/N

Compute gradient

For each $i \in L_t$, compute $\mathbf{g}_t(x_i) \leftarrow \nabla_{\theta_t}\mathcal{L}(\theta_t, x_i)$

Clip gradient

$\bar{\mathbf{g}}_t(x_i) \leftarrow \mathbf{g}_t(x_i) / \max\left(1, \frac{\|\mathbf{g}_t(x_i)\|_2}{C}\right)$

Add noise

$\tilde{\mathbf{g}}_t \leftarrow \frac{1}{L}\left(\sum_i \bar{\mathbf{g}}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})\right)$

Descent

$\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{\mathbf{g}}_t$

Output θ_T and compute the overall privacy cost (ε, δ) using a privacy accounting method.

Tightly dominating distributions for each iteration:

$$P = p\mathcal{N}(0, \sigma^2 C^2) + (1-p)\mathcal{N}(C, \sigma^2 C^2) \qquad Q = \mathcal{N}(0, \sigma^2 C^2)$$

DP-SGD

Algorithm 1 Differentially private SGD (Outline)

Input: Examples $\{x_1, \ldots, x_N\}$, loss function $\mathcal{L}(\theta) = \frac{1}{N}\sum_i \mathcal{L}(\theta, x_i)$. Parameters: learning rate η_t, noise scale σ, group size L, gradient norm bound C.

Initialize θ_0 randomly

for $t \in [T]$ **do**

Take a random sample L_t with sampling probability L/N

Compute gradient

For each $i \in L_t$, compute $\mathbf{g}_t(x_i) \leftarrow \nabla_{\theta_t}\mathcal{L}(\theta_t, x_i)$

Clip gradient

$\bar{\mathbf{g}}_t(x_i) \leftarrow \mathbf{g}_t(x_i) / \max\left(1, \frac{\|\mathbf{g}_t(x_i)\|_2}{C}\right)$

Add noise

$\tilde{\mathbf{g}}_t \leftarrow \frac{1}{L}\left(\sum_i \bar{\mathbf{g}}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})\right)$

Descent

$\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{\mathbf{g}}_t$

Output θ_T and compute the overall privacy cost (ε, δ) using a privacy accounting method.

Tightly dominating distributions for each iteration:

$$P = p\mathcal{N}(0, \sigma^2 C^2) + (1-p)\mathcal{N}(C, \sigma^2 C^2) \qquad Q = \mathcal{N}(0, \sigma^2 C^2)$$

$$\Downarrow$$

$$\delta(\varepsilon) = \mathsf{E}_\varepsilon(P \| Q) = \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)_+\right]$$

$$L = \log\left(1 - p + p \cdot e^{\frac{C(2X - C)}{2\sigma^2}}\right), \qquad X \sim P$$

Composition of DP

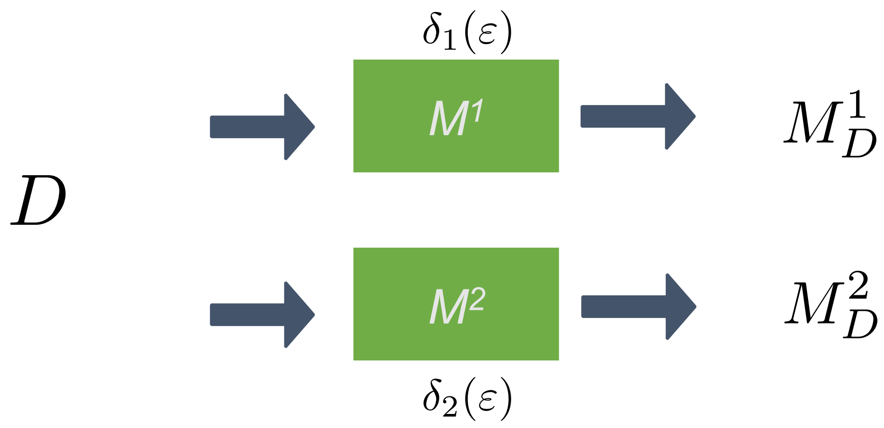

Composition of DP

$M = M^1 \circ M^2$

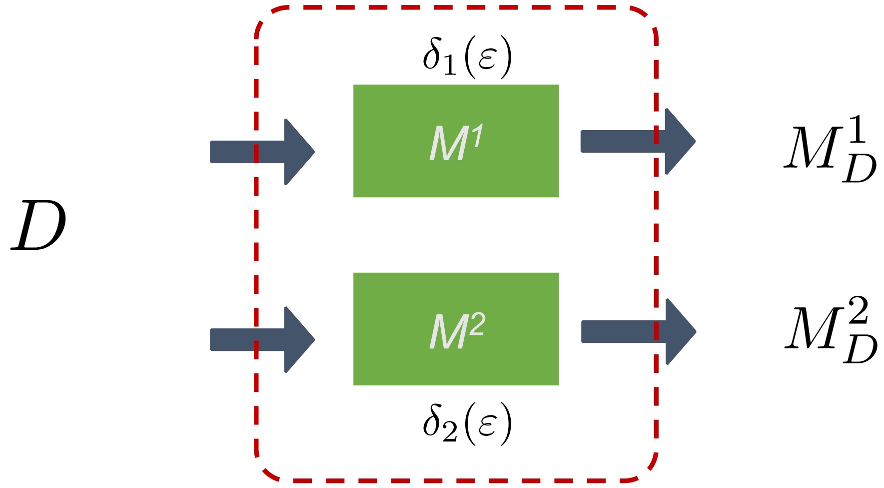

Composition of DP

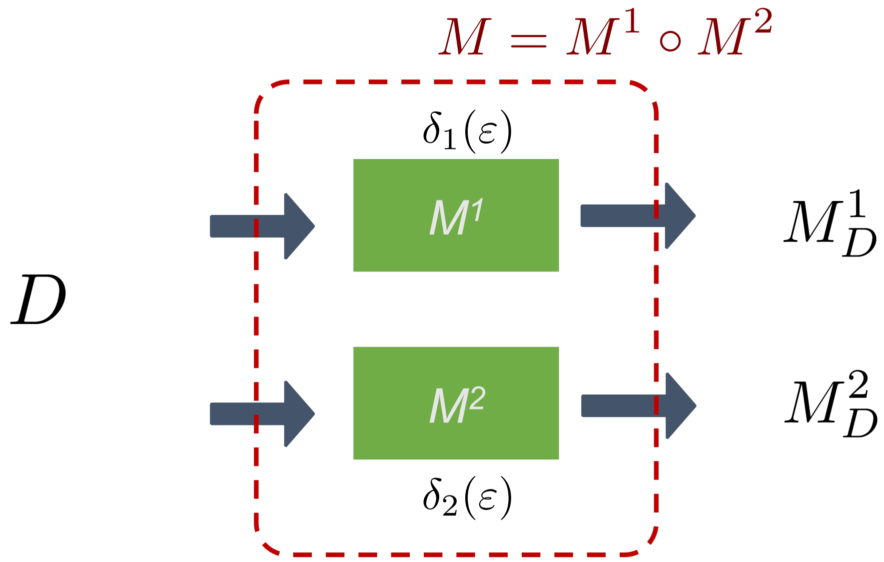

What is the privacy curve $\delta(\varepsilon)$ of M?

Composition of DP

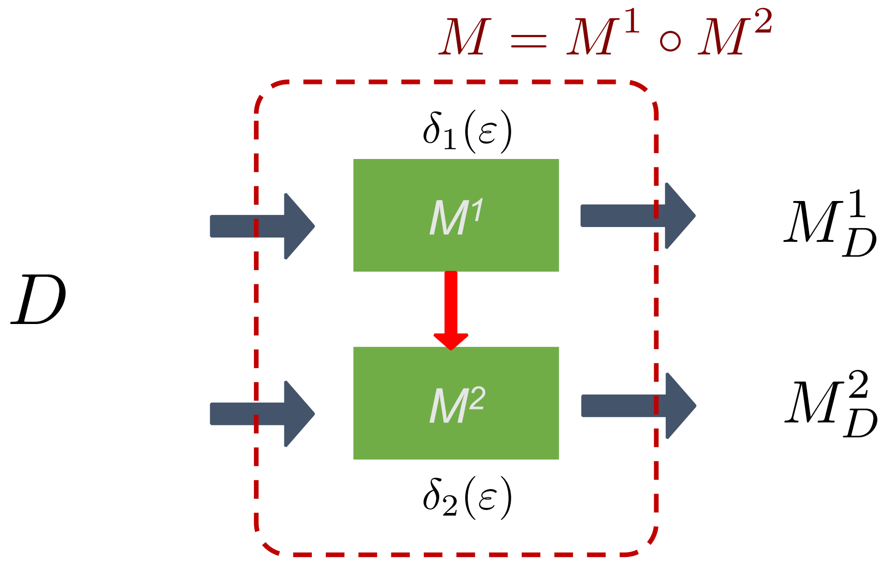

What is the privacy curve $\delta(\varepsilon)$ of M?

Composition of DP

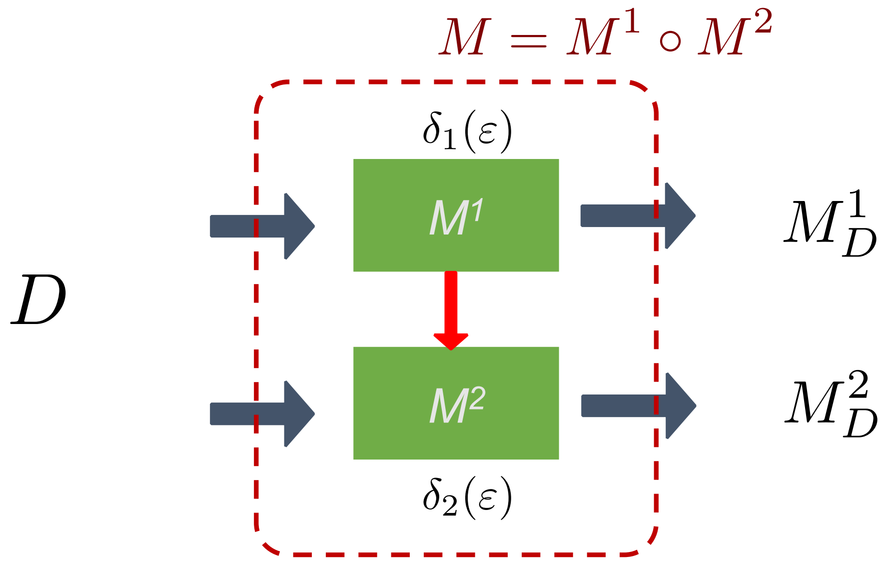

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^1, Q^1) tightly dominates M^1

(P^2, Q^2) tightly dominates M^2

Composition of DP

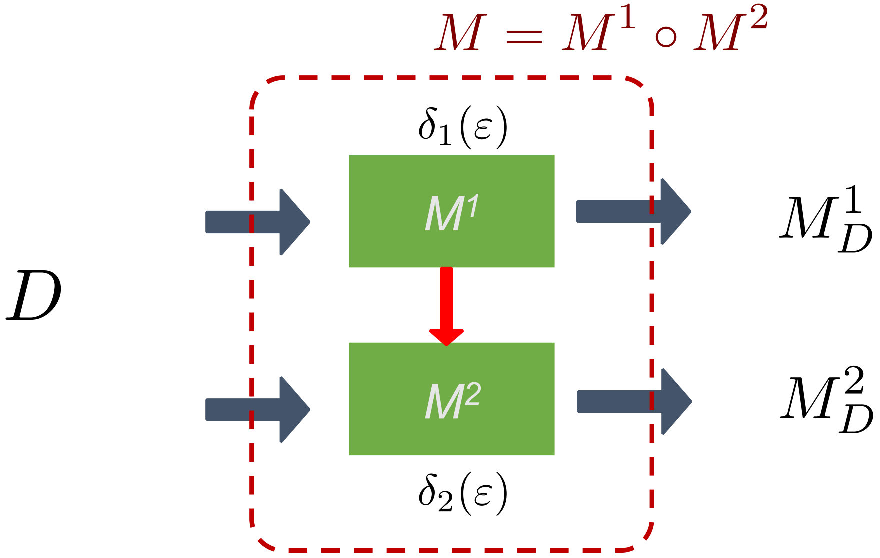

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^1, Q^1) tightly dominates M^1

(P^2, Q^2) tightly dominates M^2

$\Longrightarrow (P^1 \times P^2, Q^1 \times Q^2)$ dominates M

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

$M = M^1 \circ M^2$

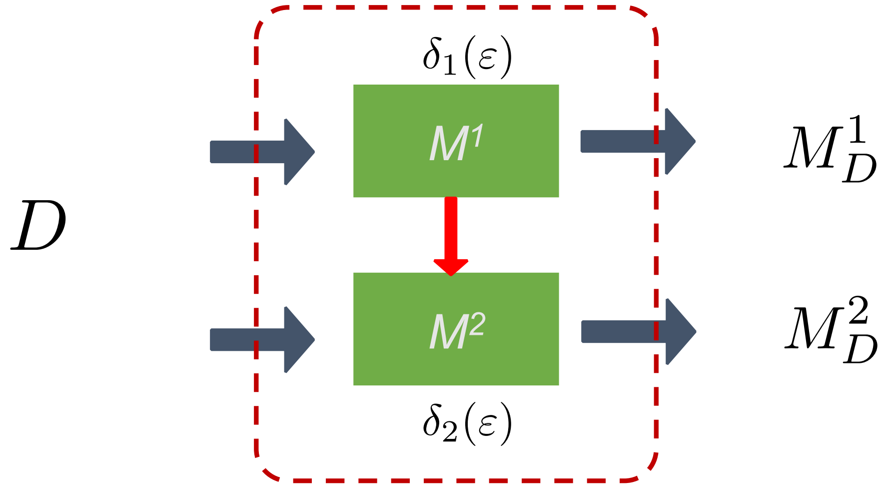

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^1, Q^1) tightly dominates M^1

(P^2, Q^2) tightly dominates M^2

$\Longrightarrow$ $(P^1 \times P^2, Q^1 \times Q^2)$ dominates M

$$\Downarrow$$

$$\delta(\varepsilon) \quad \leq \quad \mathsf{E}_\varepsilon(P^1 \times P^2 \| Q^1 \times Q^2)$$

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

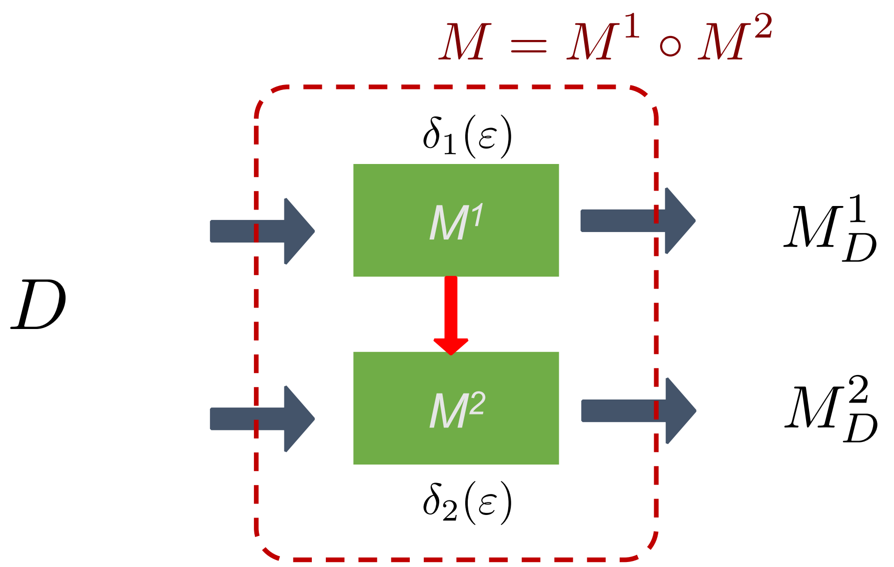

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^1, Q^1) tightly dominates M^1

(P^2, Q^2) tightly dominates M^2

$\Longrightarrow$ $(P^1 \times P^2, Q^1 \times Q^2)$ dominates M

$\Downarrow$

$$\delta(\varepsilon) \quad \le \quad \mathsf{E}_\varepsilon(P^1 \times P^2 \| Q^1 \times Q^2) = \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + L_2)}\right)_+\right]$$

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

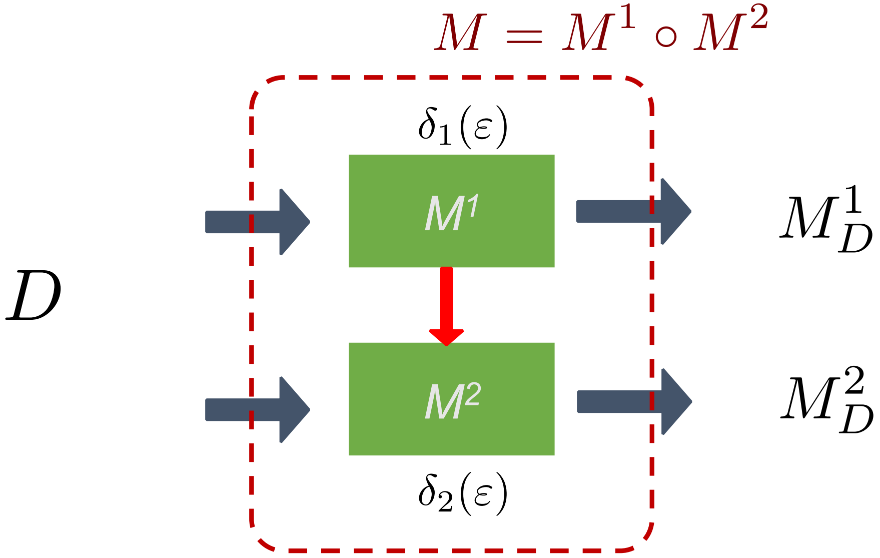

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^1, Q^1) tightly dominates M^1

(P^2, Q^2) tightly dominates M^2

$\Longrightarrow (P^1 \times P^2, Q^1 \times Q^2)$ dominates M

$\Downarrow$

PLRV for M^1 PLRV for M^2

$$\delta(\varepsilon) \leq \mathsf{E}_\varepsilon(P^1 \times P^2 \| Q^1 \times Q^2) = \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + L_2)}\right)_+\right]$$

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

$M = M^1 \circ \cdots \circ M^n$

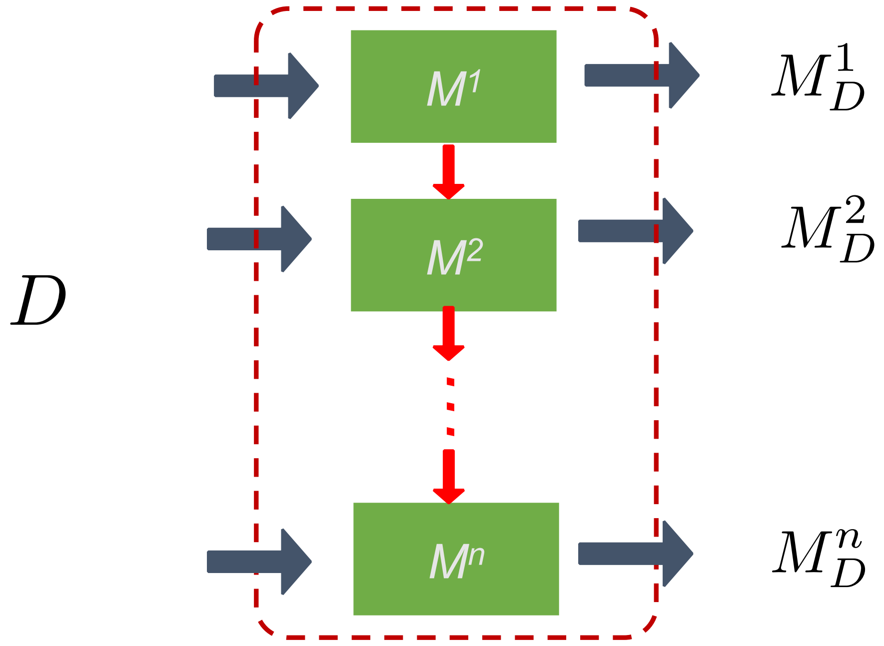

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^i, Q^i) tightly dominates M^i $\Longrightarrow$ $(P^1 \times \cdots \times P^n, Q^1 \times \cdots \times Q^n)$ dominates M

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

$M = M^1 \circ \cdots \circ M^n$

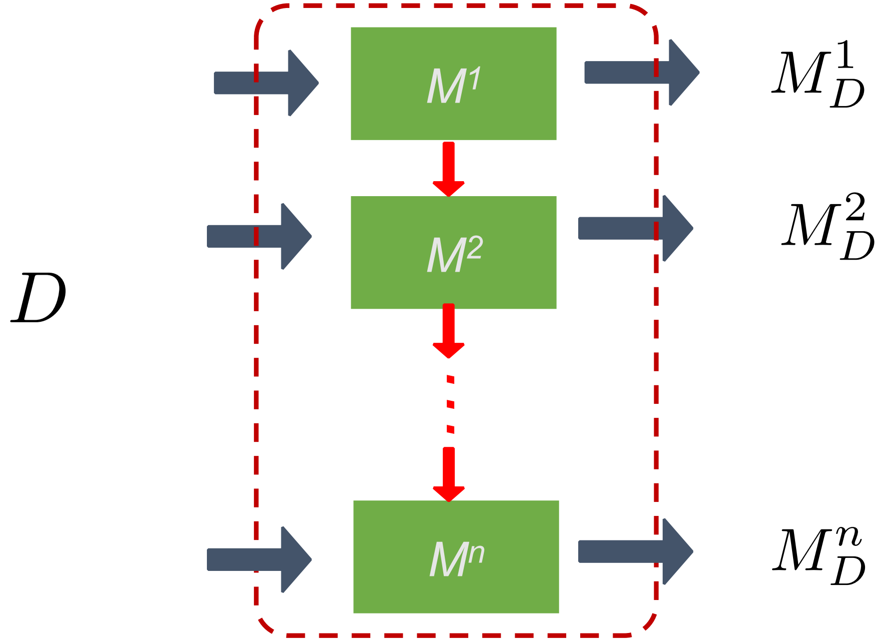

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^i, Q^i) tightly dominates M^i $\Longrightarrow$ $(P^1 \times \cdots \times P^n, Q^1 \times \cdots \times Q^n)$ dominates M

$$\Downarrow$$

$$\delta(\varepsilon) \leq \mathsf{E}_\varepsilon(P^1 \times \cdots \times P^n \| Q^1 \times \cdots \times Q^n)$$

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition of DP

$M = M^1 \circ \cdots \circ M^n$

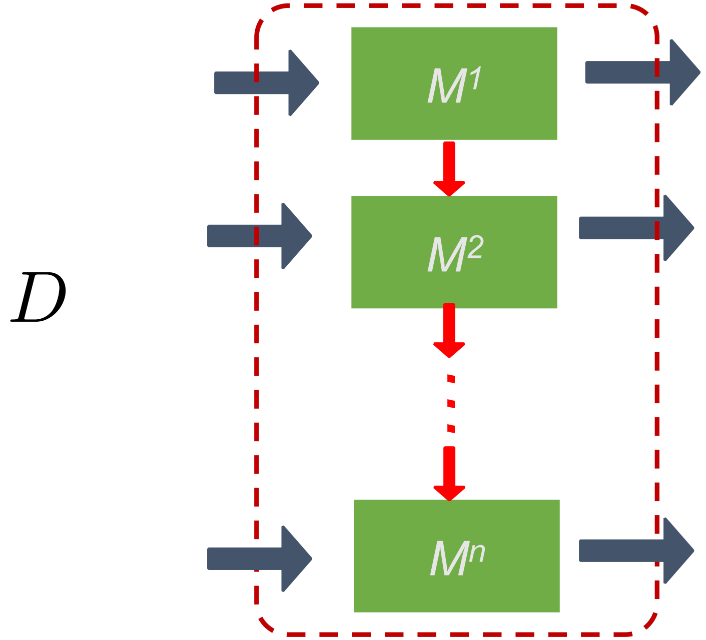

M_D^1

M_D^2

M_D^n

What is the privacy curve $\delta(\varepsilon)$ of M?

(P^i, Q^i) tightly dominates M^i $\implies$ $(P^1 \times \cdots \times P^n, Q^1 \times \cdots \times Q^n)$ dominates M

$$\Downarrow$$

$$\delta(\varepsilon) \leq \mathsf{E}_\varepsilon(P^1 \times \cdots \times P^n \| Q^1 \times \cdots \times Q^n) = \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + \cdots + L_n)}\right)_+\right]$$

Zhu, Dong, and Wang, Optimal Accounting of Differential Privacy via Characteristic Function, AISTAT 2022

Composition Results

$$\delta(\varepsilon) \;\le\; \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + \cdots + L_n)}\right)_+\right]$$

Composition Results

$$\delta(\varepsilon) \;\leq\; \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + \cdots + L_n)}\right)_+\right]$$

- Moments accountant: [Abadi et al'16], [Mironov'17]
- Central limit theorem: [Dong et al'19], [Sommer et al'19]
- Fast Fourier transform: [Koskela et al'20], [Koskela and Honkela'20], [Koskela et al'21], [Gopi et al'21], [Ghazi et al'22]
- Characteristic function: [Zhu et al'22]
- Piece-wise linearization of HS divergence: [Doroshenko et al'22]

Composition Results

$$\delta(\varepsilon) \;\leq\; \mathbb{E}\left[\left(1 - e^{\varepsilon-(L_1+\cdots+L_n)}\right)_+\right]$$

- Moments accountant: [Abadi et al'16], [Mironov'17]
- Central limit theorem: [Dong et al'19], [Sommer et al'19]
- Fast Fourier transform: [Koskela et al'20], [Koskela and Honkela'20], [Koskela et al'21], [Gopi et al'21], [Ghazi et al'22]
- Characteristic function: [Zhu et al'22]
- Piece-wise linearization of HS divergence: [Doroshenko et al'22]

Saddle-point accountant

$$\delta(\varepsilon) \; \leq \; \mathbb{E}\left[\left(1 - e^{\varepsilon - (L_1 + \cdots + L_n)}\right)_+\right]$$

Saddle-point accountant

$$\delta(\varepsilon) \;\le\; \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right]$$

Saddle-point accountant

$$\delta(\varepsilon) \;\leq\; \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell)\mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \le \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell) \mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell) \mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \;\le\; \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell) \mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell) \mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]} \mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \;\leq\; \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell)\mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell)\mathrm{d}\ell$$

MGF of L

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]}\mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \le \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell)\,\mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell)\,\mathrm{d}\ell$$

MGF of L

pdf of $\tilde{L}$ exponentially tilted of L

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]}\,\mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon-\ell}\right)_+ f_L(\ell) \mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon-\ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell) \mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon-\ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]} \mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} e^{-t\ell} \left(1 - e^{\varepsilon-\ell}\right)_+ f_{\tilde{L}}(\ell) \mathrm{d}\ell$$

Saddle-point accountant

$$\delta(\varepsilon) \le \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell) \mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell) \mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]} \mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} e^{-t\ell} \left(1 - e^{\varepsilon - \ell}\right)_+ f_{\tilde{L}}(\ell) \mathrm{d}\ell$$

Plancherel's Theorem: $\displaystyle \int_{\mathbb{R}} f(x) g(x) \mathrm{d}x = \frac{1}{2\pi} \int_{\mathbb{R}} \mathcal{F}(f)(\omega) \overline{\mathcal{F}(g)(\omega)} \mathrm{d}\omega$

Saddle-point accountant

$$\delta(\varepsilon) \le \mathbb{E}\Big[\big(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\big)_+\Big] = \int_{\mathbb{R}} \big(1 - e^{\varepsilon - \ell}\big)_+ f_L(\ell)\mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \big(1 - e^{\varepsilon - \ell}\big)_+ e^{-t\ell} e^{t\ell} f_L(\ell)\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \big(1 - e^{\varepsilon - \ell}\big)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]}\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} e^{-t\ell} \big(1 - e^{\varepsilon - \ell}\big)_+ f_{\tilde{L}}(\ell)\mathrm{d}\ell$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

Saddle-point accountant

$$\delta(\varepsilon) \le \mathbb{E}\left[\left(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\right)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell)\mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell)\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]}\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} e^{-t\ell} \left(1 - e^{\varepsilon - \ell}\right)_+ f_{\tilde{L}}(\ell)\mathrm{d}\ell$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t + is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1 + z)$$

Saddle-point accountant

$$\delta(\varepsilon) \;\le\; \mathbb{E}\left[\Big(1 - e^{\varepsilon - \overbrace{(L_1 + \cdots + L_n)}^{L}}\Big)_+\right] = \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ f_L(\ell)\mathrm{d}\ell$$

$$= \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} e^{t\ell} f_L(\ell)\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} \left(1 - e^{\varepsilon - \ell}\right)_+ e^{-t\ell} \frac{e^{t\ell} f_L(\ell)}{\mathbb{E}[e^{tL}]}\mathrm{d}\ell$$

$$= \mathbb{E}[e^{tL}] \int_{\mathbb{R}} e^{-t\ell} \left(1 - e^{\varepsilon - \ell}\right)_+ f_{\tilde{L}}(\ell)\mathrm{d}\ell$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

cumulant generating function (CGF) of L

$$K_L(z) = \log \mathbb{E}\left[e^{zL}\right]$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε:

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F_\varepsilon'(t_*) = 0$, i.e.,

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F_\varepsilon'(t_*) = 0$, i.e.,

$$\boxed{K_L'(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F_\varepsilon''(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*) \implies$$

Along the real line, F_ε is minimized at $z = t_*$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*) \implies$$

Along the real line, F_ε is <u>minimized</u> at $z = t_*$

Parallel to imaginary axis, F_ε is <u>maximized</u> at $z = t_*$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F_\varepsilon'(t_*) = 0$, i.e.,

$$\boxed{K_L'(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

First approximation:

Second approximation:

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

First approximation: Vanilla saddle-point approximation

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation: Vanilla saddle-point approximation

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_* + 1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K''_L(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

$$\Downarrow$$

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s \approx \frac{e^{F_\varepsilon(t_*)}}{\sqrt{2\pi F''_\varepsilon(t_*)}}$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

$$\Downarrow$$

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s \approx \frac{e^{F_\varepsilon(t_*)}}{\sqrt{2\pi F''_\varepsilon(t_*)}} \boxed{= \frac{e^{K_L(t_*)-\varepsilon t_*}}{\sqrt{2\pi\left[t_*^2(1+t_*)^2 K''_L(t_*) + t_*^2 + (t_*+1)^2\right]}}}$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F_\varepsilon'(t_*) = 0$, i.e.,

$$\boxed{K_L'(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F_\varepsilon''(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K_L''(t_*)$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

$$\Downarrow$$

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s \approx e^{K_L(t_*)-\varepsilon t_*}\mathbb{E}\left[e^{t_*(\varepsilon-Z)}\left(1-e^{\varepsilon-Z}\right)_+\right]$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

$$\Downarrow$$

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s \approx \boxed{e^{K_L(t_*)-\varepsilon t_*}\mathbb{E}\left[e^{t_*(\varepsilon-Z)}\left(1-e^{\varepsilon-Z}\right)_+\right]}$$

Saddle-point accountant

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Take t to be saddle-point of F_ε: Unique t_* satisfying $F'_\varepsilon(t_*) = 0$, i.e.,

$$\boxed{K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}}$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z-t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z-t_*)^2 K''_L(t_*)$$

$$\Downarrow$$

$$\delta(\varepsilon) \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)}\mathrm{d}s \approx \boxed{e^{K_L(t_*)-\varepsilon t_*}\mathbb{E}\left[e^{t_*(\varepsilon-Z)}\left(1-e^{\varepsilon-Z}\right)_+\right]}$$

$$Z \sim \mathcal{N}\left(K'_L(t_*), K''_L(t_*)\right)$$

Saddle-point accountant: Algorithm

Input: Tightly dominating pairs $\{(P_i, Q_i)\}_{i=1}^n$ for mechanisms $\{M_i\}_{i=1}^n$ and ε

- Compute (numerically estimate) $K_{L_i}(t) = \log \mathbb{E}[e^{tL_i}]$
- $K_L(t) = \sum_i K_{L_i}(t)$ (since L_i are independent)
- Find saddle-point t_* by solving $K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}$

Outputs:

Saddle-point accountant: Algorithm

Input: Tightly dominating pairs $\{(P_i, Q_i)\}_{i=1}^n$ for mechanisms $\{M_i\}_{i=1}^n$ and ε

- Compute (numerically estimate) $K_{L_i}(t) = \log \mathbb{E}[e^{tL_i}]$
- $K_L(t) = \sum_i K_{L_i}(t)$ (since L_i are independent)
- Find saddle-point t_* by solving $K'_L(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}$

Outputs:

$$\hat{\delta}_1(\varepsilon) = \frac{e^{K_L(t_*)-\varepsilon t_*}}{\sqrt{2\pi\left[t_*^2(1+t_*)^2 K''_L(t_*) + t_*^2 + (t_*+1)^2\right]}}$$

$$\hat{\delta}_2(\varepsilon) = e^{K_L(t_*)-\varepsilon t_*}\mathbb{E}\left[e^{t_*(\varepsilon - Z)}\left(1 - e^{\varepsilon - Z}\right)_+\right]$$

Saddle-point accountant: Algorithm

Input: Tightly dominating pairs $\{(P_i, Q_i)\}_{i=1}^n$ for mechanisms $\{M_i\}_{i=1}^n$ and ε

- Compute (numerically estimate) $K_{L_i}(t) = \log \mathbb{E}[e^{tL_i}]$
- $K_L(t) = \sum_i K_{L_i}(t)$ (since L_i are independent)
- Find saddle-point t_* by solving $K_L'(t_*) = \varepsilon + \frac{1}{t_*} + \frac{1}{t_*+1}$

Outputs:

$$\hat{\delta}_1(\varepsilon) = \frac{e^{K_L(t_*)-\varepsilon t_*}}{\sqrt{2\pi\left[t_*^2(1+t_*)^2 K_L''(t_*) + t_*^2 + (t_*+1)^2\right]}}$$

$$\hat{\delta}_2(\varepsilon) = e^{K_L(t_*)-\varepsilon t_*}\,\mathbb{E}\left[e^{t_*(\varepsilon - Z)}\left(1 - e^{\varepsilon - Z}\right)_+\right]$$

Both are "corrected" versions of moments accountant

Numerical experiments

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

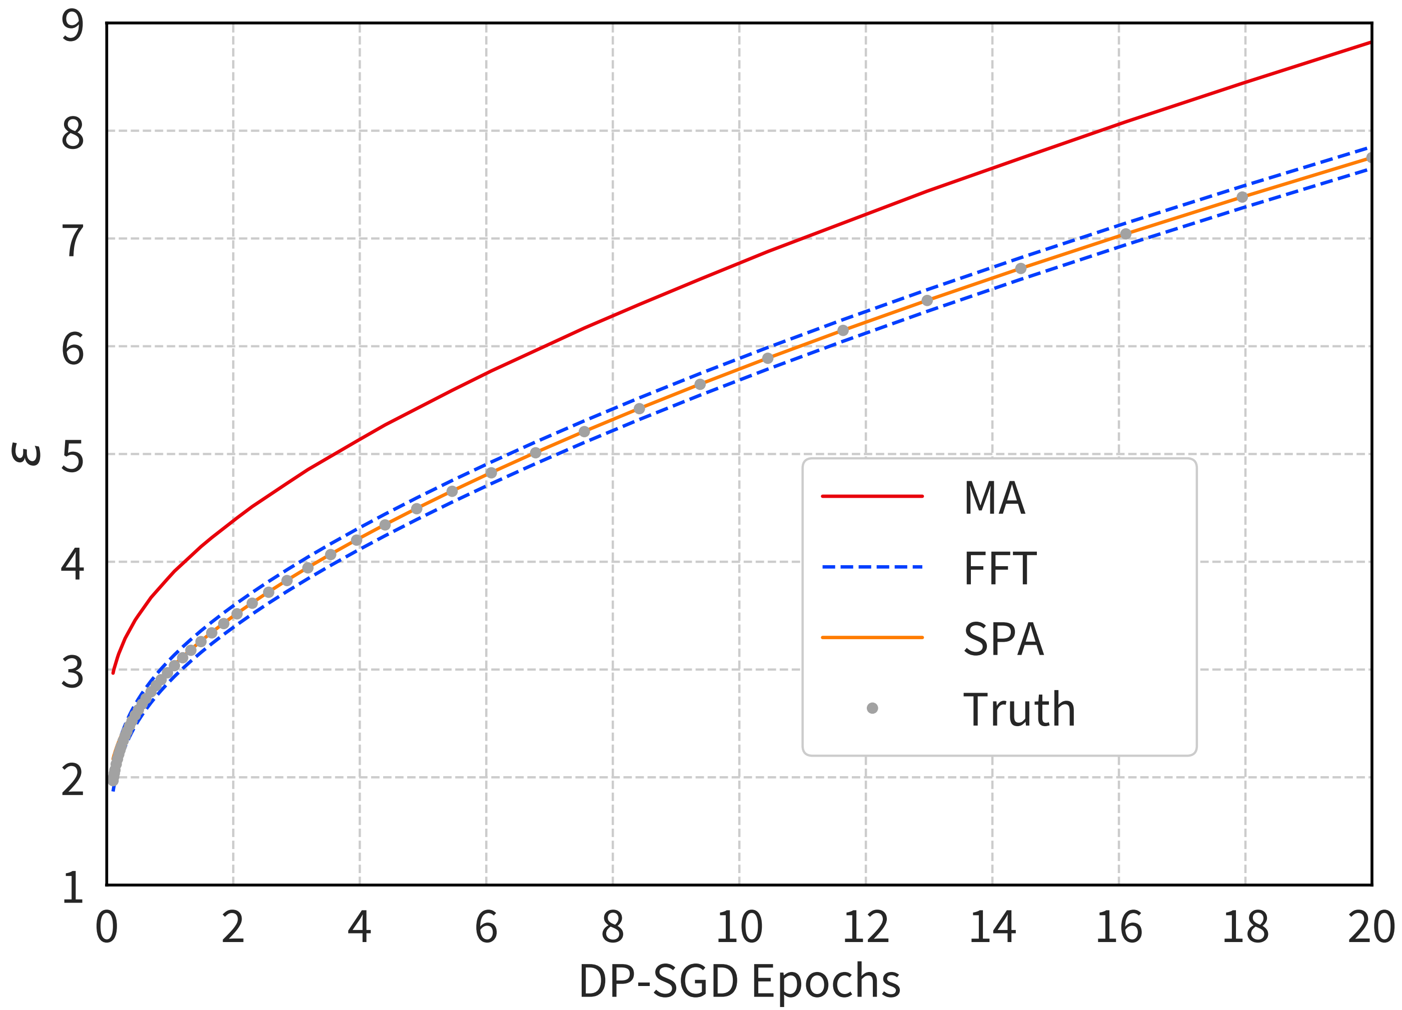

Numerical experiments

DP-SGD:

$\sigma = 0.65$

subsampling $= 10^{-2}$

$\delta = 10^{-5}$

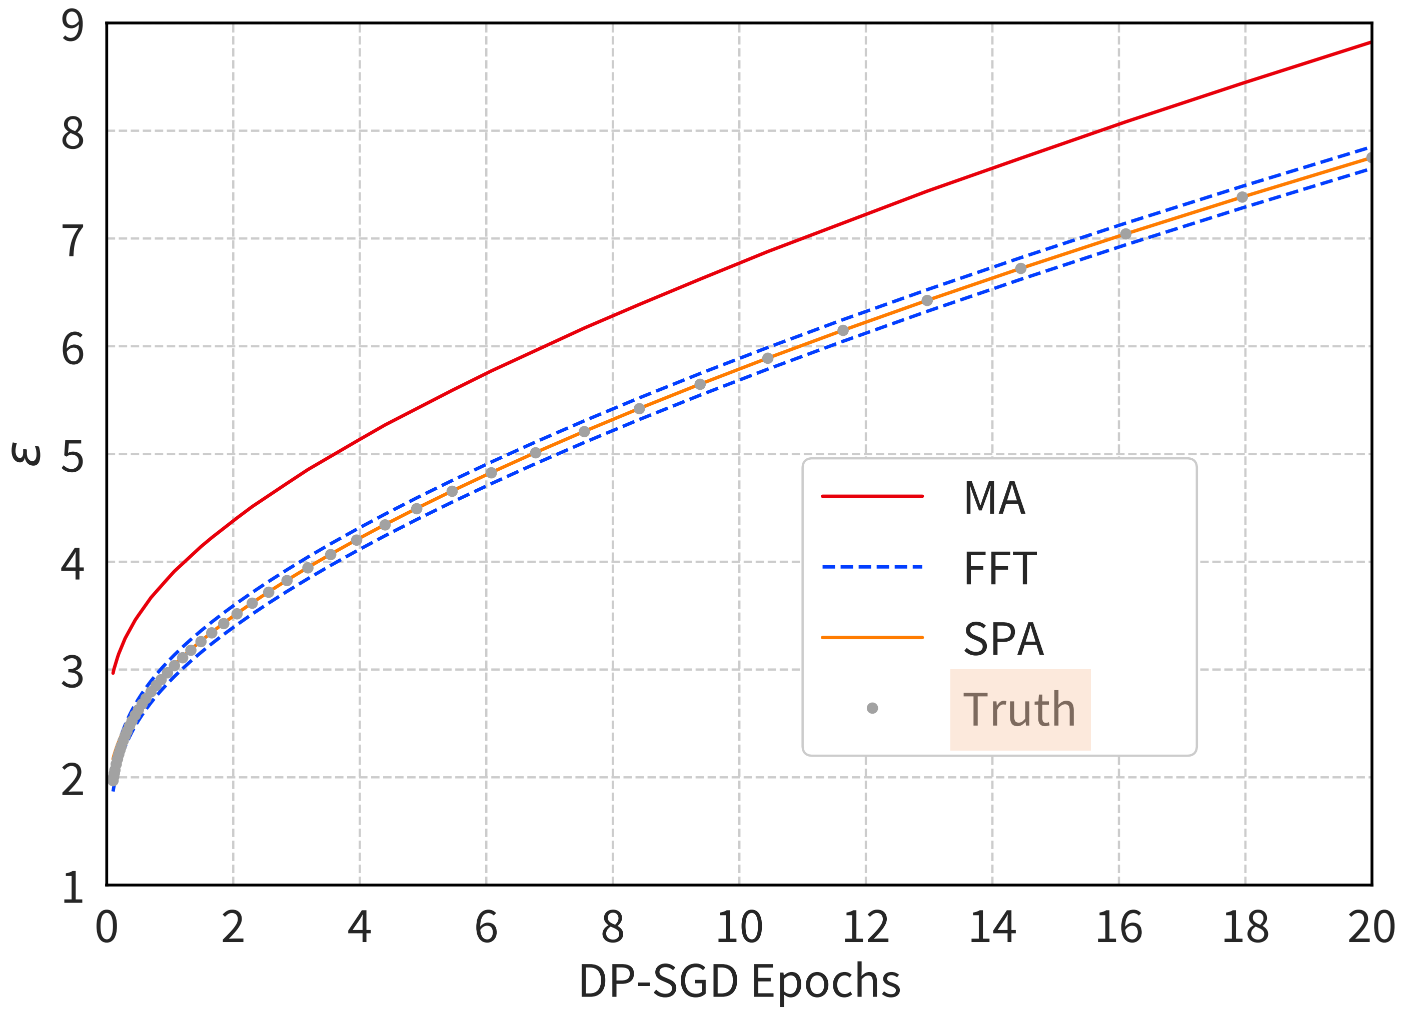

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1+z)$$

Numerical experiments

DP-SGD:

$\sigma = 1$

subsampling $= 10^{-2}$

$n = 2000$

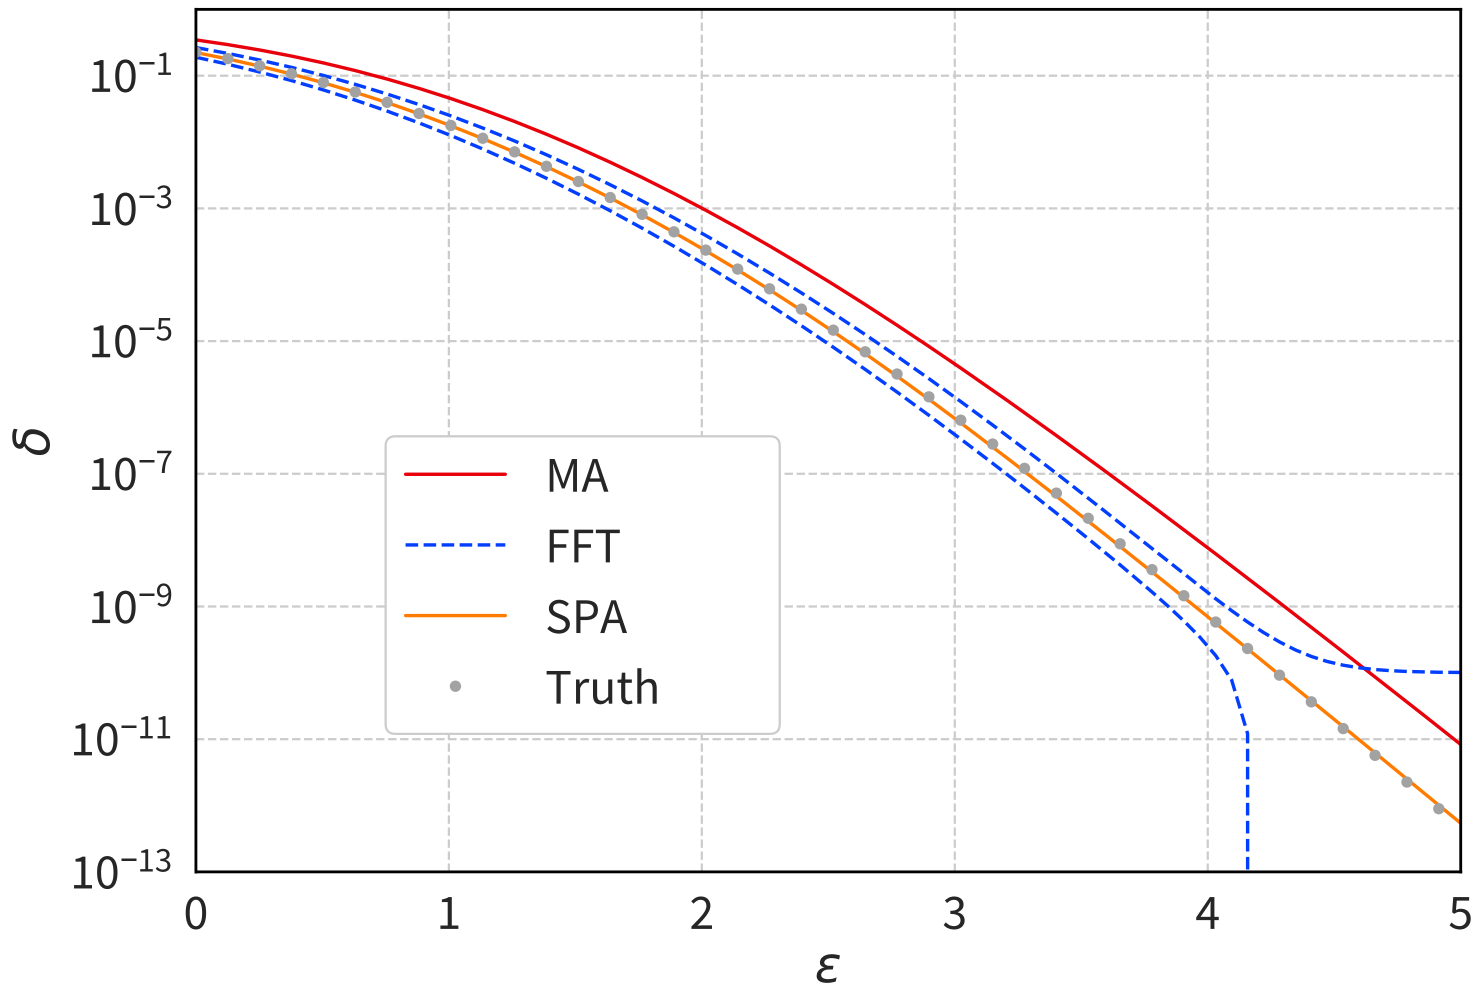

Error analysis

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K''_L(t_*)$$

Error analysis

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F_\varepsilon''(t_*)$$

vanilla saddle-point approximation

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K_L''(t_*)$$

Edgeworth expansion

Error analysis

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F_\varepsilon''(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K_L''(t_*)$$

For any $\varepsilon \geq 0$

$$\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s - \hat{\delta}_2(\varepsilon) \right| \leq e^{K_L(t_*) - \varepsilon t_*} \left(\frac{t_*}{1 + t_*} \right)^{t_*} \frac{P_{t_*}}{K_L''(t_*)^{3/2}}$$

Error analysis

$$\delta(\varepsilon) \le \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F_\varepsilon''(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K_L''(t_*)$$

$$\sum_{i=1}^{n} \mathbb{E}[|\tilde{L}_i - \mathbb{E}[\tilde{L}_i]|^3]$$

For any $\varepsilon \ge 0$

$$\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s - \hat{\delta}_2(\varepsilon) \right| \le e^{K_L(t_*) - \varepsilon t_*} \left(\frac{t_*}{1 + t_*} \right)^{t_*} \frac{P_{t_*}}{K_L''(t_*)^{3/2}}$$

Error analysis

$$\delta(\varepsilon) \leq \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s$$

First approximation:

$$F_\varepsilon(z) \approx F_\varepsilon(t_*) + \frac{1}{2}(z - t_*)^2 F''_\varepsilon(t_*)$$

Second approximation:

$$K_L(z) \approx K_L(t_*) + \frac{1}{2}(z - t_*)^2 K''_L(t_*)$$

For any $\varepsilon = \mathbb{E}[L] + b \cdot \mathrm{var}(L)$, we have

$$\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{F_\varepsilon(t+is)} \mathrm{d}s - \hat{\delta}_2(\varepsilon) \right| \leq \frac{C}{\sqrt{n}}$$

Error analysis

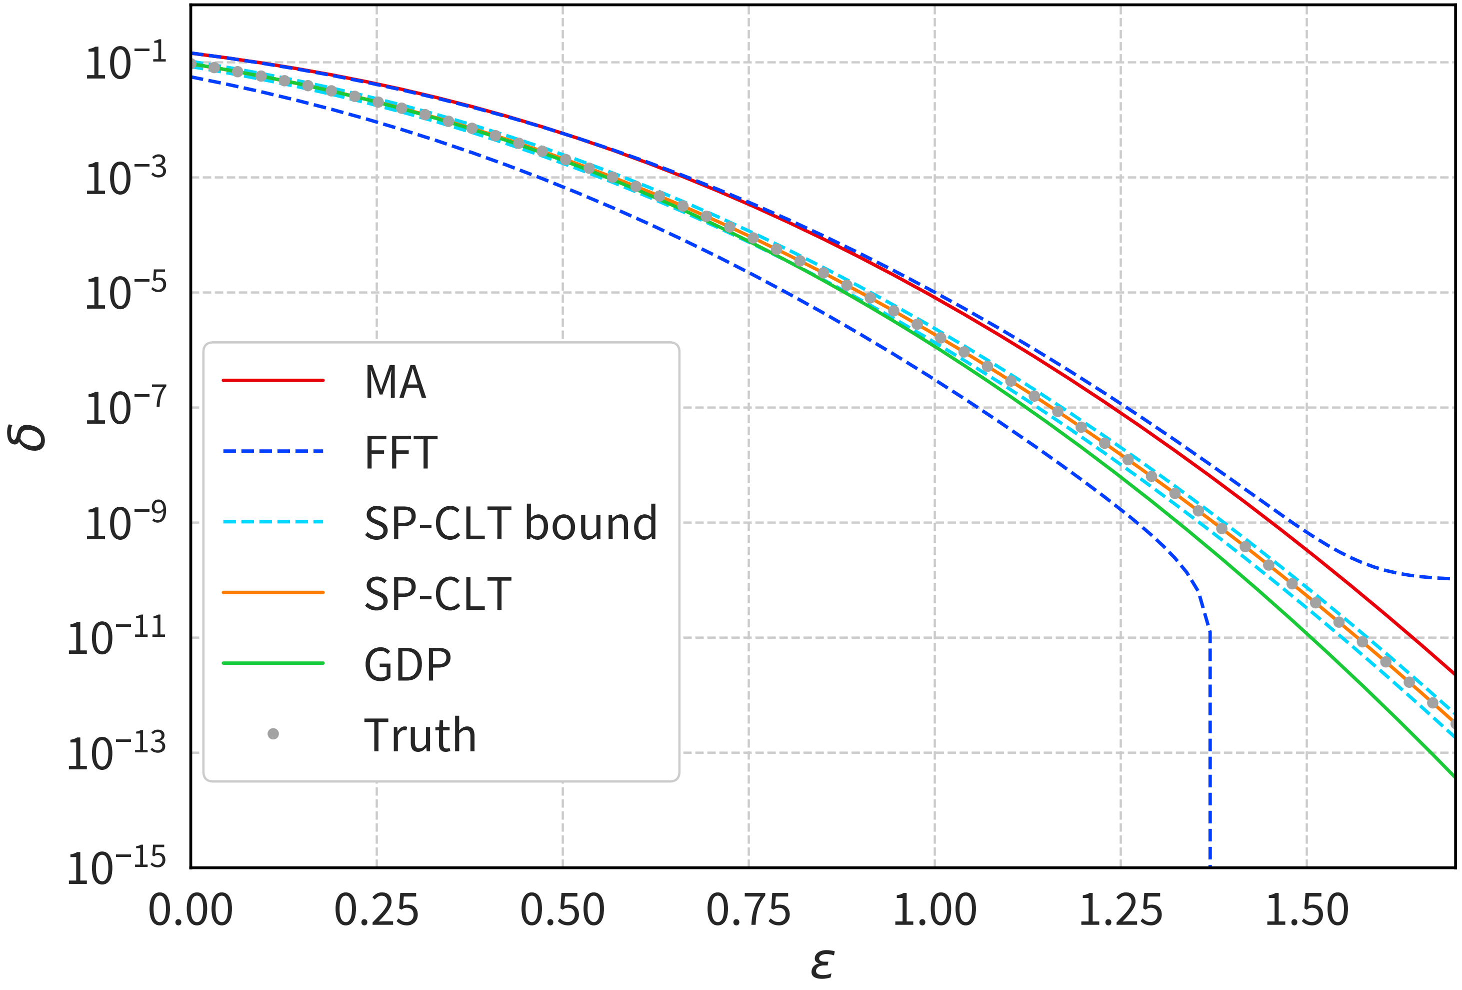

Summary

- Accountant algorithm comparable with state-of-the-art
- Runtime complexity independent of # composition
- Works for all epsilon and delta

Summary

- Accountant algorithm comparable with state-of-the-art
- Runtime complexity independent of # composition
- Works for all epsilon and delta

Ongoing work

- Tightly dominating distribution of composed mechanisms
- Tighter error analysis
- Efficient estimator for CGF

Summary

- Accountant algorithm comparable with state-of-the-art
- Runtime complexity independent of # composition
- Works for all epsilon and delta

Ongoing work

- Tightly dominating distribution of composed mechanisms
- Tighter error analysis
- Efficient estimator for CGF

Pre-print available here!