Exponential quantum space advantage for Max-kSAT in the streaming setting

Haoyu Wang (Penn State University), Guangxu Yang (University of Southern California)

July 23, 2026

Part I: introduction

Streaming Algorithm

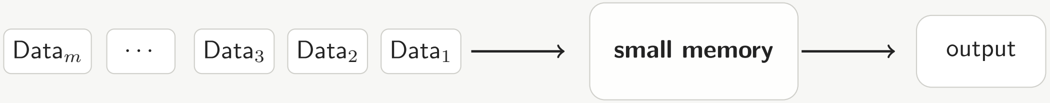

- Data arrives sequentially.
- Massive data vs. Limited memory.

Streaming Algorithm

Goal: efficient computation under **space** limitation.

Applications: Distinct elements; Heavy hitters; Frequency moments; Estimation on massive graphs(connectivity, matching, cuts, triangle counting...)

Scenarios: Network monitoring, Traffic engineering, Sensor networks, etc.

Quantum Time Advantage

Quantum time advantages:

- Shor's algorithm: exponential **speedup** on factoring.
- Grover's algorithm: $\sqrt{N}$ vs. N time advantage on unstructured search.

A major challenge: how to prove exponential quantum time advantage.

Quantum Space Advantage

Motivation:

- Space vs. Time.
- Space is crucial in quantum computation: Fault-tolerant qubits are costly.

Question: Can quantum algorithms save not only time, but also space?

Research background

1. **[F. Le Gall, SPAA 2006]**: Repeated-input Set Disjointness problem; $O(\log n)$ quantum space vs. $\Omega(n^{1/3})$ classical space.
2. **[Gavinsky, Kempe, Kerenidis, Raz, and de Wolf, SICOMP 2008]**: Boolean Hidden Matching variant problem; $O(\log n)$ quantum space vs. $\Omega(\sqrt{n})$ classical space.
3. **[Kallaugher, FOCS 2021]**: Triangle counting has polynomial quantum advantage.
4. **[Kallaugher and Parekh, FOCS 2022]**: Max-Cut has no quantum advantage.
5. **[Kallaugher, Parekh, and Voronova, STOC 2024]**: Max-DiCut has exponential quantum advantage.
6. **[Feng, Xu, Li, Guo and Lin, preprint 2026]**: Shannon entropy estimation has exponential quantum advantage.

Question: which CSPs admit exponential quantum space advantage in the streaming model?

MAX-kSAT

Given variables $\mathbf{x} := \{x_1, \dots x_n\}$ and literal $lit \in \mathbf{x} \cup \neg\mathbf{x}$, define the clause:

$$C := lit_1 \vee lit_2 \cdots \vee lit_t, \qquad t \leq k.$$

Given clauses $\Phi := \{C_1, \dots, C_m\}$, define the optimal value

$$\mathrm{OPT}(\Phi) = \max_{\mathbf{x} \in \{0,1\}^n} \sum_{i=1}^{m} C_i(\mathbf{x}).$$

To find an approximation Z with

$$\rho \cdot \mathrm{OPT}(\Phi) \leq Z \leq \mathrm{OPT}(\Phi),$$

where $0 \leq \rho \leq 1$.

Data Stream: clauses arrive one by one as $C_1,\ C_2,\ \ldots,\ C_m$.

Main theorem

Theorem 1. For every fixed $k \geq 2$, there is a one-pass quantum streaming algorithm for Max-kSAT using $\mathrm{polylog}(n)$ space, achieving 0.7172-approximation.

Theorem 2 [Chou-Golovnev-Velusamy, FOCS 2020]. Any one-pass classical streaming algorithm achieving 0.7071-approximation for Max-kSAT requires $\Omega(\sqrt{n})$ space.

quantum: $\mathrm{polylog}(n)$ qubits. vs. classical: $\Omega(\sqrt{n})$ bits

Boolean Max-2CSP Taxonomy

Thresholds for one-pass streaming by $\mathrm{polylog}(n)$ space.

type	special case	classical	quantum	advantage
TR	Trivial	1	1	No
OR	Max-2EOR	$3/4$	$3/4$	No
TR+OR	**Max-2SAT**	0.7071	$[\mathbf{0.7425}^{a}, 3/4]$	**Yes**
XOR	Max-Cut	$1/2$	$1/2$	No
AND	Max-DiCut	$4/9$	$[0.4844, 1/2]$	Yes

Classical thresholds: [CGV20]; Max-Cut/Max-2EOR: [KP22]; Max-DiCut: [KPV24].

[a] 0.7425 is optimized specially for Max-2SAT, by the same framework of general MAX-kSAT.

Part II: from Max-kSAT to snapshot

Goal: Design a statistic (snapshot) suitable for streaming algorithm.

Baby strategy: For any variable x_v, let p_v be the positive degree and q_v be the negative degree.

- If $p_v - q_v \geq 0$, let $x_v \leftarrow 1$.
- If $p_v - q_v < 0$, let $x_v \leftarrow 0$.

Teenager strategy: Allocate the variable x_v into **buckets/intervals** of $[-1, 1]$ according to its **bias**

$$b_v = (p_v - q_v)/(p_v + q_v).$$

Randomly assign x_v according to the bucket.

Weighted Bias

Definition (Weighted Bias)

For clause length $\ell = 1, 2, 3$, let p_v^ℓ and q_v^ℓ be the positive and negative degrees of variable x_v in clauses with length ℓ. Given weights $(\omega_1, \omega_2, \omega_3) = (4, 2, 1)$, let

$$p_v \leftarrow \omega_1 p_v^1 + \omega_2 p_v^2 + \omega_3 p_v^3, \quad q_v \leftarrow \omega_1 q_v^1 + \omega_2 q_v^2 + \omega_3 q_v^3.$$

Define the weighted bias of x_v as

$$b_v := (p_v - q_v)/(p_v + q_v).$$

Buckets: partition $[-1,1]$ into N intervals:

$$[-1,1] = I_1 \cup I_2 \cup \cdots \cup I_N.$$

For variable x_v with $b_v \in I_i$, let $\Pr[x_v = 1] = r_i$, $\forall i \in [N]$.

Satisfaction probability: For each literal with sign $s \in \{+,-\}$ and bucket index $i \in [N]$, let the sign-bucket pair $\alpha := (s,i)$. Define

$$q_\alpha := \begin{cases} r_i & \text{if } s = +, \\ 1 - r_i & \text{if } s = -. \end{cases}$$

Snapshot

For each clause C, let the j-th literal has sign s_j and bucket index i_j. Let the sign-bucket pair be $\alpha(C,j) := (s_j, i_j)$.

Definition (Snapshot)

Define the snapshot as

$$\begin{aligned}
U_{\ell,j,\alpha_1} &\longleftarrow \#\{C : |C| = \ell,\ \alpha(C,j) = \alpha_1\},\\
P_{\ell,(j,t),(\alpha_1,\alpha_2)} &\longleftarrow \#\{C : |C| = \ell,\ (\alpha(C,j),\alpha(C,t)) = (\alpha_1,\alpha_2)\},\\
T_{\alpha_1,\alpha_2,\alpha_3} &\longleftarrow \#\{C : |C| = 3,\ (\alpha(C,1),\alpha(C,2),\alpha(C,3)) = (\alpha_1,\alpha_2,\alpha_3)\}.
\end{aligned}$$

Definition (Raw Snapshot score)

Define the raw snapshot score as

$$L_{\le 3}^{\mathrm{raw}} = \sum_{\alpha} q_\alpha \cdot U_{1,1,\alpha} + \sum_{\alpha,\beta} \Big(1 - (1-q_\alpha)(1-q_\beta)\Big) \cdot P_{2,(1,2),(\alpha,\beta)}$$
$$+ \underbrace{\sum_{\alpha,\beta,\gamma} \Big(1 - (1-q_\alpha)(1-q_\beta)(1-q_\gamma)\Big) \cdot T_{\alpha,\beta,\gamma}}_{\text{ternary term}} .$$

Upper bound: Comparing $L_{\le 3}^{\mathrm{raw}}$ with random rounding, we have

$$L_{\le 3}^{\mathrm{raw}} \le \mathrm{OPT}(\Phi_{\le 3})$$

.

Simplify the ternary term: Define

$$H(\alpha_1, \alpha_2, \alpha_3) = c_0 + \sum_{j=1}^{3} u_j(\alpha_j) + \sum_{1 \le j < t \le 3} p_{jt}(\alpha_j, \alpha_t).$$

Choose c_0, u_j and p_{jt} properly (**To be determined later**) to ensure

$$H(\alpha_1, \alpha_2, \alpha_3) \le 1 - (1 - q_{\alpha_1})(1 - q_{\alpha_2})(1 - q_{\alpha_3}).$$

Definition (Snapshot score)

Define the snapshot score as

$$L_{\le 3} := \sum_{\alpha} q_\alpha \cdot U_{1,1,\alpha} + \sum_{\alpha,\beta} \Big(1 - (1-q_\alpha)(1-q_\beta)\Big) \cdot P_{2,1,2,\alpha,\beta} + c_0 \cdot M_3 + \sum_{j=1}^{3} \sum_{\alpha} u_j(\alpha) \cdot U_{3,j,\alpha} + \sum_{1\le j<t\le 3} \sum_{\alpha,\beta} p_{jt}(\alpha,\beta) \cdot P_{3,j,t,\alpha,\beta}.$$

Upper bound: $L_{\le 3} \le L^{\mathrm{raw}}_{\le 3} \le \mathrm{OPT}(\Phi_{\le 3})$

Lower bound

Task: $L_{\leq 3}$ is regardless of assignment. To achieve

$$L_{\leq 3} \geq \rho \cdot \mathrm{OPT}(\Phi_{\leq 3}),$$

we need to compare it with the optimal assignment.

Definition (Typed atom)

For a clause C together with an assignment $\tau \in \{0,1\}^\ell$, we define a **typed atom** c as

$$c := \mathrm{type}(C,\tau) := (\underbrace{\ell}_{\text{clause length}}; \underbrace{s_1,\ldots,s_\ell}_{\text{sign}}; \underbrace{i_1,\ldots,i_\ell}_{\text{bucket index}}; \underbrace{\tau_1,\ldots,\tau_\ell}_{\text{assignment}}).$$

Comparison on atom level: the snapshot score contribution is

$$h(c) := \begin{cases} q_{\alpha_1} & \text{if } \ell = 1, \\ 1 - (1 - q_{\alpha_1})(1 - q_{\alpha_2}) & \text{if } \ell = 2, \\ H(\alpha_1, \alpha_2, \alpha_3) & \text{if } \ell = 3. \end{cases}$$

The true contribution is

$$o(c) := 1\{\text{the clause } C \text{ is satisfied by } \mathbf{x}^*\}$$

Number of atom levels: Let N_c be the number of typed atom c from $\Phi_{\leq 3}$ with the optimal assignment:

$$\mathbf{x}^* \in \arg\max_{\mathbf{x}} \sum_{i=1}^{m} C_i(\mathbf{x}), \qquad N_c = \#\{C \in \Phi_{\leq 3} : \mathrm{type}(C, \mathbf{x}^*) = c\}.$$

The **lower bound** holds if

$$\underbrace{\sum_c N_c \cdot h(c)}_{L_{\leq 3}} - \underbrace{\rho \cdot \sum_c N_c \cdot o(c)}_{\rho \cdot \mathrm{OPT}(\Phi_{\leq 3})} \geq 0$$

Observation: It's too strong to guarantee $\forall c, \quad h(c) - \rho \cdot o(c) \geq 0$.

Idea: Allow bad atom c with

$$h(c) - \rho \cdot o(c) < 0,$$

but their mixtures are constrained. Namely, constrain N_c by the bias conditions.

Solution: choose the parameters of $L_{\leq 3}$ via a well designed **Linear Programming**.

Bias constraints

Bias constraints: Fix a bucket $I_i = [a_i, b_i]$ and an assignment $\beta \in \{0,1\}$. Each variable x_v with $b_v \in I_i$ satisfies $b_i \cdot (p_v + q_v) - (p_v - q_v) \leq 0$ and $(p_v - q_v) - a_i \cdot (p_v + q_v) \leq 0$. Thus

$$G^{\mathsf{lo}}_{i,\beta} := b_i \cdot \sum_{v: b_v \in I_i, x^*_v = \beta} (p_v + q_v) - \sum_{v: b_v \in I_i, x^*_v = \beta} (p_v - q_v) \leq 0,$$
$$G^{\mathsf{up}}_{i,\beta} := \sum_{v: b_v \in I_i, x^*_v = \beta} (p_v - q_v) - a_i \cdot \sum_{v: b_v \in I_i, x^*_v = \beta} (p_v + q_v) \leq 0$$

Bias constraints

Atom-level bias constraints: For each typed atom $c = (\ell; s_1, \ldots, s_\ell; i_1, \ldots, i_\ell; \tau_1, \ldots, \tau_\ell)$, define the atom-level bias constraint

$$g^{\mathsf{lo}}_{i,\beta}(c) := \sum_{j : j \in [\ell], i_j = i, \tau_j = \beta} b_i \cdot \omega_\ell - s_j \cdot \omega_\ell,$$

$$g^{\mathsf{up}}_{i,\beta}(c) := \sum_{j : j \in [\ell], i_j = i, \tau_j = \beta} s_j \cdot \omega_\ell - a_i \cdot \omega_\ell.$$

Take the summation over all possible typed atoms c, we have

$$\sum_c N_c \cdot g^{\mathsf{lo}}_{i,\beta}(c) = G^{\mathsf{lo}}_{i,\beta} \leq 0,$$

$$\sum_c N_c \cdot g^{\mathsf{up}}_{i,\beta}(c) = G^{\mathsf{up}}_{i,\beta} \leq 0.$$

Linear Programming

For $\{q_\alpha\}$: We find good $I_1, \ldots, I_N$ and $r_1, \ldots, r_N$ via numerical experiment. Then $\{q_\alpha\}$ are determined.

Variables: c_0, u_j p_{jt} , $\lambda^{\text{lo}}_{i,\tau}$, $\lambda^{\text{up}}_{i,\tau}$.

Maximize ρ s.t.

1. $H(\alpha_1, \alpha_2, \alpha_3) \le 1 - (1 - q_{\alpha_1})(1 - q_{\alpha_2})(1 - q_{\alpha_3})$, for all the $\alpha_1, \alpha_2, \alpha_3$;
2. $h(c) - \rho \cdot o(c) + \sum_{i,\tau} \lambda^{\text{lo}}_{i,\tau} \cdot g^{\text{lo}}_{i,\tau}(c) + \sum_{i,\tau} \lambda^{\text{up}}_{i,\tau} \cdot g^{\text{up}}_{i,\tau}(c) \ge 0$ for all typed atom c;
3. $\lambda^{\text{lo}}_{i,\tau} \ge 0$, $\lambda^{\text{up}}_{i,\tau} \ge 0$ for all i, τ.

Remark: The number of variables and constraints are constants.

Lower bound

Lower bound: Take the summation over typed atoms on Constraint-2, we have

$$\underbrace{\sum_c N_c \cdot h(c)}_{L_{\leq 3}} - \underbrace{\rho \cdot \sum_c N_c \cdot o(c)}_{\rho \cdot \mathrm{OPT}(\Phi_{\leq 3})} + \sum_{i,\tau} \lambda^{\mathsf{lo}}_{i,\tau} \cdot G^{\mathsf{lo}}_{i,\tau} + \sum_{i,\tau} \lambda^{\mathsf{up}}_{i,\tau} \cdot G^{\mathsf{up}}_{i,\tau} \geq 0.$$

Because $G^{\mathsf{lo}}_{i,\tau} \leq 0$, $G^{\mathsf{up}}_{i,\tau} \leq 0$, $\lambda^{\mathsf{lo}}_{i,\tau} \geq 0$ and $\lambda^{\mathsf{up}}_{i,\tau} \geq 0$, we have

$$L_{\leq 3} - \rho \cdot \mathrm{OPT}(\Phi_{\leq 3}) \geq 0$$

By choosing proper r_i, we have

$$1 - \prod_{j=1}^{4}(1 - q_{\alpha_j}) \geq \rho.$$

Thus

$$\mathrm{OPT}(\Phi) \geq L_{\leq 3} + \rho \cdot M_{\geq 4} \geq \rho \cdot \mathrm{OPT}(\Phi)$$

.

What remains

Recall the snapshot score

$$L_{\leq 3} = \sum_{\alpha} q_\alpha \cdot U_{1,1,\alpha} + \sum_{\alpha,\beta} \Big(1 - (1-q_\alpha)(1-q_\beta)\Big) \cdot P_{2,(1,2),(\alpha,\beta)} + c_0 \cdot M_3 + \sum_{j=1}^{3} \sum_{\alpha} u_j(\alpha) \cdot U_{3,j,\alpha} + \sum_{1\leq j<t\leq 3} \sum_{\alpha,\beta} p_{jt}(\alpha,\beta) \cdot P_{3,(j,t),(\alpha,\beta)},$$

where $\{q_\alpha\}$, c_0, $\{u_j\}$ and $\{p_{jt}\}$ have been determined.

Fact: M_3 and $M_{\geq 4}$ can be counted exactly by $O(\log m)$ bits.

Part III: Quantum Sketches for Snapshot

Goal: estimate $U_{\ell,j,\alpha}$ and $P_{\ell,(j,t),(\alpha,\beta)}$ by $\mathrm{polylog}(n)$ qubits, with additive error at most $\varepsilon \cdot m$ and success probability at least 0.9.

Core idea: take $U_{\ell,j,\alpha}$, $\alpha = (s,i)$ for example. Sample each clause independently with probability σ. If the sampled clause contributes to $U_{\ell,j,(s,i)}$, update

$$U_{\ell,j,(s,i)} \leftarrow U_{\ell,j,(s,i)} + \frac{1}{\sigma}.$$

Challenge: The obstacle is checking the bias constraint, while the bias

$$b_v = \frac{p_v - q_v}{p_v + q_v}$$

can only be calculated after knowing the degrees, which is global information.

Instead of

counting via classical register,

we apply

threshold query via quantum register.

Strategy:

1. Partition the degree of (C, j) into prefix one and suffix one.
2. A quantum sketch will check whether the prefix degree lies in a specific level.
3. If it does, end the quantum sketch and count the suffix degree exactly.

Definition (Prefix and suffix degree)

Define the total degree

$$d_v := p_v + q_v.$$

Upon the arrival the of a clause C, partition the Φ into two parts.

- Define $d_v^{\leq C}$ ($d_v^{>C}$) as the prefix (suffix) total degree;
- Define $p_v^{\leq C}$ ($p_v^{>C}$) as the prefix (suffix) positive degree.

Degree level: For $a = 1, 2, \ldots, A$,

$$D_0 = 0, D_a := (1 + \varepsilon^3)^a, \text{ and } D_A \geq 4m,$$

where $A = O(\log m)$.

Decomposition via degree level: let (C, j) be the variable of the j-th literal of C. Let

$$U_{\ell,j,\alpha}^{(a)} := \#\{C : |C| = \ell, \alpha(C, j) = \alpha, d_{(C,j)} \in (D_a, D_{a+1}]\},$$

Therefore $U_{\ell,j,\alpha} = \sum_a U_{\ell,j,\alpha}^{(a)}$.

Choose a large enough constant $\kappa = \mathrm{poly}(1/\varepsilon)$. we focus on the **high degree** case

$$D_a \geq \kappa/2.$$

The low degree case $D_a < \kappa/2$ follows by the similar argument.

Approximate $d_v^{\leq C}$: Define

$$\widetilde{d_v^{\leq C}} := D_a$$

as the approximation of $d_v^{\leq C} \in (D_a, D_{a+1}]$.

Approximate $p_v^{\leq C}$: Let $s_v^{\leq C} := \mathsf{Bin}(p_v^{\leq C}, \frac{\kappa}{2D_a})$. Define

$$\widetilde{p_v^{\leq C}} := \frac{2D_a}{\kappa} \cdot s_{v^{\leq C}}$$

as the approximation of $p_v^{\leq C}$.

Fact 1: $|\widetilde{d}_v^{\leq C} - d_v^{\leq C}| \leq D_{a+1} - D_a \leq O(\varepsilon^3 \cdot d_v^{\leq C})$

Fact 2: $\mathrm{E}[s_v^{\leq C}] = p_v^{\leq C} \cdot \frac{\kappa}{2D_a} \leq D_{a+1} \cdot \frac{\kappa}{2D_a} \leq (1+\varepsilon^3)/2 \cdot \kappa$. And $s_v^{\leq C} \leq \kappa$ w.h.p.

Fact 3: By Chernoff's bound, w.h.p. $|s_v^{\leq C} - p_v^{\leq C} \cdot \frac{\kappa}{2D_a}| \leq O(\varepsilon^3 \kappa)$. Thus

$$|\widetilde{p}_v^{\leq C} - p_v^{\leq C}| \leq O(\varepsilon^3 \cdot D_a) \leq O(\varepsilon^3 \cdot d_v^{\leq C}).$$

Fact: Replacing $(d_v^{\leq C}, p_v^{\leq C})$ with $(\widetilde{d}_v^{\leq C}, \widetilde{p}_v^{\leq C})$ costs $O(\varepsilon \cdot m)$ error w.h.p. (Actually, we apply a ϵ-Noise on $\widetilde{b}_v$ and employ a ε-smoothed LP)

Quantum sketch framework:

1. When a clause C arrives, update the prefix degree information.
2. Use threshold queries to determine whether the prefix degree lies in a specific level.

Lower-q-sketch: for all $q = 0, 1, 2, \ldots, \kappa$, check

$$1\left\{d_v^{\leq C} > D_a\right\} \text{ and } 1\left\{s_v^{\leq C} \geq q\right\}$$

.

Upper-q-sketch: for all $q = 0, 1, 2, \ldots, \kappa$, check

$$1\left\{d_v^{\leq C} > D_{a+1}\right\} \text{ and } 1\left\{s_v^{\leq C} \geq q\right\}$$

.

Quantum register: Let T be the set of alive basic states.

$$|Q(T)\rangle = \frac{|\perp\rangle + \sum_{x \in T} |x\rangle}{\sqrt{1 + |T|}}.$$

Basic states: Let $m' = C_\varepsilon \cdot m$ with large enough C_ε.

1. Spare states: $|S_1\rangle, \ldots, |S_{m'}\rangle$; $\{|G_H(v,t)\rangle\}_{v \in [n], t \in [m']}$; $\{|G_E(v,t)\rangle\}_{v \in [n], t \in [m']}$
2. H-states: $\{|H(v,t)\rangle\}_{v \in [n], t \in [m']}$.
3. E-states: $\{|E(v,t)\rangle\}_{v \in [n], t \in [m']}$.

Initiate T by $\{|S_1\rangle, \ldots, |S_{m'}\rangle\}$.

Pointer: Let p be a pointer on $\{|S_1\rangle, \ldots, |S_{m'}\rangle\}$. Initiate it by $p \leftarrow 1$.

Basic operation: define inc$(H, u, 1)$ as

1. Apply the permutation

$$|S_p\rangle \to |H(u,1)\rangle \to |H(u,2)\rangle \to \cdots \to |H(u,M)\rangle \to$$
$$\to |G_H(u,M)\rangle \to \cdots \to |G_H(u,1)\rangle \to |S_p\rangle .$$

2. Increase the pointer

$$p \leftarrow (p+1) \mod m'$$

.

Let $\mathrm{inc}(H, u, r) := \underbrace{\mathrm{inc}(H, u, 1) \circ \cdots \circ \mathrm{inc}(H, u, 1)}_{r \text{ times}}$. Define $\mathrm{inc}(E, u, r)$ in the same way.

Lower-q-sketch

Update: For all the variables v in C,

1. inc$(H, v, 1)$;
2. inc$(E, v, 1)$;
3. If the literal of v is positive, with probability $\kappa/(2D_a)$, inc(E, v, D_{a+1}).

Lower-q-sketch

Query: Let $u = (C, j)$. If the clause length ℓ and sign s are saturated, measure $|Q(T)\rangle$ by

$$|\Phi^+\rangle = \frac{|H(u, D_a + 1)\rangle + |E(u, qD_{a+1})\rangle}{\sqrt{2}}, \quad |\Phi^-\rangle = \frac{|H(u, D_a + 1)\rangle - |E(u, qD_{a+1})\rangle}{\sqrt{2}}$$

- Let $\chi_C^{\mathsf{lo}} = 1$ if $|\Phi^+\rangle$; $\chi_C^{\mathsf{lo}} = -1$ if $|\Phi^-\rangle$; $\chi_C^{\mathsf{lo}} = 0$ otherwise.
- If $\chi_C^{\mathsf{lo}} = 0$, continue with the next clause.
- If $\chi_C^{\mathsf{lo}} \in \{-1, 1\}$, end the quantum sketch and count $d_u^{>C}$, $p_u^{>C}$ exactly.

Lower-q-sketch

Output: Let

$$F_q(C) := 1\left[\mathsf{clip}_{[-1,1]}\left(\underbrace{\frac{2\left(\frac{2D_a}{\kappa}\cdot q + p_v^{>C}\right)}{D_a + d_v^{>C}} - 1}_{\widetilde{b}_v} + \underbrace{\mathsf{Uni}[-\varepsilon,\varepsilon]}_{\varepsilon-\text{Noise}}\right) \in I_i\right].$$

And the Lower-q-sketch outputs

$$L_q(C) := \underbrace{\frac{m'}{2}}_{\text{contribution}} \cdot \underbrace{\chi_C^{\mathsf{lo}} \cdot (F_q(C) - F_{q-1}(C))}_{\text{threshold indicator}}.$$

Combine the lower and upper sketch:

$$\mathrm{E}[L_q(C) - U_q(C)] = \mathsf{1}[\ell, s; d_v^{\leq C} \in (D_a, D_{a+1}]] \cdot (F_q - F_{q-1}) \cdot \mathsf{1}[s_v^{\leq C} \geq D_{a+1}].$$

Take the summation

$$\widetilde{U}^{(a)}_{\ell,j,\alpha}(C) := \sum_{q=0}^{\kappa}(L_q(C) - U_q(C)).$$

Then

$$\mathrm{E}[\widetilde{U}^{(a)}_{\ell,j,\alpha}(C)] = 1\{\ell, s; d_v^{\leq C} \in (D_a, D_{a+1}]\} \cdot \underbrace{\left(F_0(C) + \sum_{q=1}^{\kappa}(F_q(C) - F_{q-1}(C))1\{s_v^{\leq C} \geq q\}\right)}_{= 1\{b_{(C,j)} \in I_i\}}$$

Estimation for $U_{\ell,j,(s,i)}$:

$$\widetilde{U}_{\ell,j,(s,i)} = \sum_{C} \sum_{a=0}^{A} \widetilde{U}^{(a)}_{\ell,j,(s,i)}(C).$$

Expectation: Because

$$\mathrm{E}[\widetilde{U}^{(a)}_{\ell,j,(s,i)}(C)] = \mathbf{1}\{\ell, s; d_v^{\leq C} \in (D_a, D_{a+1}]; b_{(C,j)} \in I_i\},$$

we have

$$\mathrm{E}[\widetilde{U}_{\ell,j,(s,i)}] = U_{\ell,j,(s,i)}.$$

Variance: For each $\mathrm{Var}\{\widetilde{U}^{(a)}_{\ell,j,\alpha}\} = O(m^2)$. Hence for $\widetilde{U}_{\ell,j,\alpha} = \sum_a \widetilde{U}^{(a)}_{\ell,j,\alpha}$, we have

$$\mathrm{Var}\{\widetilde{U}_{\ell,j,\alpha}\} = O(m^2 \log m).$$

We apply parallel repetition $O(\log(m))$ times to reduce it to $O(m^2)$.

Expectation $\&$ Variation $\&$ Chebyshev's inequality:

$$|\widetilde{U}_{\ell,j,\alpha} - U_{\ell,j,\alpha}| \leq \varepsilon \cdot m$$

with probability 0.9.

Space complexity:

1. For $\widetilde{U}^{(a)}_{\ell,j,(s,i)}$: for each $0 \le q \le \kappa = \mathrm{poly}(1/\varepsilon)$, Lower-$q$-Sketch or Upper-$q$-Sketch costs $O(\log(mn))$ qubits and $O(\log m)$ classical bits.
2. For $\widetilde{U}_{\ell,j,(s,i)}$: since $0 \le a \le A = O(\log m)$, there are $O(\log m)$ degree levels.
3. Multiplied by $O(\log m)$ parallel repetitions, the total space complexity is $\mathrm{polylog}(n)$.

Open Problem

1. How to classify the quantum advantage for more general CSPs, like $k \geq 3$ or large alphabet?
2. Does multipass streaming model admit quantum advantage?

Thanks for listening!