Computational geometry: notes 2*

Computing volumes: Metropolis chains

Our goal is to sample from a distribution with probability distribution whose density proportional to f, where f will be a function related to the convex body C. We shall introduce the rectangular lattice $\delta\mathbb{Z}^d$, and shall sample a point x from the grid with probability proportional to $f(x)$. Then we will "smudge" the point x, by choosing a random point from the cube of size δ centered at x. The parameter δ will be chosen small enough so we introduce only little error doing so, but as small as to make the number of lattice points too large.

The random walk we will employ will be a lazy random walk on points of $\delta\mathbb{Z}^d$, which we now define. To avoid the technical complications with the infinite state space, we restrict the random to $\delta[-N,N]^d$ for some large N (to be specified later). At each step with probability 1/2, the state of the walk does not change (this ensures laziness). With remaining probability 1/2, one chooses one of $2d$ coordinate directions uniformly. If the walk is currently at the point $x \in \delta\mathbb{Z}^d$, and the chosen direction is v, then the walk attempts to move to $y = x + \delta v$. If $y \notin [-N,N]^d$, then the attempt fails. Otherwise, the attempt succeeds with probability

$$\min\left(1, \frac{f(y)}{f(x)}\right). \tag{2}$$

Therefore, the transition probabilities of the random walk are

$$P(x,y) = \begin{cases} 0 & \text{if } xandy \text{ are not adjacent,} \\ \frac{1}{4d}\min\left(1, \frac{f(y)}{f(x)}\right) & \text{if } x \text{ and } y \text{ are adjacent} \\ 1 - \sum_{y'} P(x,y') & \text{if } x = y. \end{cases}$$

The chain is reversible. Indeed, for any constant c the expression $\big(cf(x)\big)P(x,y) = (c/4d)\min(f(x), f(y))$ is symmetric in x and y. Therefore, if c is chosen so that $cf(x)$ is probability distribution, then $cf(x)$ is the stationary distribution of the random walk. Note that we do need to know only f, but not the normalizing constant c to perform this random walk.

A chain that has transition probabilities (2) is called a Metropolis chain. It is powerful tool that allows us to sample according to an arbitrary distribution f.

*These notes are from `http://www.borisbukh.org/CompGeomEaster11/notes2.pdf`.

Computing volumes: sampling from log-concave distributions

A function $f\colon \mathbb{R}^d \to \mathbb{R}_+$ is called *log-concave* if $F = \log f$ is concave, i.e. if

$$F\big(\lambda x + (1-\lambda)\big) \geq \lambda F(x) + (1-\lambda)F(y) \qquad \text{for all } x, y \in \mathbb{R}^d \text{ and } 0 \leq \lambda \leq 1.$$

Arithmetic-geometric means inequality immediately implies that positive concave functions are log-concave.

The log-concave functions are the cousins of convex sets, and many theorems that holds for convex sets have an analogue for log-concave functions. There is a formal connection as well. If C is convex, then its characteristic function is log-concave (with the convention that $\log 0 = -\infty$). Conversely, if f is log-concave then the set $\{x : f(x) \geq 1\}$ is convex.

A function F is called *α-Lipschitz* if

$$|F(x) - F(y)| \leq \alpha \|x - y\|.$$

We shall assume that $F = \log f$ is α-Lipschitz. This condition eliminates the problems with the boundary of a convex body that we alluded to in the previous notes. We can already see the relevance of the Lipschitz condition from (2): If the ratio $f(y)/f(x)$ is small, then the random walk will move slowly.

Computing volumes: an isoperimetric inequality

Originally the term "isoperimetric inequality" referred to the statement that among all bodies of the same perimeter the disk has the smallest area. Gradually the term came to describe not only the obvious extension to the volume and surface area of bodies in higher dimension, but also virtually any inequality between geometric quantities that quantify "size", in any sense.

We shall use the following isoperimetric inequality relating the diameter of the convex body to the size of the cuts that it admits. The *diameter* of a set C is the maximum distance between points of C:

$$\operatorname{diam} C = \sup\{\operatorname{dist}(x-y) : x, y \in C\}.$$

Associate to the log-concave function f a measure μ defined by $\mu(S) = \int_S f(x)\,dx$.

Theorem 10 (Applegate–Kannan, Dyer–Frieze)**.** *Suppose C is a convex body, and μ is measure associated to a continuous log-concave function f. Let $C = S_1 \cup S_2 \cup B$ be any partition of C into three sets. Suppose $D \geq \operatorname{diam} C$ and $T \leq \operatorname{dist}(S_1, S_2)$.*

$$\min\big(\mu(S_1), \mu(S_2)\big) \leq \tfrac{1}{2} \cdot \frac{D}{T}\mu(B). \tag{3}$$

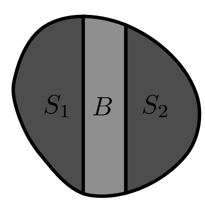

Cut B is large.

The reduction will preserve D and T, but might change $\operatorname{diam} C$ and $\operatorname{dist}(S_1, S_2)$. We call $\frac{1}{2}\frac{D}{T}$ the *isoperimetric constant* of the inequality. Note that it suffices to prove that (3) holds for every isoperimetric constant greater than $\frac{1}{2}\frac{D}{T}$. The isoperimetric constant is a function of D and T.

The idea of the proof is that it is possible to cut C by a hyperplane, as to reduce (3) to two similar inequalities. Via repeated cutting we shall replace C by a convex body that will be "needle-like", i.e. long in one direction and very narrow in the other directions. That case will be treated (up to a small error term) as a 1-dimensional problem.

To do the cutting we need the following.

Lemma 11 (Ham-sandwich theorem)**.** *Let $\mu_1, \cdots, \mu_d$ be finite continuous measures on $\mathbb{R}^d$. Then there exists a hyperplane H such that the halfspace H^+ and H^- satisfy*

$$\mu_i(H^+) = \mu_i(H^-), \qquad \textit{for all } i = 1, \ldots, d.$$

Corollary 12. *Let μ be a finite continuous measure on $\mathbb{R}^d$ and $S_1, \ldots, S_d$ be measurable. Then there exists a hyperplane H such that the halfspace H^+ and H^- satisfy*

$$\mu(H^+ \cap S_i) = \mu(H^- \cap S_i), \qquad \textit{for all } i = 1, \ldots, d.$$

The ham-sandwich theorem is a consequence of Borsuk–Ulam theorem, which properly belong to an algebraic topology course, and will not be proved here. The case $d = 2$, that we shall use, can be proved directly via the intermediate value theorem.

Fix a direction v (by "direction" we always mean a unit vector in $\mathbb{R}^d$). We denote the cross-section of a body C by hyperplane normal to v by

$$C(s) \stackrel{\text{def}}{=} C \cap \{x : \langle x, v \rangle = s\}.$$

A body is *ε-needle-like* in direction v if $\operatorname{diam} C(s) \leq \varepsilon$ for all s. A partition $C = S_1 \cup S_2 \cup B$ is a *straight* in direction v if

$$S_i = \left\{ \bigcup C(s) : S_i \cap C(s) \neq \emptyset \right\} \qquad \text{for } i = 1, 2.$$

Furthermore, the partition is called *fully straight* if S_1 and S_2 are intersections of C with a halfspace.

Recall that the *width* of the body C in direction v is

$$\operatorname{wd}(C; v) = \sup\{\langle x - y, v \rangle : x, y \in C\}.$$

Lemma 13. *There is a constant $c_d > 0$ such that for every convex body C there is a direction v in which the width is $\operatorname{wd}(C; v) \leq c_d \operatorname{vol}(C)^{1/d}$.*

Proof. From the ellipsoid method, we know that for every C there is an ellipsoid E containing C, and satisfying $\operatorname{vol}(E) \leq (8d\sqrt{d})^d \operatorname{vol}(C)$. Since $\operatorname{vol}(E) = 2^{-d} \operatorname{vol} B(0,1) \prod_{i=1}^{d} \operatorname{wd}(E; v_1)$ where v_i are the axes of the ellipsoid, the lemma follows. $\square$

Lemma 14. *To prove Theorem 10 with isoperimetric constant $r(D,T)$ it suffices to prove it for ε-needle-like bodies with the same isoperimetric constant $r(D,T;\varepsilon)$, where $r(\varepsilon) \to r$ as $\varepsilon \to 0$.*

Proof. Let M be sufficiently large so that $f(x) \geq 1/M$ for all x. Suppose theorem 10 fails, and let $C = S_1 \cup S_2 \cup B$ be a counterexample. Choose ε to be sufficiently small that

$$\min\bigl(\mu(S_1), \mu(S_2)\bigr) > r(D,T;\varepsilon)\mu(B).$$

Let $v_1, \dots, v_j$ be j orthonormal directions such that

$$\operatorname{wd}(C; v_i) \leq \varepsilon/\sqrt{d-1} \qquad \text{for each } i = 1, \dots, j.$$

Suppose furthermore that $C = S_1 \cup S_2 \cup B$ is a counterexample with the largest value of j among all counterexamples. If $j \geq d-1$, then C is ε-needle-like with respect to the direction orthogonal to $v_1, \dots, v_{d-1}$. That contradicts the assumption of the lemma. Thus, $j \leq d-2$. Since $\dim\bigl(\operatorname{span}(v_1, \dots, v_j)^\perp\bigr) \geq 2$, by Corollary 12 there is a hyperplane H with normal direction $v \in \operatorname{span}(v_1, \dots, v_j)^\perp$ such that $\mu(H^+ \cap S_i) = \mu(H^- \cap S_i)$ for $i = 1, 2$. Without loss of generality $\mu(H^+ \cap B) \leq \frac{1}{2}\mu(B)$.

Let $C' = C \cap H^+$, $S_1' = S_1 \cap H^+$, $S_2' = S_2' \cap H^+$ and $B' = B \cap H^+$. Since the partition $C = S_1 \cup S_2 \cup B$ is a counterexample to the theorem 10, so is the partition $C' = S_1' \cap S_2' \cup B'$ because $\operatorname{diam} C' \leq \operatorname{diam} C$ and $\operatorname{dist}(S_1', S_2') \geq \operatorname{dist}(S_1, S_2)$. Note that $\mu(C') \leq (1/2)\mu(C)$. We can iterate the bisection, and obtain an infinite sequence of convex bodies

$$C = C_0 \supset C_1 \supset \cdots \supset C_m \supset$$

Since $\mu(C_m) \leq 2^{-m}\mu(C)$, and the function f satisfies $f(x) \geq 1/M$, it follows that $\operatorname{vol}(C_m) \to 0$.

Let $\pi \colon \mathbb{R}^d \to \operatorname{span}(v_1, \dots, v_j)$ be the projection. If $\operatorname{vol}(\pi(C_m); v) \to 0$, then by Lemma 13 there $\operatorname{wd}(\pi(C_m)) \to 0$, which contradicts minimality of j. Otherwise, if we let $\bar{C} = \bigcap_m C_m$, it follows from $\operatorname{vol}(\bar{C}) = 0$ that all the fibers are one-element sets. As, the fibers of C agree with those of $\bar{C}$, we reach a contradiction. $\square$

Our next step is to reduce to the case of a fully straight partition. We shall do it in two steps. First, we reduce it to a straight partition, and then to a fully straight partition.

Lemma 15. *If Theorem 10 is true for straight partitions of ε-needle-like bodies with the isoperimetric constant $r_s(D,T;\varepsilon)$, then it holds for an arbitrary partition of ε-needle-like bodies with the isoperimetric constant $r(D,T;\varepsilon) = r_s(D, T-2\varepsilon; \varepsilon)$.*

Proof. Let C be ε-needle-like. Let v be the "long" direction of C. Let

$$\hat{S}_i = \left\{ \bigcup C(s) : S_i \cap C(s) \neq \emptyset \right\}.$$

Then $\operatorname{dist}(\hat{S}_1, \hat{S}_2) \geq \operatorname{dist}(\hat{S}_1, \hat{S}_2) - 2\varepsilon$, as well as $\mu(S_i) \leq \hat{S}_i$ and $\mu(\hat{B}) \geq \mu(B)$. Hence, it follows that

$$\min\bigl(\mu(S_1), \mu(S_2)\bigr) \leq \min\bigl(\mu(\hat{S}_1), \mu(\hat{S}_2)\bigr) \leq r_{\mathrm{s}}(D, T-2\varepsilon; \varepsilon)\mu(\hat{B}) \leq r_{\mathrm{s}}(D, T-2\varepsilon; \varepsilon)\mu(B),$$

and we are done in the light of the previous lemma. $\square$

Lemma 16. *If Theorem 10 holds for fully straight partitions ε-needle-like bodies with the isoperimetric constant $r_{fs}(D, T; \varepsilon)$, then it holds for straight partitions of ε-needle-like bodies with the same constant.*

Proof. Fix D and T and write $r = r_{\mathrm{fs}}(D, T; \varepsilon)$. Let $C = S_1 \cup S_2 \cup B$ be a straight partition. Set $T_i = \{s \in \mathbb{R} : C(s) \cap S_i \neq \emptyset\}$. We may assume that the connected components of T_1 and T_2 are alternating intervals, for otherwise we could increase either $\mu(S_1)$ or $\mu(S_2)$. Let the connected components be $[a_1, b_1], [a_2, b_2], \ldots, [a_n, b_n]$ in that order. For brevity we also write $I_m = [a_m, b_m]$. The complement of the union of these intervals is consists of intervals $B_1, \ldots, B_{n-1}$ in that order. See the picture below.

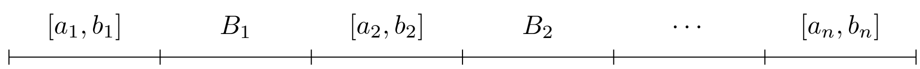

Without loss of generality odd-numbered intervals are in T_1 and even-numbered are in T_2.

We shall abuse the notation and write $\mu(I_m)$ for $\mu\bigl(I_m \times \mathbb{R}^{d-1}) \cap C\bigr)$, and similarly for $\mu(B_m)$. The assumption of the lemma tells us that for each m we have

$$\min(\mu(I_m), \mu(I_{m+1})) \leq r\mu(B_m).$$

Our goal is to show that $\min(\mu(T_1), \mu(T_2)) \leq r\mu(B)$.

Suppose for some even m we have $\mu(I_m) \geq \mu([a_1, a_m])$ and $\mu(I_m) \geq \mu([b_m, b_n])$. Then

$$\mu(S_1) \leq \mu([a_1, a_m]) + \mu([b_m, b_n]) \leq r\mu(B_{m-1}) + r\mu(B_m) \leq r\mu(B),$$

as desired. Otherwise, for each even m either $\mu(I_m) \leq \mu([a_1, a_m])$ or $\mu(I_m) \leq \mu([b_m, b_n])$. In either case

$$\mu(I_m) \leq \max(r\mu(B_m), r\mu(B_{m+1}) \leq r\mu(B_m) + r\mu(B_{m+1}).$$

Summing over all even m we obtain $\mu(S_2) \leq r\mu(B)$. $\square$

Computing volumes: Brunn–Minkowski inequality and the conclusion of the proof

To dispose of the fully straight case, we need the classical Brunn–Minkowski inequality.

Theorem 17 (Brunn–Minkowski inequality)**.** *If A and B are compact and nonempty, then* $\operatorname{vol}(A+B)^{1/d} \geq \operatorname{vol}(A)^{1/d} + \operatorname{vol}(B)^{1/d}$.

Proof. A *brick set* is a union of finitely many axis-parallel boxes with disjoint interiors. First note that, it suffices to prove the inequality only for brick sets. Indeed, let $A_1, A_2, \ldots$ be a sequence of brick sets containing A, such that $A = \bigcap_k A_k$. Similarly define $B_1, B_2, \ldots$. If $x \in \bigcap_k (A_k + B_k)$ then $x = a_k + b_k$ for some $a_k \in A_k$ and $b_k \in B_k$. By passing to a subsequence, we may assume that $a_k \to a$ and $b_k \to b$. As $a \in A$ and $b \in B$, it follows that $x \in A + B$. Since x is arbitrary, $\bigcap_k (A_k + B_k) \subset A + B$. Since $\operatorname{vol} A_k \to \operatorname{vol} A$ and $\operatorname{vol} B_k \to \operatorname{vol} B$, Brunn–Minkowski indeed needs to be proved only for the brick sets.

The proof is by induction on the total number of bricks. To dispose of the case when A and B consist of a single brick each, we need to show the inequality $\prod x_i^{1/n} + \prod_i y_i^{1/n} \leq \prod_i (x_i + y_i)^{1/n}$ for all positive x_i, y_i. After division by $\prod_i x_i$ the inequality reduces to

$$1 + \prod_i x_i^{1/n} \geq \prod_i (1 + x_i)^{1/n}. \tag{4}$$

Since $(1 + \sqrt{x_i x_j})^2 \leq (1 + x_i)(1 + x_j)$, we might assume that all x's are equal. However, in that case (4) holds vacuously.

More interesting is the induction step. Suppose A consists of two bricks. Let H be a coordinate hyperplane such that there is one whole brick of A on each side of H. Without loss of generality $H = \{x_d = 0\}$. Let $A^- = A \cap \{x_d \leq 0\}$, $A^+ = A \cap \{x_d \geq 0\}$, and similarly for B^- and B^+. Translate the set B so that

$$\lambda \stackrel{\text{def}}{=} \frac{\operatorname{vol} A^-}{\operatorname{vol} A} = \frac{\operatorname{vol} B^-}{\operatorname{vol} B}$$

for some ratio λ. Note that translation of B does not affect the volume of $A+B$. Since $A \cap \{x_d \leq 0\}$ and $A \cap \{x_d \geq 0\}$ consist of fewer bricks than A, by the induction hypothesis

$$\begin{aligned}
\operatorname{vol}(A+B) &\geq \operatorname{vol}(A^- + B^-) + \operatorname{vol}(A^+ + B^+) \\
&\geq [\operatorname{vol}(A^-)^{1/d} + (B^-)^{1/d}]^d + [\operatorname{vol}(A^+)^{1/d} + (B^+)^{1/d}]^d \\
&= \lambda[\operatorname{vol}(A)^{1/d} + (B)^{1/d}]^d + (1-\lambda)[\operatorname{vol}(A)^{1/d} + (B)^{1/d}]^d \\
&= [\operatorname{vol}(A)^{1/d} + (B)^{1/d}]^d.
\end{aligned}$$

□

If we apply Brunn–Minkowski to sections of a single convex body by hyperplanes normal to a given vector v, we obtain Brunn's inequality.

Theorem 18 (Brunn's inequality)**.** *Suppose $C \subset \mathbb{R}^{d+1}$ is compact convex body, and v any direction. Then the function $s \mapsto (\operatorname{vol} C(s))^{1/d}$ is concave on its support.*

Proof. It suffices to deal with a special case. Let $A = C(0)$ and $B = C(1)$, and $\lambda \in (0,1)$ arbitrary. We wish to show that $\operatorname{vol} C(\lambda)^{1/d} \geq \operatorname{vol}(A)^{1/d} + (1 - \lambda)\operatorname{vol}(B)^{1/d}$. As $C(\lambda)$ contains $\lambda A + (1-\lambda)B$, it suffices to show

$$\operatorname{vol}(\lambda A + (1-\lambda)B)^{1/d} \geq \lambda \operatorname{vol}(A)^{1/d} + (1-\lambda)\operatorname{vol}(B)^{1/d}.$$

That follows from the Brunn–Minkowski inequality since $\operatorname{vol}(\lambda A) = \lambda^d \operatorname{vol}(A)$ and $\operatorname{vol}\bigl((1-\lambda)B\bigr) = (1-\lambda)^d \operatorname{vol}(B)$. $\square$

Proof of theorem 10. The function $\log f$ is concave on compact set, it is Lipschitz. Let M be the Lipschitz constant of $\log f$. We shall show that the theorem 10 holds for fully straight partitions of ε-needle-like bodies with the isoperimetric constant of $\frac{1}{2} \cdot e^{M\varepsilon} \frac{D}{T}$. Let C be an ε-needle-like body, and $C = S_1 \cup S_2 \cup B$ be a fully straight partition.

Let s_1 and s_2 be the largest numbers s.t. $K(-s_1)$ and $K(s_2)$ are non-empty. Let L be any line connecting a point of $K(s_1)$ and a point of $K(s_2)$. For a real number $-s_1 \le s \le s_2$ let $\bar f(s)$ for the value of f in the unique point of $L \cap K(s)$. By the Lipschitz condition, we have $e^{-\varepsilon M} \bar f(s) \le f(x) \le e^{\varepsilon M} \bar f(s)$ for any $x \in K(s)$.

Since the partition $C = S_1 \cup S_2 \cup B$ is fully straight, there are numbers u_1 and u_2 such that $S_1 = \bigcup_{s \le u_1} C(s)$ and $S_2 = \bigcup_{s \ge u_2} C(s)$. Let $v(s) = \operatorname{vol} C(s)$. By the mean value theorem,

$$\int_{u_1}^{u_2} \bar f(s) v(s) = (u_2 - u_1) \bar f(\zeta) v(\zeta)$$

for some $u_1 \le \zeta \le u_2$. Without loss of generality $\zeta = 0$. Since $u_2 - u_1 \ge t$, this implies that

$$\mu(B) \ge e^{-M\varepsilon} \bar f(0) v(0).$$

By Brunn–Minkowski inequality $v(s)^{1/(d-1)}$ is concave, and hence so is $\log v(s)$. Since $\ln f$ is log-concave, it follows that $G(s) \stackrel{\text{def}}{=} \ln f(s) + \ln v(s)$ is concave. By scaling orthogonal to v we may assume that $G(0) = 0$. Since G is concave, there is a $\gamma \in \mathbb{R}$ such that $G(s) \le \gamma s$ for all s.

If $\gamma = 0$, then $\mu(S_1) \le e^{M\varepsilon} s_1$ and $\mu(S_2) \le e^{M\varepsilon} s_2$. Hence

$$\min\bigl(\mu(S_1), \mu(S_2)\bigr) \le \tfrac{1}{2}\bigl(\mu(S_1) + \mu(S_2)\bigr) \le \tfrac{1}{2} e^{M\varepsilon} D \le \tfrac{1}{2} e^{2M\varepsilon} \frac{D}{T}.$$

Suppose $\gamma \ne 0$. We may assume that $\gamma > 0$ by reflection of C if necessary. We can also assume that $\gamma = 1$ by scaling in v direction. Thus, $G(s) \le s$, giving

$$\mu(S_1) \le e^{M\varepsilon} \int_{-s_1}^{0} e^s \, ds = e^{M\varepsilon}(1 - e^{-s_1}),$$
$$\mu(S_2) \le e^{M\varepsilon} \int_{0}^{s_2} e^s \, dx = e^{M\varepsilon}(e^{s_2} - 1).$$

Let $\tilde\mu = e^{-M\varepsilon} \min\bigl(\mu(S_1), \mu(S_2)\bigr)$. We have $s_1 \ge -\ln(1 - \tilde\mu)$ and $s_2 \ge \ln(1 + \tilde\mu)$. Thus,

$$\tfrac{1}{2} D \ge \tfrac{1}{2}(s_2 + s_1) \ge \tfrac{1}{2}\bigl(\ln(1 + \tilde\mu) + \ln(1 - \tilde\mu)\bigr) > \tilde\mu,$$

where we used the Taylor expansion of $\ln(1+x)$ to deduce the last inequality. $\square$

Computing volumes: the end

Recall that we may assume that convex body C is sandwiched between $B(0,1)$ and $B(0,8d\sqrt{d})$. For $x \neq 0$ let

$$r'(x) = \inf\{t > 0 : x \in tC\} \text{ and } r(x) = \max(0, r'(x) - 1).$$

The function $r(x)$ is convex because if $x \in r(x)C$ and $y \in r(y)C$, then $\lambda x + (1-\lambda)y \in \lambda r(x)C + (1-\lambda)r(y)C$. The function $r(x)$ is convex because it is a maximum of two convex functions. The smoothening of C that we will work with is the function

$$f(x) = \exp\bigl(-2d\, r(x)\bigr).$$

The function $f(x)$ is log-concave. Its integral is comparable to the volume of C. The inequality $\int f \geq \operatorname{vol} C$ is clear. In the other direction, by slicing $\mathbb{R}^d$ into level sets of $r(x)$ we obtain

$$\begin{aligned}
\int f - \operatorname{vol} C &\leq \int_0^\infty \Bigl[\operatorname{vol}\bigl((1+r+\,dr)K\bigr) - \operatorname{vol}\bigl((1+r)K\bigr)\Bigr] e^{-2dr} \\
&= d \operatorname{vol}(C) \int_0^\infty (1+r)^{d-1} e^{-2dr}\, dr \\
&\leq d \operatorname{vol}(C) \int_0^\infty e^{r(d-1)} e^{-2dr}\, dr \\
&= \operatorname{vol}(C)/e.
\end{aligned}$$

The function r' is 1-Lipschitz because $r'(x) \leq r'(y) + |x-y|$ because $B(0,1) \subset C$. Thus $r(x)$ is also 1-Lipschitz, and so $\log f(x)$ is $2d$-Lipschitz. We shall perform the Metropolis random walk (2) on $\delta\{-N,\dots,N\}^d$ with $N = 800\delta^{-1} d\sqrt{d}\log d$. We shall choose $\delta = 1/d^2$. We start the walk from the origin.

Lemma 19. *If $f(x)$ is log-concave, and $\log f(x)$ has Lipschitz constant α, then the conductance of the Metropolis random walk (2) on $\delta\{-N,\dots,N\}^d$ is at most $8d^{3/2}\exp(3\alpha\delta\sqrt{d})N$.*

Proof. By scaling f as necessary, we may assume that f is the stationary distribution of the Markov chain. Note that scaling does not change the Lipschitz constant of $\log f$.

Write $\Omega = \delta\{-N,\dots,N\}^d$ for the state space of the random walk. Let Q_δ denote the cube of side length δ centered at the origin. For every set $X \subset \Omega$ the notation $\bar{X}$ denotes the se $X + Q_\delta$. Note that the capacity $C_X = \sum_{x\in X} f(x)$ satisfies

$$e^{-\alpha\delta\sqrt{d}}\mu(\bar{X}) \leq C_X \leq e^{\alpha\delta\sqrt{d}}\mu(\bar{X}).$$

Suppose S is any subset of Ω of capacity $C_S \leq 1/2$. The flow from S to

$\Omega \setminus S$ is

$$F_S = \sum_{\substack{x \in S \\ y \notin S}} f(x)P(x,y) = \sum_{\substack{x \in S \\ y \notin S \\ x \sim y}} \frac{1}{4d} \min(f(x), f(y))$$
$$\geq \tfrac{1}{4d} e^{-\alpha\delta\sqrt{d}} \sum_{\substack{x \in S \\ y \notin S \\ x \sim y}} f(y).$$

Introduce the sets

$$B = \{x \in S : \exists y \in \Omega \setminus S \text{ s.t. } x \sim y\}.$$

Then

$$F_S \geq \tfrac{1}{4d} e^{-\alpha\delta\sqrt{d}} C_B.$$

Let $S' = \Omega \setminus (B \cup S)$. Since $C_S \leq 1/2$

$$\min\bigl(\mu(\bar{S}), \mu(\bar{S}')\bigr) \geq e^{-\alpha\delta\sqrt{d}} \min(C_S, C_{S'}) = e^{-\alpha\delta\sqrt{d}} \min(C_S, 1 - C_S - C_B)$$
$$\geq e^{-\alpha\delta\sqrt{d}} (C_S - C_B).$$

By the isoperimetric theorem 10 applied to $\delta[-N,N]^d$ with $T = \delta$ and $D = \delta\sqrt{d}N$ we have

$$e^{-\alpha\delta\sqrt{d}}(C_S - C_B) \leq \min\bigl(\mu(\bar{S}), \mu(\bar{S}')\bigr) \leq \tfrac{1}{2} N \mu(\bar{B}) \leq \tfrac{1}{2}\sqrt{d} N e^{\alpha\delta\sqrt{d}} C_B.$$

Therefore,

$$C_S \leq \bigl(1 + \tfrac{1}{2}\sqrt{d}Ne^{2\alpha\delta\sqrt{d}}\bigr)C_B \leq 4de^{2\alpha\delta\sqrt{d}}\bigl(1 + \tfrac{1}{2}\sqrt{d}Ne^{2\alpha\delta\sqrt{d}}\bigr)F_S \leq 8d^{3/2}\exp(3\alpha\delta\sqrt{d})NF_s.$$

□

With the choice of α, δ, N we made obtain that the conductance of the random walk is

$$\Phi \geq 1/8d^{3/2}\exp(3\alpha\delta\sqrt{d})N = \Omega(d^{-11/2}\log^{-1} d)$$

By Jerrum–Sinclair bound on the mixing of the random walk it follows that

$$\Delta'(x)^2\pi(x) \leq d_2(t) \leq d_2(0)(1 - \Phi^2)^t.$$

It remains to estimate $d_2(0)$. We will be wasteful. We have

$$d_2(0) = \sum_x \Delta_0(x)^2/\pi(x) = \sum_x \bigl(\mathbf{P}[X_t = x] - \pi(x)\bigr)^2/\pi(x)$$
$$\leq \sum_x \bigl(\mathbf{P}[X_0 = x]^2 + \pi(x)^2\bigr)/\pi(x)$$
$$= 1 + \sum_x \mathbf{P}[X_0 = x]^2/\pi(x) = 1 + 1/\pi(0)$$

Since $\log f$ is $2d$-Lipschitz and $\delta = 1/d\sqrt{d}$, it follows that $e^{-2} \le f(x)/f(y) \le e^2$ for $y \in x + Q_\delta$. Thus for every $x \in C$ we have

$$\pi(x) = \frac{f(x)}{\sum_y f(y)} \ge \frac{f(0)}{e^2 \delta^{-d} \int f} = d^{-O(d)}$$

Thus we choose the number of steps to be $t = \log\bigl(1/\epsilon\pi(0)^2\Phi^2\bigr) = O(d^{13}\log^4 d)$, then $\Delta'(x)^2 \le \epsilon$.

We have thus constructed a polynomial-time sampler from $C \cap \delta\mathbb{Z}^d$ with $\delta = 1/d\sqrt{d}$. From it we can easily obtain a sampler from C as follows. Pick a random point $x \in C \cap \delta\mathbb{Z}^d$, and than $y \in x + Q_\delta$. With probability $f(y)/e^2 f(x)$ output y, otherwise restart the random walk from scratch.