






Practical Volume Estimation of Zonotopes by a New Annealing Schedule for Cooling Convex Bodies

Apostolos Chalkis¹, Ioannis Z. Emiris^{1,2}, and Vissarion Fisikopoulos¹

¹ Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece

{achalkis,emiris,vfisikop}@di.uoa.gr

² Athena Research and Innovation Center, Maroussi, Greece

Abstract. We study the problem of estimating the volume of convex polytopes, focusing on zonotopes. Although a lot of effort is devoted to practical algorithms for polytopes given as an intersection of halfspaces, there is no such method for zonotopes. Our algorithm is based on Multiphase Monte Carlo (MMC) methods, and our main contributions include: (i) a new uniform sampler employing Billiard Walk for the first time in volume computation, (ii) a new simulated annealing generalizing existing MMC by making use of adaptive convex bodies which fit to the input, thus drastically reducing the number of phases. Extensive experiments on zonotopes show our algorithm requires sub-linear number of oracle calls in the dimension, while the best theoretical bound is cubic. Moreover, our algorithm can be easily generalized to any convex body. We offer an open-source, optimized C++ implementation, and analyze its performance. Our code tackles problems intractable so far, offering the first efficient algorithm for zonotopes which scales to high dimensions (e.g. one hundred dimensions in less than 1 h).

Keywords: Volume approximation · Zonotope · Simulated annealing · Billiard Walk · Mathematical software

1 Introduction

Volume computation is a fundamental problem with many applications. It is #P-hard for explicit polytopes [7, 11], and APX-hard [9] for convex bodies in the oracle model. Therefore, a significant effort has been devoted to randomized approximation algorithms, starting with the celebrated result in [8] with complexity $O^*(d^{23})$ oracle calls, where $O^*(\cdot)$ suppresses polylog factors and dependence on error parameters, and d is the dimension. Improved algorithms reduced the exponent to 5 [13] and further results [5, 14] reduced the exponent to 3. Current theoretical results consider either the general oracle model or polytopes given as an intersection of halfspaces (i.e. H-polytopes). Regarding implementations, the approach of [13] led to the first practical implementation in [10]

for high dimensions, followed by another practical implementation [6] based on [5, 14]. However, both implementations can handle only H-polytopes.

An important class of convex polytopes are zonotopes [15]. A zonotope is the Minkowski sum of k d -dimensional segments. Equivalently, given a matrix $G \in \mathbb{R}^{d \times k}$ a zonotope can be seen as the affine projection of the hypercube $[-1, 1]^k$ to \mathbb{R}^d using the matrix G , while the columns of G are the corresponding segments (or generators). Zonotopes are centrally symmetric and each of their faces are again zonotopes. We call the *order* of a zonotope P the ratio between the number of generators of P over the dimension. For a nice introduction to zonotopes we refer to [20].

Volume approximation for zonotopes is of special interest in several applications in smart grids [1], in autonomous driving [2] or human-robot collaboration [16]. The complexity of algorithms that work on zonotopes strongly depends on their order. Thus, to achieve efficient computations, a solution that is common in practice is to over-approximate P , as tight as possible, with a second zonotope P_{red} of smaller order, while $\text{vol}(P_{red})$ is given by, an easy to compute, closed formula. A good measure for the quality of the approximation is the ratio of fitness, $\rho = (\text{vol}(P_{red})/\text{vol}(P))^{1/d}$, which involves a volume computation problem [3]. Existing work (e.g. in [12]) uses exact - deterministic volume computation [11], and thus ρ can not be computed for $d > 10$ in certain applications.

A typical randomized algorithm uses a Multiphase Monte Carlo (MMC) technique, which reduces volume approximation of convex P to computing a telescoping product of ratios of integrals. Then each ratio is estimated by means of random walks sampling from a proper multivariate distribution. In this paper we rely on MMC of [13] which specifies a sequence of convex bodies $P_m \subseteq \dots \subseteq P_0 = P$, assuming P is well-rounded, i.e. $B_d \subseteq P \subseteq C\sqrt{d}B_d$, where C is constant and B_d is the unit ball. We define a sequence of scaled copies of B_d , and let $P_i = (2^{(m-i)/d}B_d) \cap P$, $i = 0, \dots, m$. One computes $\text{vol}(P_m)$ and applies:

$$\begin{aligned} \text{vol}(P) &= \text{vol}(P_m) \frac{\text{vol}(P_{m-1})}{\text{vol}(P_m)} \dots \frac{\text{vol}(P_0)}{\text{vol}(P_1)}, \\ m &= O(d \lg d), P_0 = C\sqrt{d}B_d \cap P. \end{aligned} \quad (1)$$

There is a closed-form expression to compute $\text{vol}(P_m) = \text{vol}(B_d)$. Each ratio $r_i = \text{vol}(P_{i+1})/\text{vol}(P_i)$ in Eq. (1) can be estimated within arbitrary small error ϵ_i by sampling uniformly distributed points in P_i and accept/reject points in P_{i+1} so $\text{vol}(P)$ can be derived after m multiplications. The estimation of r_i shows how sampling comes into the picture. In [10], assuming $rB_d \subseteq P \subseteq RB_d$ for $r < R$, they get $m = \lceil d \lg(R/r) \rceil$. The issue is to minimize m while each ratio remains bounded by a constant, and to use a random walk that converges, after a minimum number of steps, to the uniform distribution. The first would permit a larger approximation error per ratio without compromising overall error, while it would require a smaller uniform sample to estimate each ratio. The second would reduce the cost per sample point. Total complexity is determined by the number of ratios, or phases, multiplied by the number of points, or steps, to

estimate each ratio, multiplied by the cost to generate a point. The first two factors are determined by the MMC and the third by the random walk.

Previous Work. Exact volume computation for zonotopes can be reduced to a sum of absolute values of determinants, with an exponential number of summands in d [11]. Practical algorithms for volume computation of zonotopes are limited to low dimensions (typically ≤ 10 in [6]). This is due to two main reasons: current algorithms create a long sequence of phases in MMC for zonotopes, and the boundary and membership oracles are costlier than for H-polytope, as they both reduce to Linear Programs (LP). In [6], they consider low dimensional zonotopes: in \mathbb{R}^{10} with $k = 20$ generators, the algorithm performs 1.92×10^5 Boundary Oracle Calls (BOC), whereas our algorithm requires only 8.50×10^3 BOCs, and for $d = 100, k = 200$ it performs 6.51×10^4 BOCs (see Table 1). In [19] they present an implementation of an efficient algorithm that computes Minkowski sums of polytopes (generalization of zonotopes). In [18] they propose a randomized algorithm for enumerating the vertices of a zonotope.

Our Contribution. We focus on zonotopes and introduce crucial algorithmic innovations to overcome the existing barriers, by reducing significantly the number of oracle calls. Thus, our method scales to high dimensions ($d = 100$ in ≤ 1 h), performing computations which were intractable till now.

We use a new simulated annealing method in order to define a sequence of appropriate convex bodies, instead of balls, in MMC, and we exploit the fast convergence of Billiard Walk (BW) [17] to the uniform distribution. We experimentally analyze complexity by counting the number of BOCs, since BW uses boundary reflections.

The new simulated annealing specifies the P_i 's by exploiting the statistical properties of the telescoping ratios to drastically reduce the number of phases. In particular, we bound each ratio $r_i = \text{vol}(P_{i+1})/\text{vol}(P_i)$ to a given interval $[r, r + \delta]$ with high probability, for some real r . Moreover, our MMC generalizes balls, used in [13] and previous papers, by taking as input any convex body C and constructing the sequence by only scaling C . It does not need an enclosing body of P nor an inscribed ball (or body), unlike [10, 13].

Most of the previous algorithms use a rounding step before volume computation, as preprocessing, to reduce the number of phases in MMC. However, rounding requires uniform sampling from P which makes it costly for zonotopes because of the expensive oracle calls. Our approach is to exploit the fact that the schedule uses any body C and skip rounding by letting C be an H-polytope that fits well to P . The idea is to construct C fast and reduce the number of phases and the total runtime more than a rounding preprocessing would do in practice.

We prove that the number of bodies defined in MMC is, with high probability, $m = O(\lg(\text{vol}(P)/\text{vol}(P_m)))$, where $P_m = qC \cap P$, for some $q \in \mathbb{R}$, is the body with minimum volume, and $\frac{\text{vol}(qC \cap P)}{\text{vol}(qC)} \in [r, r + \delta]$. The bound on m is not surprising, as it does not improve worst-case complexity [5], if C is a ball, but

offers crucial advantages in practice. First, the hidden constant is small. More importantly, if C is a good fit to P , $\text{vol}(P_m)$ increases and m decreases (Fig. 1).




Fig. 1. Different selection of body in our algorithm’s MMC; $r = 0.8$ and $r + \delta = 0.85$. Body C : left is the unit ball; right is the centrally symmetric H-polytope of Sect. 2.3.

We also show that, for constant d , and k (number of generators) increasing, m decreases to 1, when we use ball in MMC, since the schedule constructs an enclosing ball of P . Intuitively, while order increases for constant d , a random zonotope approximates the hypersphere. The latter can be approximated up to ϵ in the Hausdorff metric by a zonotope with $k \leq c(d)(\epsilon^2 |\lg \epsilon|)^{(d-1)/(d+2)}$, $c(d)$ being a constant [4]. This does not directly prove our claim on m but strengthens it intuitively. So, in our experiments, the number of phases is $m \leq 3$ for any order, without rounding for $d \leq 100$.

Considering uniform sampling, BW defines a linear trajectory starting at the current point, using boundary reflections [17]. No theoretical mixing time exists. We show that with the right selection of parameters, BW behaves like an almost perfect uniform sampler even if the walk length is 1. In particular, for this walk length, it generates just $O^*(1)$ points per phase, with sub-linear number of reflections per point, and provides the desired accuracy. To stop sampling when estimating ratio r_i we modify the binomial proportion confidence interval. We use the standard deviation of a sliding window of the last l ratios, thus defining a new empirical convergence criterion; $l = O(1)$ suffices with BW.

Our software contributions build upon and enhance `volesti`¹ a C++ open source library for high dimensional sampling and volume computation with an R interface. We experimentally show that the total number of oracle calls grows as $o^*(d)$ for random zonotopes; the best available theoretical bound is $O^*(d^3)$ [5].

2 Volume Algorithm

The algorithm first constructs a sequence of convex bodies $C_1 \supseteq \dots \supseteq C_m$ intersecting the zonotope P ; the C_i ’s are determined by simulated annealing.

¹ https://github.com/GeomScale/volume_approximation.

A typical choice of C_i 's in this paper is co-centric balls, or centrally symmetric H-polytopes. C_m is chosen for its volume to be computed faster than $\text{vol}(P)$ and easily sampled. Then,

$$\text{vol}(P) = \frac{\frac{\text{vol}(P_m)}{\text{vol}(C_m)}}{\frac{\text{vol}(P_1)}{\text{vol}(P_0)} \frac{\text{vol}(P_2)}{\text{vol}(P_1)} \dots \frac{\text{vol}(P_m)}{\text{vol}(P_{m-1})}} \text{vol}(C_m), \quad P_i = C_i \cap P, i = 1, \dots, m,$$

where $P_0 = P$. Let $r_i = \frac{\text{vol}(P_{i+1})}{\text{vol}(P_i)}$, $i = 0, \dots, m - 1$, $r_m = \frac{\text{vol}(P_m)}{\text{vol}(C_m)}$.

2.1 Uniform Sampling and Oracles for Zonotopes

We use BW to sample approximately uniform points in P_i at each phase i . BW picks a uniformly distributed line ℓ through the current point. It walks on a linear trajectory of length $L = -\tau \ln \eta$, $\eta \sim \mathcal{U}(0, 1)$, reflecting at the boundary. BW can be used to sample only uniform points; in [17] they experimentally show that BW converges fast to the uniform distribution when $\tau \approx \text{diam}(P)$.

The membership oracle is a feasibility problem. A point $p \in P$ iff the following region is feasible: $\sum_{i=1}^k x_i g_i = p$, $-1 \leq x_i \leq 1$, where g_i are the generators of P . Let the uniformly distributed vector on the boundary of the unit ball v define the line ℓ through the current point. The boundary oracle for the intersection $\ell \cap \partial P$ is expressed as a LP. One extreme point of the segment can be computed as follows: $\min -\lambda$, s.t. $p + \lambda v = \sum_{i=1}^k x_i g_i$ $-1 \leq x_i \leq 1$. The second extreme point which corresponds to a negative value of λ is not used by BW. For the BW we need the normal of the facet that intersects ℓ to compute the reflection of the trajectory if needed. We keep the generators that corresponds to $x_i \neq -1, 1$ and then the normal vector is computed straightforwardly.

2.2 Annealing Schedule for Convex Bodies

Given P , the annealing schedule generates the sequence of convex bodies $C_1 \supseteq \dots \supseteq C_m$ defining $P_i = C_i \cap P$ and $P_0 = P$. The main goal is to restrict each ratio r_i in the interval $[r, r + \delta]$ with high probability. We define the following two statistical tests, which can be reduced to t-tests:

[U-test (P_1, P_2)] $H_0: \text{vol}(P_2)/\text{vol}(P_1) \geq r + \delta$
[L-test (P_1, P_2)] $H_0: \text{vol}(P_2)/\text{vol}(P_1) \leq r$

The **U-test** and **L-test** are successful iff null hypothesis H_0 is rejected, namely r_i is upper bounded by $r + \delta$ or lower bounded by r , with high probability, respectively. If we sample N uniform points from P_i then r.v. X that counts points in P_{i+1} , follows $X \sim b(N, r_i)$, the binomial distribution, and $Y = X/N \sim \mathcal{N}(r_i, r_i(1 - r_i)/N)$. Then each sample proportion that counts successes in P_{i+1} over N is an unbiased estimator for the mean of Y , which is r_i .

Perform L-test and U-test

Input : convex bodies P_1, P_2 , cooling parameters r, δ , s.l. $\alpha, \nu, N \in \mathbb{N}$

Sample νN uniform points from P_1

Partition νN points to lists S_1, \dots, S_ν , each of length N

Compute ratios $\hat{r}_i = |\{p \in P_2 : p \in S_i\}| / N, i = 1, \dots, \nu$

Compute the mean, $\hat{\mu}$, and st.d., s , of the ν ratios

if $\hat{\mu} \geq r + t_{\nu-1, \alpha} \frac{s}{\sqrt{\nu}}$ **then** L-test holds, **otherwise** L-test fails

if $\hat{\mu} \leq r + \delta - t_{\nu-1, \alpha} \frac{s}{\sqrt{\nu}}$ **then** U-test holds, **otherwise** U-test fails

Let us now describe the annealing schedule: Each C_i in $C_1 \supseteq \dots \supseteq C_m$ is a scalar multiple of a given body C . Since our algorithm does not use an inscribed body, initialization computes the body with minimum volume, denoted by C' or C_m . This is the last body in the sequence. The algorithm sets $P_0 = P$ and employs C' to decide stopping at the i -th phase.

Initialization. Given C , and interval $[q_{\min}, q_{\max}]$, one employs binary search to compute $q \in [q_{\min}, q_{\max}]$ s.t. both **U-test**($qC, qC \cap P$) and **L-test**($qC, qC \cap P$) are successful. Let $q = (q_{\min} + q_{\max})/2$. If **U-test**($qC, qC \cap P$) succeeds and **L-test**($qC, qC \cap P$) fails, we continue to the left-half of the interval. With inverse outcomes, we continue to the right-half of the interval. If both succeed, stop and set $C' = qC$. The output is C' , denoted by C_m at termination.

Regular Iteration. At iteration i , the algorithm determines P_{i+1} s.t. volume ratio $r_i \in [r, r + \delta]$ with high probability. The schedule samples νN points from P_i and binary searches for a q_{i+1} in an updated interval $[q_{\min}, q_{\max}]$ s.t. both **U-test**($P_i, q_{i+1}C \cap P$) and **L-test**($P_i, q_{i+1}C \cap P$) are successful. Then set $P_{i+1} = q_{i+1}C \cap P$.

Stopping and Termination. The algorithm uses $C' \cap P$ in the i -th iteration for checking whether $\text{vol}(C' \cap P) / \text{vol}(P_i) > r$ with high probability, using only **L-test**, and then stops if **L-test**($P_i, C' \cap P$) holds. Then, set $m = i + 1$, and $P_m = C' \cap P$.

In the t-tests, errors of different types may occur, thus, binary search may enter intervals that do not contain ratios in $[r, r + \delta]$. Hence, there is a probability that annealing schedule fails to terminate. Let β capture the power of a t-test: $\text{pow} = \Pr[\text{reject } H_0 \mid H_0 \text{ false}] = 1 - \beta$.

Theorem 1. Let J be the minimum number of steps by annealing schedule, corresponding to no errors occurring in the t-tests. Let the algorithm perform $M \geq J$ iterations. Let $\beta_{\max}, \beta_{\min}$ be the maximum and minimum among all β 's in the M pairs of t-tests in the **U-test** and **L-test**, respectively. Then, annealing schedule terminates with constant probability, namely:

$$\Pr[\text{an. sched. terminates}] \geq 1 - 2 \frac{\alpha(1 - \beta_{\min}) + \beta_{\max}}{1 - \alpha(1 - \beta_{\min}) + \beta_{\max}} - \frac{2\beta_{\max} - \beta_{\min}^2}{1 - 2\beta_{\max} - \beta_{\min}^2}.$$




Fig. 2. Number of bodies in MMC. For each dimension we generate 10 random zonotopes and we compute the number of bodies, m , in MMC when C is ball. We keep the zonotope with the larger m and then, for that one, we compute m when C is the P-approx.

2.3 Rounding and Convex Bodies in MMC

The annealing schedule allows as to use any C which must (a) be a good fit to P , (b) allow for more efficient sampling than in P , and (c) for faster volume calculation than of $\text{vol}(P)$. For low order ones C shall be an enclosing H-polytope that fits well to P . Indeed it is possible that with certain choices for C rounding is not needed. We define a centrally symmetric H-polytope with $\leq 2k$ facets:

Construct P-approx

Input : The generator matrix $G \in \mathbb{R}^{d \times k}$ of zonotope P

Output: An H-polytope $C \supset P$

compute the eigenvectors of $G^T G$ (has $k - d$ zero eigenvalues)

let the eigenvectors of $k - d$ zero eigenvalues of $G^T G$ form $E \in \mathbb{R}^{k \times (k-d)}$.

compute an orthonormal basis for E , and the orthogonal complement W_{\perp}

Let $Ay \leq b_0$, $A \in \mathbb{R}^{2k \times k}$ be an H-representation of $[-1, 1]^k$

$C := \{x | Mx \leq b_0\}$, $M = AW_{\perp}^T (GW_{\perp}^T)^{-1} \in \mathbb{R}^{2k \times d}$

return C ;

2.4 Experimental Complexity

We perform extended experiments analyzing practical complexity. We use `eigen`² for linear algebra and `lpSolve`³ for LPs. All experiments were performed on a PC with Intel® Core™ i7-6700 3.40 GHz 8 CPU and 32 GB RAM. We use three zonotope generators. All of them pick uniformly a direction for each one of the k segments. Then, (a) $Z_{\mathcal{U}-d-k}$: the length of each segment is uniformly sampled from $[0, 100]$, (b) $Z_{\mathcal{N}-d-k}$: the length of each segment is random from

² eigen.tuxfamily.org.

³ lpsolve.sourceforge.net.




Fig. 3. Experimental complexity for order = 2. Total number of oracle calls is given by the #phases (bodies) \times #steps (points) per phase \times #reflections per step.

$\mathcal{N}(50, (50/3)^2)$ truncated to $[0, 100]$, (c) $Z_{Exp-d-k}$: the length of each segment is random from $Exp(1/30)$ truncated to $[0, 100]$. Total number of boundary oracle calls of our algorithm:

$$\#BOCs = \#phases(bodies) \times \#steps(points) / phase \times \#reflections/point.$$

Figure 2 denotes the best choice between ball and P-approx in MMC. It moreover shows that for order ≤ 5 , the number of phases $m \leq 3$ for $d \leq 100$. In particular, when we use P-approx, m is smaller for order ≤ 4 compared to using balls without rounding. For order equal to 5 the number of balls in MMC is smaller compared to the number of bodies when the choice is the P-approx. Notice that when we use balls in MMC, m decreases for constant d as k increases. Table 1 shows that, for high-order zonotopes, $m = 1$, which implies one or two rejection steps, while the run-time is smaller when we use ball in MMC. It also reports the difference in the run-time for random zonotopes of order = 2 between the cases of using ball and the P-approx in MMC. In all our experiments, BW performs only $O^*(1)$ steps per phase with just a factor of ϵ_i hidden in the complexity. The plot that counts the BW reflections per point in Fig. 3 imply this number grows sub-linearly in d . Hence, the total number of BOCs grows sub-linearly in d .

Table 1. Volume estimation for zonotopes. For each $Z-d-k$ we approximate its volume using ball and the P-approx in MMC. *Body* stands for the type of body in MMC; *order* = k/d , *Vol* the average of volumes over 10 runs; *m* the average number of bodies in MMC; *OracleCalls* is the average number of BOCs; *time* is average time in seconds. We set the error parameter $\epsilon = 0.1$ in all cases.

$Z-d-k$	<i>Body</i>	<i>order</i>	<i>Vol</i>	<i>m</i>	<i>OracleCalls</i>	<i>time</i>
$Z_N-20-2000$	Ball	100	3.69e+83	1	3.52e+03	1 442
$Z_N-20-2000$	P-approx	100	3.54e+83	1	4.10e+03	1 647
$Z_U-30-600$	Ball	20	3.93e+104	1	5.26e+03	451
$Z_U-30-600$	P-approx	20	3.84e+104	1	5.34e+03	554
$Z_U-60-120$	Ball	2	4.31e+139	6	7.94e+04	694
$Z_U-60-120$	P-approx	2	4.18e+139	2	3.39e+04	361
$Z_{Exp}-80-160$	Ball	2	1.68e+187	9	1.67e+05	3 045
$Z_{Exp}-80-160$	P-approx	2	1.82e+187	2	4.22e+04	950
$Z_N-100-200$	Ball	2	9.77e+233	12	2.81e+05	12 223
$Z_N-100-200$	P-approx	2	1.03e+234	3	6.51e+04	2 815

References

1. Althoff, M.: Formal and compositional analysis of power systems using reachable sets. *IEEE Trans. Power Syst.* **29**(5), 2270–2280 (2014)
2. Althoff, M., Dolan, J.M.: Online verification of automated road vehicles using reachability analysis. *IEEE Trans. Robotics* **30**(4), 903–918 (2014). <https://doi.org/10.1109/TRO.2014.2312453>
3. Althoff, M.: *Reachability Analysis and Its Application to the Safety Assessment of Autonomous Cars* (2010). Google-Books-ID: ndN3vgAACAAJ
4. Bourgain, J., Lindenstrauss, J.: Approximating the ball by a Minkowski sum of segments with equal length. *Discrete Comput. Geom.* **9**(2), 131–144 (1993). <https://doi.org/10.1007/BF02189313>
5. Cousins, B., Vempala, S.: Bypassing KLS: Gaussian cooling and an $O^*(n^3)$ volume algorithm. In: *Proceedings of the ACM STOC*, pp. 539–548 (2015)
6. Cousins, B., Vempala, S.: A practical volume algorithm. *Math. Prog. Comp.* **8**(2), 133–160 (2015). <https://doi.org/10.1007/s12532-015-0097-z>
7. Dyer, M., Frieze, A.: On the complexity of computing the volume of a polyhedron. *SIAM J. Comput.* **17**(5), 967–974 (1988). <https://doi.org/10.1137/0217060>
8. Dyer, M., Frieze, A., Kannan, R.: A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM* **38**(1), 1–17 (1991). <http://doi.acm.org/10.1145/102782.102783>
9. Elekes, G.: A geometric inequality and the complexity of computing volume. *Discrete Comput. Geom.* **1**(4), 289–292 (1986). <https://doi.org/10.1007/BF02187701>
10. Emiris, I., Fisikopoulos, V.: Practical polytope volume approximation. *ACM Trans. Math. Soft.* **44**(4), 38:1–38:21 (2018). <https://doi.org/10.1145/3194656>, Prelim. version: Proc. SoCG 2014
11. Gover, E., Krikorian, N.: Determinants and the volumes of parallelotopes and zonotopes. *Linear Algebra Appl.* **413**, 28–40 (2010)

12. Kopetzki, A., Schürmann, B., Althoff, M.: Methods for order reduction of zonotopes. In: Proceedings of the IEEE Annual Conference on Decision & Control (CDC), pp. 5626–5633 (2017). <https://doi.org/10.1109/CDC.2017.8264508>
13. Lovász, L., Kannan, R., Simonovits, M.: Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algor.* **11**, 1–50 (1997)
14. Lovász, L., Vempala, S.: Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithms. *J. Comput. Syst. Sci.* **72**, 392–417 (2006)
15. McMullen, P.: On zonotopes. *Trans. Am. Math. Soc.* **159**, 91–109 (1971). <http://www.jstor.org/stable/1996000>
16. Pereira, A., Althoff, M.: Safety control of robots under computed torque control using reachable sets. In: Proceedings of the IEEE International Conference on Robotics Automation (ICRA), pp. 331–338 (2015). <https://doi.org/10.1109/ICRA.2015.7139020>
17. Polyak, B., Gryazina, E.: Billiard walk - a new sampling algorithm for control and optimization. *IFAC Proc. Vol.* **47**(3), 6123–6128 (2014). <https://doi.org/10.3182/20140824-6-ZA-1003.02312>. 19th IFACWorld Congress
18. Stinson, K., Gleich, D.F., Constantine, P.G.: A randomized algorithm for enumerating zonotope vertices. *CoRR abs/1602.06620* (2016). <http://arxiv.org/abs/1602.06620>
19. Weibel, C.: Implementation and parallelization of a reverse-search algorithm for Minkowski sums. In: *ALENEX* (2010)
20. Ziegler, G.: *Lectures on Polytopes*. Springer, New York (1995). <https://doi.org/10.1007/978-1-4613-8431-1>