International Journal of Computational Geometry & Applications

THE SMALLEST ENCLOSING BALL OF BALLS: COMBINATORIAL STRUCTURE AND ALGORITHMS

KASPAR FISCHER*

Institut für Theoretische Informatik
ETH Zürich, CH-8092 Zürich, Switzerland
fischerk@inf.ethz.ch

BERND GÄRTNER†

Institut für Theoretische Informatik
ETH Zürich, CH-8092 Zürich, Switzerland
gaertner@inf.ethz.ch

Received (received date)
Revised (revised date)
Communicated by (Name)

We develop algorithms for computing the exact smallest enclosing ball of a set of n balls in d-dimensional space. Unlike previous methods, we explicitly address small cases ($n \leq d+2$), derive the necessary primitive operations and show that they can efficiently be realized with rational arithmetic. An implementation (along with a fast[a] and robust floating-point version) is available as part of the CGAL library.[b]

Our algorithms are based on novel insights into the combinatorial structure of the problem. For example, we show that Welzl's randomized linear-time algorithm for computing the ball spanned by a set of points fails to work for balls. Consequently, the existing adaptations of the method to the ball case are incorrect.

In solving the small cases we may assume that the ball centers are affinely independent; in this case, the problem is surprisingly well-behaved: via a geometric transformation and suitable generalization, it fits into the combinatorial model of *unique sink orientations* whose rich structure has recently received considerable attention. One consequence is that Welzl's algorithm *does* work for small instances; moreover, there is a variety of *pivoting* methods for unique sink orientations which have the potential of being fast in practice even for high dimensions.

As a by-product, we show that the problem of finding the smallest enclosing ball of balls with a *fixed point on the boundary* is equivalent to the problem of finding the minimum-norm point in the convex hull of a set of balls; we give algorithms to solve both problems.

*Supported by the Berlin/Zürich joint graduate program "Combinatorics, Geometry, and Computation" (CGC).
†Partly supported by the IST Programme of the EU and the Swiss Federal Office for Education and Science as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG — Effective Computational Geometry for Curves and Surfaces).
[a]For $d = 3$, a set of 1,000,000 balls is processed in less than two seconds on a modern PC.
[b]http://www.cgal.org

1. Introduction

In this paper, we study the problem of finding the closed ball of smallest radius that contains a given set of n closed balls in d-dimensional Euclidean space. This problem—which we denote by SEBB—generalizes the well-understood problem SEBP of finding the smallest enclosing ball of n given points. Applications include collision detection, the computation of bounding sphere hierarchies for clustering or efficient rendering of complex scenes, culling (e.g. for visualization of molecular models[26]), automated manufacturing,[12] and similarity search in feature spaces.[15]

Exact algorithms. The SEBB problem can be solved in time $\mathcal{O}(n)$ for fixed dimension d, where the constant factor is exponential in d. For SEBP, the first (deterministic) linear-time algorithm—based on the *prune-and-search* paradigm—is due to Megiddo,[17,18] with subsequent improvements by Dyer;[5] Welzl's randomized algorithm[25] is the first 'practical' algorithm to achieve this bound. Extending the applicability of prune-and-search, Megiddo[19] and Dyer[6] later showed that the $\mathcal{O}(n)$ bound also applies to SEBB. Linear-time algorithms for SEBB also result from the observation that the problem is of *LP-type*,[16] in which case generic $\mathcal{O}(n)$-methods for this class of problems can be applied.[16,3]

When it comes to actual implementations of these methods, it turns out that none of them work out of the box; the prune-and-search approaches require routines for solving systems of constant-degree algebraic equations, while the LP-type approach asks us to provide a method of computing SEBB for small instances ($n \leq d + 2$). To the best of our knowledge, there are no results that address exact solutions for these primitive operations. Restricting attention to the SEBP problem, we also have Welzl's algorithm at our disposal which has the attractive feature that the primitive operations are very easy to implement. It is therefore tempting to generalize Welzl's algorithm to the case of balls as input. This has indeed been done by Xu *et al.*[27] for $d = 2$, and by David White for general d. Links to White's code appear in many web directories of computational geometry software.[c]

Approximation algorithms. The above exact approaches from computational geometry are complemented with iterative solution methods based on general optimization techniques, see for example the papers by Xu *et al.*[27] and Zhou *et al.*[28]. The most important advantage of these algorithms is that they scale well with d in practice, allowing the problem to be efficiently solved even in high dimensions. A different kind of approach has been pursued by Kumar *et al.*[14]. Their method efficiently finds a small subset of the input balls—a *core set*—approximately spanning the same smallest enclosing ball as the input itself; for this, it repeatedly calls an (approximate or exact) solver for small instances.

Our contribution. As we show in this paper, the extension of Welzl's algorithm to balls as input does not work. For $d = 2$, we exhibit a concrete, nondegenerate

[c]Meanwhile, his web site has disappeared; a version of White's code is available to us.

example of five input balls for which Welzl's algorithm (and in particular, the codes by Xu *et al.* and White) may crash.[d] In other instances, it might compute balls which are too large. The reason behind these failures becomes apparent in the next section when we discuss basic properties of the SEBB problem (in particular, properties of the SEBP problem that do not generalize).

Having said this, it might come as a surprise that Welzl's method *does* extend to the case where the ball centers are affinely independent, and this becomes an important special case, if we want to solve the small instances within the LP-type framework mentioned above. However, already in the point case there are inputs (e.g. the vertices of a regular simplex) for which Welzl's algorithm has complexity $\Omega(2^d)$—it enumerates all candidate solutions before finding the optimal one—so this approach is limited to small or only moderately high dimensions. For such d, we can even solve the small cases (up to $d+2$ balls) by an explicit exhaustive search, circumventing the affine independence requirement. In this approach, the primitive operations need to be done carefully but are otherwise as easy as in Welzl's method; the resulting code is very fast for dimensions up to 10.

For dimensions beyond that, it is known that for SEBP, the performance can be improved under affine independence of the input points: in theory, there are subexponential methods,[7] and there is also a very practical approach[10] that reduces the expected number of candidate solutions to $(3/2)^d$. Using the recent concept of *unique sink orientations* (USO), this can be improved even further.[24] We show that the SEBB problem with affinely independent ball centers, too, exhibits enough structure to fit into the USO framework—even though we have to overcome some nontrivial obstacles on the way. The resulting procedure for small instances can be plugged into the LP-type algorithm which invokes it only a *subexponential* number of times.[9] This means that the *exponential* complexity of the procedure itself is the bottleneck of the algorithm, and savings in the exponent pay off.

On the practical side, the USO approach allows for *pivoting* methods for which we may not be able to give performance guarantees, but which have the potential of being fast for almost any input. Candidates are Murty's method[21] and approaches based on random walks in the orientation. Even though this approach requires affine independence, the gain in runtime would justify an embedding of our (at most $(d+2)$-element) ball set into dimension $d+1$, followed by a symbolic perturbation to attain affine independence.[e] On the theoretical side, our USO approach improves over the previously best complexity for small instances, that is, it breaks the aforementioned $\Omega(2^d)$-barrier resulting from complete enumeration of all candidate solutions.

[d] For points as input, Xu *et al.*'s implementation of Welzl's algorithm works but is not state-of-the-art: the *move-to-front*[25] and the *farthest-first*[8] heuristic (which is also employed in our code) reduce the runtime in practice significantly.

[e] If done properly, this requires only minor modifications of the formulas for the affinely independent case.[11]

Still, our techniques remain impractical for large d. The pivoting approach might change this, but currently, for $d > 30$, our results are more significant with respect to the geometric and combinatorial properties of the SEBB problem. In particular, our use of the *inversion* transform relates SEBB to the problem of finding the distance of a point to the convex hull of a set of balls. Moreover, the formulation of the latter geometric problem in terms of well-behaved mathematical programs might prove useful in optimization-based techniques for both problems.

From a practical point of view, our main focus is the case of small d. Here, we are not aware of any (exact or approximate) methods which outperform our code.

2. Basics

A d-dimensional ball with center $c \in \mathbb{R}^d$ and nonnegative radius $\rho \in \mathbb{R}$ is the point set $B(c, \rho) = \left\{x \in \mathbb{R}^d \mid \|x - c\|^2 \leq \rho^2\right\}$, and we write c_B and ρ_B to denote the center and radius, respectively, of a given ball B. We say that a ball is *proper* if its radius is nonzero.

Ball $B' = B(c', \rho')$ is contained in ball $B = B(c, \rho)$ if and only if

$$\|c - c'\| \leq \rho - \rho', \tag{1}$$

with equality if and only if B' is *internally tangent* to B.

We define the *miniball* $\textsc{mb}(U)$ of a finite set U of balls set U in $\mathbb{R}^d$ to be the unique ball of smallest radius which contains all balls in U (Fig. 1). We also set $\textsc{mb}(\emptyset) = \emptyset$ (note that $\textsc{mb}(\emptyset)$ is not a ball). The next lemma shows that $\textsc{mb}(U)$ is well-defined.

Lemma 2.1. *For a finite nonempty set U of balls, there is a unique ball of smallest radius that contains all balls of U.*

Proof. A standard compactness argument shows that some enclosing ball of smallest radius exists. If this radius is zero, the lemma easily follows. Otherwise, we use *convex combinations of balls*,[25] a concept we will also need later on: a proper ball $B = B(c, \rho)$ can be written as the set of points $x \in \mathbb{R}^d$ satisfying $f_B(x) \leq 1$

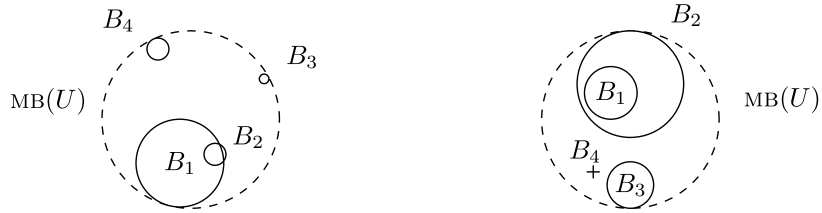

Fig. 1: Two examples in $\mathbb{R}^2$ of the miniball $\textsc{mb}(U)$ for $U = \{B_1, \ldots, B_4\}$. Throughout the paper, points and balls of zero radius are drawn as '+'.

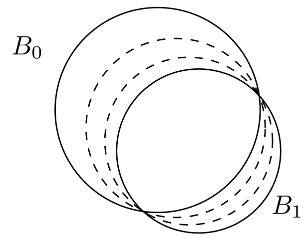

Fig. 2. Convex combinations B_λ (dashed) of the two balls B_0 and B_1 (solid), for $\lambda \in \{1/3, 2/3\}$.

for $f_B(x) = \|x - c\|^2/\rho^2$. For any $\lambda \in [0,1]$, the *convex combination* B_λ of two intersecting balls B, B' is the set of points x fulfilling

$$f_{B_\lambda}(x) = (1-\lambda) f_B(x) + \lambda f_{B'}(x) \le 1.$$

It is easily verified that B_λ is a ball again, that it contains $B \cap B'$ and that the points on the boundary of both B and B' lie on the boundary of B_λ, too. Moreover, if B and B' are distinct then the radius of the ball B_λ is, for any $\lambda \in (0,1)$, strictly smaller than the maximum of the radii of B, B' (Fig. 2). The latter property immediately proves that $\text{MB}(U)$ is well-defined: assuming there are two distinct smallest enclosing balls, a proper convex combination of them is still enclosing, but has smaller radius, a contradiction. □

The following optimality criterion generalizes a statement for points due to Seidel.[23] Recall that a point $q \in \mathbb{R}^d$ lies in the convex hull $\text{conv}(P)$ of a finite point set $P \subseteq \mathbb{R}^d$ if and only if $\min_{p \in P} (p-q)^T u \le 0$ for all unit vectors u.

Lemma 2.2. *Let V be a nonempty set of balls, all internally tangent to some ball D. Then $D = \text{MB}(V)$ if and only if $c_D \in \text{conv}(\{c_B \mid B \in V\})$.*

Proof. For direction ($\Leftarrow$), assume $D \neq \text{MB}(V)$, i.e., there exists an enclosing ball D' with radius $\rho_{D'} < \rho_D$. Write its center (which must be different from c_D by the internal tangency assumption) as $c_{D'} = c_D + \lambda u$ for some unit vector u and $\lambda > 0$. Then the distance from $c_{D'}$ to the farthest point in a ball $B \in V$ is

$$\begin{aligned}\delta_B &= \|c_{D'} - c_B\| + \rho_B \\ &= \sqrt{(c_D + \lambda u - c_B)^T (c_D + \lambda u - c_B)} + \rho_B \\ &= \sqrt{\|c_D - c_B\|^2 + \lambda^2 u^T u - 2\lambda (c_B - c_D)^T u} + \rho_B \\ &= \sqrt{(\rho_D - \rho_B)^2 + \lambda^2 - 2\lambda (c_B - c_D)^T u} + \rho_B, \end{aligned} \qquad (2)$$

because (1) holds with equality by our tangency assumption. Since D' is enclosing, we must have

$$\rho_{D'} \geq \max_{B \in V} \delta_B. \tag{3}$$

Furthermore, the observation preceding the lemma yields the existence of $B' \in V$ such that $(c_{B'} - c_D)^T u \leq 0$, for c_D lies in the convex hull of the centers of V. Consequently,

$$\delta_{B'} > \sqrt{(\rho_D - \rho_{B'})^2} + \rho_{B'} = \rho_D > \rho_{D'}$$

by equation (2), a contradiction to (3).

For direction ($\Rightarrow$), suppose that c_D does not lie in the convex hull of the centers of V. By the above observation there exists a vector u of unit length with $(c_B - c_D)^T u > 0$ for all $B \in V$. Consider the point $c_{D'} := c_D + \lambda u$, for some strictly positive $\lambda < 2 \min_{B \in V} (c_B - c_D)^T u$. According to (2), $\delta_B < (\rho_D - \rho_B) + \rho_B = \rho_D$ for all B, and consequently, the ball D' with center $c_{D'}$ and radius $\max_B \delta_B < \rho_D$ is enclosing, contradiction. □

Another property we will use for our algorithm in Sec. 3 is the following intuitive statement which has been proved by Welzl for points.[25]

Lemma 2.3. *If the ball $B \in U$ is properly contained in the miniball* $\mathrm{MB}(U)$ *(i.e., not internally tangent to it) then* $\mathrm{MB}(U) = \mathrm{MB}(U \setminus \{B\})$*, equivalently,* $B \subseteq \mathrm{MB}(U \setminus \{B\})$.

Proof. Consider the convex combination D_λ of the balls $D = \mathrm{MB}(U)$ and $D' = \mathrm{MB}(U \setminus \{B\})$; it continuously transforms D into D' as λ ranges from 0 to 1 and contains all balls in $U \setminus \{B\}$. Since B is not tangent to $\mathrm{MB}(U)$, there is a $\lambda' > 0$ such that $D_{\lambda'}$ still encloses all balls from U. But if D and D' do not coincide, $D_{\lambda'}$ has smaller radius than D, a contradiction to the minimality of $D = \mathrm{MB}(U)$. □

Motivated by this observation, we call a set $U' \subseteq U$ a *support set of* U if all balls in U' are internally tangent to $\mathrm{MB}(U)$ and $\mathrm{MB}(U') = \mathrm{MB}(U)$. An inclusion-minimal support set of U is called *basis of* U (see Fig. 3), and we call a ball set V a *basis* if it is a basis of itself. A standard argument based on Helly's Theorem reveals that the miniball is determined by a support set of size at most $d + 1$.

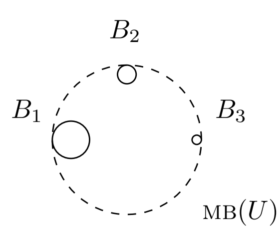

Fig. 3. $U = \{B_1, B_2, B_3\}$ is a support set (but not a basis) of U; $V = \{B_1, B_3\}$ is a basis.

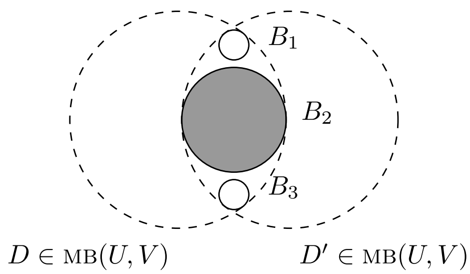

Fig. 4. $\mathrm{MB}(U,V)$ may contain several balls (left) or none (right): set $U = \{B_1, B_2, B_3\}$, $V = \{B_2\}$.

Lemma 2.4. *Let U be a set of at least $d+1$ balls in $\mathbb{R}^d$. Then there exists a subset $U' \subseteq U$ of $d+1$ balls such that $\mathrm{MB}(U) = \mathrm{MB}(U')$.*

Proof. Let $D = \mathrm{MB}(U)$ and consider the set $I = \bigcap_{B \in U} B(c_B, \rho_D - \rho_B)$. Observe that $B(c_B, \rho_D - \rho_B)$ is the set of all centers which admit a ball of radius ρ_D that encloses B. By the existence and uniqueness of $\mathrm{MB}(U)$, I thus contains exactly one point, namely c_D. It follows that $\bigcap_{B \in U} \operatorname{int} B(c_B, \rho_D - \rho_B) = \emptyset$, where $\operatorname{int} B'$ denotes the interior of ball B'. Helly's Theorem[f] yields a set $U' \subseteq U$ of $d+1$ elements such that $\bigcap_{B \in U'} \operatorname{int} B(c_B, \rho_D - \rho_B) = \emptyset$. Consequently, no ball of radius $< \rho_D$ encloses the balls U', and thus $\mathrm{MB}(U)$ and $\mathrm{MB}(U')$ have the same radius. This however implies $\mathrm{MB}(U) = \mathrm{MB}(U')$, since we would have found two different miniballs of U' otherwise. □

Lemma 2.5. *The centers of a basis V of U are affinely independent.*

Proof. The claim is obvious for $V = \emptyset$. Otherwise, by Lemma 2.2, the center c_D of the miniball $D = \mathrm{MB}(V) = \mathrm{MB}(U)$ can be written as $c_D = \sum_{B \in V} \lambda_B c_B$ for some coefficients $\lambda_B \geq 0$ summing up to 1. Observe that $\lambda_B > 0$, $B \in V$, by minimality of V. Suppose that the centers $\{c_B \mid B \in V\}$ are affinely dependent, or, equivalently, that there exist coefficients μ_B, not all zero, such that $\sum_{B \in V} \mu_B c_B = \mathbf{0}$ and $\sum \mu_B = 0$. Consequently,

$$c_D = \sum_{B \in V} (\lambda_B + \alpha \mu_B)\, c_B \qquad \text{for any } \alpha \in \mathbb{R}. \tag{4}$$

Change α continuously, starting from 0, until $\lambda_{B'} + \alpha \mu_{B'} = 0$ for some B'. At this moment all nonzero coefficients $\lambda'_B = \lambda_B + \alpha \mu_B$ of the combination (4) are strictly positive, sum up to 1, but $\lambda'_{B'} = 0$, a contradiction to the minimality of V. □

[f] *Helly's Theorem*[4] states that if $C_1, \ldots, C_m \subset \mathbb{R}^d$ are $m \geq d+1$ convex sets such that any $d+1$ of them have a common point then also $\bigcap_{i=1}^m C_i$ is nonempty.

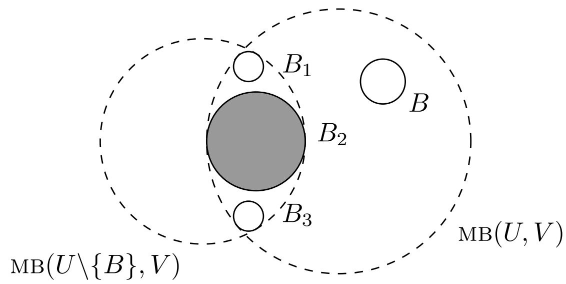

Fig. 5. Ball B cannot be dropped although it is properly contained in $\text{MB}(U, V)$.

We conclude this section with some basic properties of $\text{MB}(U, V)$ which is the following generalization of $\text{MB}(U)$. For sets $U \supseteq V$ of balls, we denote by $\text{B}(U, V)$ the set of balls B which contain the balls U and to which the balls in V are internally tangent (we set $\text{B}(\emptyset, \emptyset) = \{\emptyset\}$). Based on this, we define $\text{MB}(U, V)$ to be the set of smallest balls in $\text{B}(U, V)$; in case $\text{MB}(U, V)$ contains exactly one ball D, we abuse notation and refer to D as $\text{MB}(U, V)$. Observe that $\text{MB}(U) = \text{MB}(U, \emptyset)$ and hence any algorithm for computing $\text{MB}(U, V)$ solves the SEBB problem. However, several intuitive properties of $\text{MB}(U)$ do not carry over to $\text{MB}(U, V)$: the set $\text{MB}(U, V)$ can be empty, or there can be *several* smallest balls in $\text{B}(U, V)$, see Fig. 4. Furthermore, properly contained balls cannot be dropped as in the case of $\text{MB}(U)$ (Lemma 2.3): for a counterexample refer to Fig. 5, where $\text{MB}(U, V) \neq \text{MB}(U \setminus \{B\}, V)$ for $V = \{B_2\}$ and $U = \{B_1, B_2, B_3, B\}$, although B is properly contained in $\text{MB}(U, V)$.

In the sequel we will also deal with

$$\text{MB}_p(U) := \text{MB}(U \cup \{p\}, \{p\}), \tag{5}$$

where $p \in \mathbb{R}^d$ is some point and U as usual is a set of balls.[g] Again $\text{MB}_p(U)$ may be empty (place p inside the convex hull $\text{conv}(U) := \text{conv}(\bigcup_{B \in U} B)$), but in the nonempty case it contains a unique ball. This follows from

Lemma 2.6. *Let $U \supseteq V$ be two sets of balls, V being a set of points (balls of zero radius). Then $\text{MB}(U, V)$ consists of at most one ball.*

Proof. If $D, D' \in \text{MB}(U, V)$, their convex combination D_λ contains U and in addition has the points V on the boundary. Thus, $D_\lambda \in \text{B}(U, V)$ for any $\lambda \in [0, 1]$. If D and D' were distinct, a proper convex combination would have smaller radius than D' or D, a contradiction to the minimality of D, D'. □

Combining a compactness argument as in the proof of Lemma 2.1 with the techniques from the previous lemma, we can also show the following.

[g] In writing $U \cup \{p\}$ we abuse notation and identify the ball $B(p, 0)$ with the point p.

Lemma 2.7. *Let U be a set of balls and $p \in \mathbb{R}^d$ such that no ball in U contains p. Then* $\textsc{mb}_p(U)$ *is empty if and only if* $p \in \operatorname{conv}(U)$.

Without the assumption on U and p, it may happen that $\textsc{mb}_p(U) \neq \emptyset$ although $p \in \operatorname{conv}(U)$ (take a single ball, $U = \{B\}$, and a point p on its boundary).

3. Algorithms

Our algorithms are combinatorial in nature and based on the notion of $\textsc{mb}(U,V)$. Of particular importance is $\textsc{mb}(V,V)$, which will take the role of a 'base' case and for which we therefore need explicit formulas; further properties of $\textsc{mb}(U,V)$ are developed in Secs. 4, 5 and 6.

Lemma 3.1. *Let V be a basis of U. Then* $\textsc{mb}(V,V) = \textsc{mb}(U)$, *and this ball can be computed in time* $\mathcal{O}(d^3)$.

Proof. For $V = \emptyset$, the claim is trivial, so assume $V \neq \emptyset$. As a basis of U, V satisfies $\textsc{mb}(V) = \textsc{mb}(U)$. Since the balls in V must be tangent to $\textsc{mb}(U)$ (Lemma 2.3), we have $\textsc{mb}(V) \in \textsc{mb}(V,V)$. But then *any* ball in $\textsc{mb}(V,V)$ is a smallest enclosing ball of V, so Lemma 2.1 guarantees that $\textsc{mb}(V,V)$ is a singleton.

Let $V = \{B_1, \ldots, B_m\}$, $m \leq d+1$, and observe that $B(c,\rho) \in \textsc{b}(V,V)$ if and only if $\rho \geq \rho_{B_i}$ and $\|c - c_{B_i}\|^2 = (\rho - \rho_{B_i})^2$ for all i. Defining $z_{B_i} = c_{B_i} - c_{B_1}$ for $1 < i \leq m$ and $z = c - c_{B_1}$, these conditions are equivalent to $\rho \geq \max_i \rho_{B_i}$ and

$$\begin{aligned} z^T z &= (\rho - \rho_{B_1})^2, \\ (z_{B_i} - z)^T (z_{B_i} - z) &= (\rho - \rho_{B_i})^2, \qquad 1 < i \leq m. \end{aligned} \tag{6}$$

Subtracting the latter from the former yields the $m-1$ linear equations

$$2 z_{B_i}^T z - z_{B_i}^T z_{B_i} = 2\rho\,(\rho_{B_i} - \rho_{B_1}) + \rho_{B_1}^2 - \rho_{B_i}^2, \qquad 1 < i \leq m.$$

If $B(c,\rho) = \textsc{mb}(V,V)$ then $c \in \operatorname{conv}(\{c_{B_1}, \ldots, c_{B_m}\})$ by Lemma 2.2. Thus we get $c = \sum_{i=1}^m \lambda_i c_{B_i}$ with the λ_i summing up to 1. Then, $z = \sum_{i=2}^m \lambda_i (c_{B_i} - c_{B_1}) = Q\lambda$, where $Q = (z_{B_2}, \ldots, z_{B_m})$ and $\lambda = (\lambda_2, \ldots, \lambda_m)^T$. Substituting this into our linear equations results in

$$2 z_{B_i}^T Q\lambda = z_{B_i}^T z_{B_i} + \rho_{B_1}^2 - \rho_{B_i}^2 + 2\rho\,(\rho_{B_i} - \rho_{B_1}), \qquad 1 < i \leq m. \tag{7}$$

This is a linear system of the form $A\lambda = e + f\rho$, with $A = 2Q^T Q$. So $B(c,\rho) = \textsc{mb}(V,V)$ satisfies $c - c_{B_1} = z = Q\lambda$ with (λ,ρ) being a solution of (6), (7) and $\rho \geq \max_i \rho_{B_i}$. Moreover, the columns of Q are linearly independent as a consequence of Lemma 2.5, which implies that A is in fact regular.

Hence we can in time $\mathcal{O}(d^3)$ compute A^{-1}, find the solution space of the linear system (which is one-dimensional, parameterized by ρ) and substitute this into the quadratic equation (6). From the possible solutions (λ,ρ) we select one such that $\rho \geq \max_i \rho_{B_i}$, $\lambda \geq 0$ and $\lambda_1 = 1 - \sum_{i=2}^m \lambda_i \geq 0$; by $\textsc{mb}(V) = \textsc{mb}(V,V)$ and Lemma 2.2 such a pair (λ,ρ) exists, and in fact, there is only one such pair because the ball determined by any (λ, ρ) with the above properties is tangent, enclosing and by Lemma 2.2 equal to $\text{MB}(V)$. □

We note that the existing, robust formulas for computing $\text{MB}(V, V)$ in the point case[8] can be generalized to balls (and are employed in our code).

```
procedure welzl(U,V)
{ Intended to compute MB(U,V) but does not work }
{ Precondition: U ⊇ V, |MB(U,V)| = 1 }
begin
    if U = V then
        return MB(V,V)
    else
        choose B ∈ U\V uniformly at random
        D:= welzl(U\{B},V)
        if B ⊄ D then
            return welzl(U,V ∪ {B})
        end
    end
end welzl
```

Fig. 6. Welzl's algorithm for balls.

3.1. *Welzl's algorithm*

Welzl's algorithm[25] for the SEBP problem can easily be 'rewritten' for balls (Fig. 6). However, the resulting procedure does not work anymore in general. The reason for this is that Welzl's Lemma,[25] underlying the algorithm's correctness proof in the point case, fails for balls:

Dilemma 3.2. Let $U \supseteq V$ be sets of balls such that $\text{MB}(U, V)$ and $\text{MB}(U \setminus \{B\}, V)$ contain unique balls each. If

$$B \not\subseteq \text{MB}(U \setminus \{B\}, V)$$

for some $B \in U \setminus V$ then B is tangent to $\text{MB}(U, V)$, so $\text{MB}(U, V) = \text{MB}(U, V \cup \{B\})$.

A counterexample is depicted in Fig. 7: the point B_5 is not contained in the ball $D = \text{MB}(\{B_1, B_3, B_4\}, \{B_1, B_3, B_4\})$, but B_5 is not tangent to

$$D' = \text{MB}(\{B_1, B_3, B_4, B_5\}, \{B_1, B_3, B_4\}).$$

As a matter of fact, feeding the procedure welzl with the balls from Fig. 7 produces incorrect results from time to time, depending on the outcomes of the internal random choices.[h] If in each call, B is chosen to be the ball of lowest index in $U \setminus V$, the algorithm eventually gets stuck when it tries to find the ball $\text{MB}(\{B_1, B_3, B_4, B_5\}, \{B_1, B_3, B_4, B_5\})$, which does not exist (see Fig. 8). Observe that this counterexample is free of degeneracies, and that no set $\text{MB}(U, V)$ contains more than one ball.

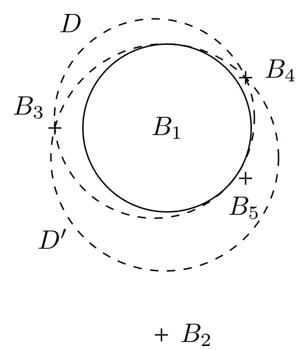

Fig. 7. Five balls $\{B_1, \dots, B_5\}$ in $\mathbb{R}^2$ for which Welzl's algorithm may fail.

3.2. *LP-type algorithm*

From the combinatorial point of view, we want to find an inclusion-minimal set $V \subseteq U$ spanning the same miniball as U; since we then have $\text{MB}(U) = \text{MB}(V, V)$, it is straightforward to compute $\text{MB}(U)$ from V (Lemma 3.1). This formulation of the problem fits nicely into the framework of so-called *LP-type problems*.[16]

In general, an *LP-type problem* is a pair (T, w) with $w : 2^T \to \Omega$ for Ω some ordered set, satisfying the following two conditions for all $U' \subseteq U \subseteq T$.

(i) $w(U') \leq w(U)$ (*monotonicity*), and
(ii) $w(U') < w(U)$ implies the existence of $B \in U$ with $w(U') < w(U' \cup \{B\})$ (*locality*).

In LP-type terminology, a *basis of* U is an inclusion-minimal subset $V \subseteq U$ such that $w(V) = w(U)$, and the goal is to find a basis of T.

LP-type problems can be solved by the *MSW-algorithm*[16] of Fig. 9, provided two primitives are available: for the *violation test* we are given a basis $V \subseteq U$ (meaning V is a basis of itself but not necessarily of U) and $B \in U$, and we need

[h]The balls in Fig. 5 already constitute a counterexample to Dilemma 3.2 but cannot be used to fool Welzl's algorithm.

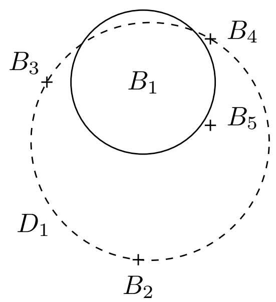

$B_1 \not\subseteq D_1 = \mathsf{welzl}(\{B_2, \ldots, B_5\}, \{\})$
$\rightsquigarrow \mathsf{welzl}(\{B_1, \ldots, B_5\}, \{B_1\})$

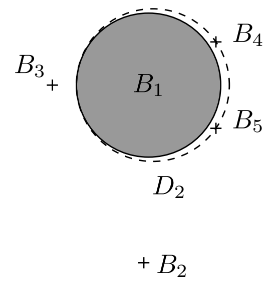

$B_3 \not\subseteq D_2 = \mathsf{welzl}(\{B_1, B_4, B_5\}, \{B_1\})$
$\rightsquigarrow \mathsf{welzl}(\{B_1, B_3, B_4, B_5\}, \{B_1, B_3\})$

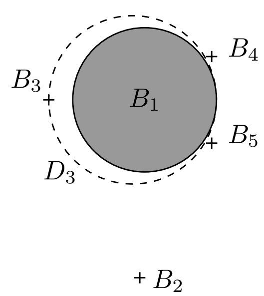

$B_5 \not\subseteq D_3 = \mathsf{welzl}(\{B_1, B_3\}, \{B_1, B_3\})$
$\rightsquigarrow \mathsf{welzl}(\{B_1, B_3, B_5\}, \{B_1, B_3, B_5\})$

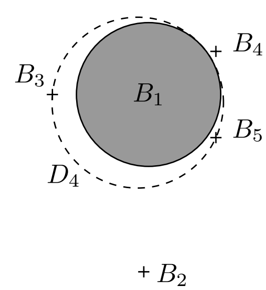

$B_4 \not\subseteq D_4 = \mathsf{welzl}(\{B_1, B_3, B_5\}, \{B_1, B_3\})$
$\rightsquigarrow \mathsf{welzl}(\{B_1, B_3, B_4, B_5\}, \{B_1, B_3, B_4\})$

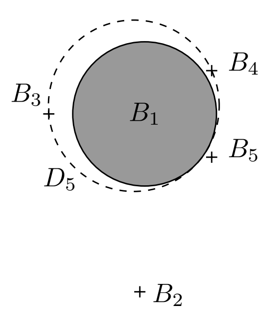

$B_5 \not\subseteq D_5 = \mathsf{welzl}(\{B_1, B_3, B_4\}, \{B_1, B_3, B_4\})$
$\rightsquigarrow \mathsf{welzl}(V, V)$, for $V = \{B_1, B_3, B_4, B_5\}$

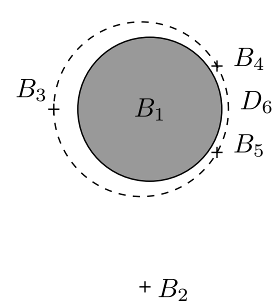

$\textsc{mb}(V, V) = \emptyset$ because B_1 is not tangent to
$\textsc{b}(\{B_3, B_4, B_5\}, \{B_3, B_4, B_5\}) = \{D_6\}$.

Fig. 8. A failing run of Welzl's algorithm on the circles from Fig. 7.

```
procedure msw(U,V)
{ Computes a basis of U }
{ Precondition: V ⊆ U is a basis }
begin
    if U = V then
        return V
    else
        choose B ∈ U\V uniformly at random
        V:= msw(U\{B},V)
        if B violates V then
            return msw(U,basis(V,B))
        end
    end
end msw
```

Fig. 9. The MSW-algorithm.

to test whether $w(V) < w(V \cup \{B\})$. In the *basis computation* basis(V, B), V is a basis such that $B \in U$ violates V, and the task is to find a basis of $V \cup \{B\}$. If the *combinatorial dimension* (size of the largest basis) is constant then both primitive operations can be realized in constant time.

For SEBB, we define the function $w : 2^T \to \mathbb{R}^+ \cup \{-\infty\}$ to map a subset $U \subseteq T$ to the radius of $\textsc{mb}(U)$, with the convention that the radius of $\textsc{mb}(\emptyset)$ is $-\infty$. Using the properties from Sec. 2, it is easily shown that any instance of SEBB in this formulation is an LP-type problem of combinatorial dimension at most $d + 1$.

In this case, the primitive operations are the following. The violation test needs to check whether $B \not\subseteq \textsc{mb}(V)$ (as V is always a basis, Lemma 3.1 can be used to compute $\textsc{mb}(V) = \textsc{mb}(V, V)$). In the *basis computation* we have a basis V and a violating ball B (i.e., $B \not\subseteq \textsc{mb}(V)$), and we are to produce a basis of $V \cup \{B\}$. By Lemma 2.3, the ball B is internally tangent to $\textsc{mb}(V \cup \{B\})$. A basis of $V \cup \{B\}$ can then be computed in a brute-force manner[i] by using Lemma 3.1, as follows.

We generate all subsets V', $B \in V' \subseteq V \cup \{B\}$, in increasing order of size. For each V' we test whether it is a support set of $V \cup \{B\}$. From our enumeration order it follows that the first set V' which passes this test constitutes a basis of $V \cup \{B\}$.

We claim that V' is a support set of $V \cup \{B\}$ if and only if the computations from Lemma 3.1 go through and produce a ball that in addition encloses the balls in $V \cup \{B\}$: if V' is a support set of $V \cup \{B\}$ then it is, by our enumeration order,

[i]We will improve on this in Sec. 6. Also, Welzl's algorithm could be used here, by lifting and subsequently perturbing the centers, but this will not be better than the brute-force approach, in the worst case.

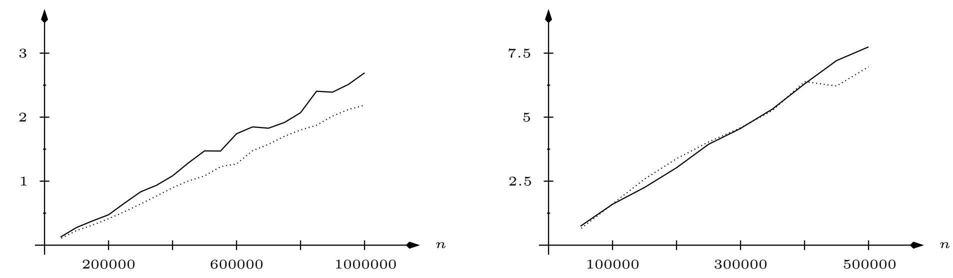

Fig. 10. Running times (in seconds) of our heuristic in $\mathbb{R}^3$ (left) and $\mathbb{R}^{10}$ (right) for n balls with center coordinates and radii uniformly distributed in $[0,1]$ (solid) and for n balls with centers uniformly distributed on the unit sphere and radius uniformly drawn from $[0,1/2]$ (dotted). All times were measured on a 480 MHz Sun Ultra 4 workstation with double arithmetic.

a basis and hence the lemma applies. Conversely, a successful computation yields a ball $D \in \mathrm{B}(V', V')$ (enclosing $V \cup \{B\}$) whose center is a convex combination of the centers of V'; by Lemma 2.2, $D = \mathrm{MB}(V') = \mathrm{MB}(V \cup \{B\})$.

Plugging these primitives into algorithm msw yields an expected $\mathcal{O}(d^3 2^{2d} n)$ algorithm for computing the miniball $\mathrm{MB}(U)$ of any set of n balls in d-space.[16] Moreover, it is possible to do all computations in rational arithmetic (provided the input balls have rational coordinates and radii): although the center and the radius of the miniball may have irrational coordinates, the calculations in the proof of Lemma 3.1 show that they actually are of the form $\alpha_i + \beta_i \sqrt{\gamma}$, where $\alpha_i, \beta_i, \gamma \in \mathbb{R}$ and where $\gamma \geq 0$ is the discriminant of the quadratic equation (6). Therefore, we can represent the coordinates and the radius by pairs $(\alpha_i, \beta_i) \in \mathbb{R}^2$, together with the discriminant γ. Since the only required predicate is the containment test, which boils down to determining the sign of an algebraic number of degree 2, all computations can easily be done in $\mathbb{Q}$.

We have implemented the algorithm in C++. The code follows the generic programming paradigm and has been released with CGAL 3.0. It is parameterized with a number type F; choosing F to be a type realizing rational numbers of arbitrary precision,[j] no roundoff errors occur and the computed ball is the exact smallest enclosing ball of the input balls.

Under a naive floating-point implementation, numerical problems may arise when balls are 'almost' tangent to the current miniball. In order to overcome these issues, we also provide a (deterministic) variant of algorithm msw. In this heuristic—it comes without any theoretical guarantee on the running time—we maintain a basis V (initially consisting of a single input ball) and repeatedly add to it, by an invocation of the basis computation, a ball farthest away from the basis,

[j]Efficient implementations of such types are available, see for instance the CORE or the LEDA library.[13,20]

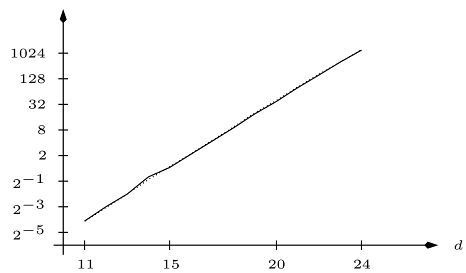

Fig. 11. Running time (in seconds) of algorithm msw (dotted) and our heuristic (solid) when fed with the vertices e_i, $i \in \{1, \ldots, d+1\}$, of the d-dimensional simplex embedded in $\mathbb{R}^{d+1}$. All times were measured on a 480 Mhz Sun Ultra 4 workstation with double arithmetic.

that is, a ball B' satisfying

$$\|c - c_{B'}\| + \rho_{B'} = \max_{B \in U} (\|c - c_B\| + \rho_B) =: \chi_V,$$

with c being the center of $\textsc{mb}(V)$. The algorithm stops as soon as χ_V is smaller or equal to the radius of $\textsc{mb}(V)$, i.e., when all balls are contained in $\textsc{mb}(V)$. This method, together with a suitable adaptation of efficient and robust methods for the point case,[8] handles degeneracies in a satisfactory manner. An extensive testsuite containing various degenerate configurations of balls is passed without problems. Some running times are shown in Figs. 10 and 11.

4. Signed balls and shrinking

In this section we show that under a suitable generalization of SEBB, one of the input balls can be assumed to be a point, and that SEBB can be reduced to the problem of finding the miniball with some point *fixed on the boundary*. With this, we prepare the ground for the more sophisticated material we are going to develop in Secs. 5 and 6.

Recall that a ball $B = B(c, \rho)$ encloses a ball $B' = B(c', \rho')$ if and only if relation (1) holds. Now we are going to use this relation for *signed* balls. A signed ball is of the form $B(c, \rho)$, where—unlike before—ρ can be *any* real number, possibly negative. $B(c, \rho)$ and $B(c, -\rho)$ represent the same ball $\{x \in \mathbb{R}^d \mid \|x - c\|^2 \leq \rho^2\}$, meaning that a signed ball can be interpreted as a regular ball with a sign attached to it; we simply encode the sign into the radius. If $\rho \geq 0$, we call the ball *positive*, otherwise *negative*.

Definition 4.1. Let $B = B(c, \rho)$ and $B' = B(c', \rho')$ be signed balls. B *dominates* B' if and only if

$$\|c - c'\| \leq \rho - \rho'. \tag{8}$$

B *marginally dominates* B' if and only if (8) holds with equality.

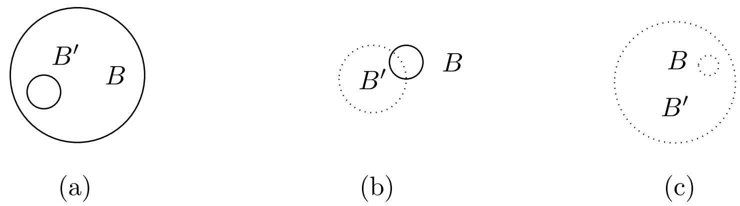

Fig. 12. B dominates B' (a) if both balls are positive and $B \supseteq B'$, (b) if B is positive and B' negative and the two intersect, or (c) if both are negative and $B \subseteq B'$. (Negative balls are drawn dotted, positive ones solid as usual.)

Figure 12 depicts three examples of the dominance relation. Furthermore, marginal dominance has the following geometric interpretation: if both B, B' are positive, B' is internally tangent to B; if B is positive and B' is negative then B and B' are *externally tangent* to each other. Finally, if both B, B' are negative then B is internally tangent to B'.

We generalize SEBB to the problem of finding the ball of smallest signed radius that dominates a given set of signed balls. For two sets $U \supseteq V$ of signed balls, we denote by $\textsc{b}(U,V)$ the set of signed balls that dominate the balls in U and that marginally dominate the balls in V. Then, $\textsc{mb}(U,V)$ is the set of smallest signed balls in $\textsc{b}(U,V)$, and we again set $\textsc{mb}(\emptyset,\emptyset) = \{\emptyset\}$ and abuse notation in writing $\textsc{mb}(U,V)$ for the ball D in case $\textsc{mb}(U,V)$ is a singleton $\{D\}$.

Figure 13 depicts some examples of $\textsc{mb}(U) := \textsc{mb}(U,\emptyset)$. In particular, Fig. 13(c) illustrates that this generalization of SEBB covers the problem of computing a ball of largest volume (equivalently, smallest negative radius) contained in the intersection $I = \bigcap_{B \in U} B$ of a set of balls U; for this, simply encode the members of U as negative balls.

At this stage, it is not yet clear that $\textsc{mb}(U)$ is always nonempty and contains a unique ball. With the following argument, we can easily show this. Fix any ball O and define $s_O : B \mapsto B(c_B, \rho_B - \rho_O)$ to be the map which 'shrinks' a ball's radius by ρ_O while keeping its center unchanged.[k] We set $s_O(\emptyset) := \emptyset$ and extend s_O to sets T of signed balls by means of $s_O(T) = \{s_O(B) \mid B \in T\}$. From Eq. (8) it follows that dominance and marginal dominance are invariant under shrinking and we get the following

Lemma 4.2. *Let $U \supseteq V$ be two sets of signed balls, O any signed ball. Then $B \in \textsc{b}(U,V)$ if and only if $s_O(B) \in \textsc{b}(s_O(U), s_O(V))$ for any ball B.*

Obviously, also the 'smaller' relation between signed balls is invariant under shrinking, from which we obtain

[k]Actually, s_O only depends on one real number, but in our application, this number will always be the radius of an input ball.

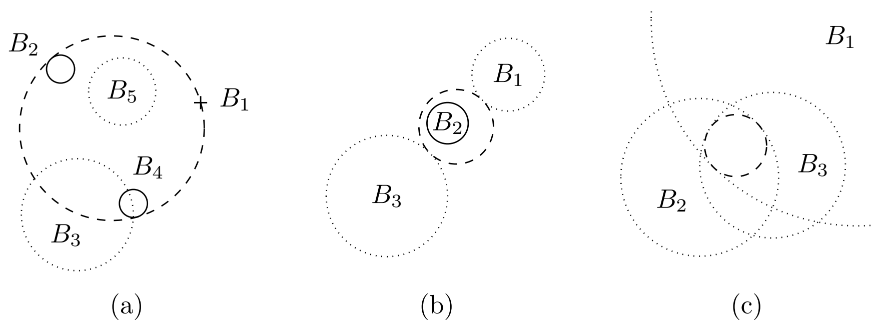

Fig. 13. The miniball $\text{MB}(U)$ (dashed) for three sets U of signed balls: (a) $\text{MB}(U)$ is determined by three positive balls, (b) $\text{MB}(U)$ is determined by two negative balls, (c) the miniball $\text{MB}(U)$ of intersecting, negative balls is the ball of largest volume contained in $\bigcap_{B\in U} B$; its radius is negative.

Corollary 4.3. $\text{MB}(s_O(U), s_O(V)) = s_O(\text{MB}(U,V))$ *for any two sets* $U \supseteq V$ *of signed balls,* O *any signed ball.*

This leads to the important consequence that an instance of SEBB defined by a set of signed balls U has the same combinatorial structure as the instance defined by the balls $s_O(U)$: Most obviously, Corollary 4.3 shows that both instances have the same number of miniballs, the ones in $\text{MB}(s_O(U))$ being shrunken copies of the ones in $\text{MB}(U)$. In fact, replacing the 'positive' concepts of containment and internal tangency with the 'signed' concepts of dominance and marginal dominance in Sec. 2, we can define support sets and bases for sets of signed balls. It then holds (Corollary 4.3) that U and $s_O(U)$ have the same support sets and bases, i.e., the combinatorial structure only depends on parameters which are invariant under shrinking: the ball centers and the differences in radii.

In particular, if the ball O in the corollary is a smallest ball in U then $s_O(U)$ is a set of positive balls, and all material we have developed for this special case in Sec. 2 carries over to the general case (most prominently, this shows that $\text{MB}(U)$ for signed balls U is well-defined, and that SEBB over signed balls is of LP-type and thus solvable by algorithm msw). In this sense, any instance of SEBB over signed balls is combinatorially equivalent to an instance over positive balls, and from now on, we refer to SEBB as the problem of finding $\text{MB}(U)$ for signed balls U.

Reconsidering the situation, it becomes clear that this extension to signed balls is not a real generalization; instead it shows that any instance comes with a 'slider' to simultaneously change all radii.

One very useful slider placement is obtained by shrinking w.r.t. some ball $O \in U$. In this case, we obtain a ball set $s_O(U)$ where at least one ball is a *point*. Consequently, when we solve SEBB using algorithm msw of Fig. 9, we can also assume that the violating ball[1] B entering the basis computation is actually a

[1]A ball now *violates* V if it is not dominated by $\text{MB}(V)$.

point. We can therefore focus on the problem SEBB_p of finding the smallest ball that dominates a set U of signed balls, with an additional point p marginally dominated. More precisely, for given $p \in \mathbb{R}^d$ we define

$$\text{B}_p(U,V) := \text{B}(U \cup \{p\}, V \cup \{p\})$$

and denote the smallest balls in this set by $\text{MB}_p(U,V)$. Then SEBB_p is the problem of finding $\text{MB}_p(U) := \text{MB}_p(U,\emptyset)$ for a given set U of signed balls and a point $p \in \mathbb{R}^d$. Observe that this generalizes the notion $\text{MB}_p(U)$ (Eq. (5)) from only positive balls to signed balls. In contrast to the case of positive balls (Lemma 2.6), the set $\text{MB}_p(U)$ may contain more than one ball when U is a set of signed balls (to see this, shrink Fig. 4 (left) w.r.t. B_2). Throughout the paper, we only compute $\text{MB}_p(U)$ in case the set contains at most one ball. We note that all balls in $\text{B}_p(U,V)$ are positive (they dominate p) and that we can always achieve $p = \mathbf{0}$ through a suitable translation, i.e., reduce problem SEBB_p to $\text{SEBB}_{\mathbf{0}}$.

Using the currently best algorithm for general LP-type problems,[9] the previous discussion yields the following result.

Theorem 4.4. *Problem SEBB over a set of n signed balls can be reduced to problem $\text{SEBB}_{\mathbf{0}}$ over a set of at most $d+1$ signed balls: given an algorithm for the latter problem of runtime $f(d)$, we get an $\mathcal{O}((dn+e^{\sqrt{d\log d}})f(d))$-algorithm for the former.*

In the sequel (Secs. 5 and 6), we concentrate on methods for solving problem $\text{SEBB}_{\mathbf{0}}$ with the goal of improving over the complete enumeration approach which has $f(d) = \Omega(2^d)$. Any ball is assumed to be a signed ball, unless stated otherwise.

5. Inversion

In this section we present a 'dual' formulation of the $\text{SEBB}_{\mathbf{0}}$ problem for (signed) balls. We derive this by employing the *inversion* transform to obtain an 'almost' linear[m] program that describes $\text{MB}_{\mathbf{0}}(U,V)$. This program is also the basis of our approach to small cases of $\text{SEBB}_{\mathbf{0}}$ (Sec. 6). As a by-product, this section links $\text{SEBB}_{\mathbf{0}}$ to the problem of finding the distance from a point to the convex hull of a set of balls.

5.1. *A dual formulation for $\text{SEBB}_{\mathbf{0}}$*

We use the *inversion transform* $x^* := x/\|x\|^2$, $x \neq \mathbf{0}$, to map a ball $B \in \text{B}_{\mathbf{0}}(U,V)$, to some simpler object. To this end, we exploit the fact that under inversion, balls through the origin map to halfspaces while balls not containing the origin simply translate to balls again.

We start by briefly reviewing how balls and halfspaces transform under inversion. For this, we extend the inversion map to point sets via $P^* := \text{cl}(\{p^* \mid p \in P \setminus \{\mathbf{0}\}\})$,

[m] In contrast to the convex but far-from-linear programs obtained by Megiddo and Dyer.[19,6].

where cl(Q) denotes the closure of set Q, and to sets S of balls or halfspaces by means of $S^* := \{P^* \mid P \in S\}$.[n] Consider a halfspace $H \subset \mathbb{R}^d$; H can always be written in the form

$$H = \left\{x \mid v^T x + \alpha \geq 0\right\}, \qquad v^T v = 1. \tag{9}$$

If H does not contain the origin (i.e., $\alpha < 0$) then H maps to the positive ball

$$H^* = B(-v/(2\alpha), -1/(2\alpha)). \tag{10}$$

In this case, the number $-\alpha$ is the distance of the halfspace H to the origin.

Since $(P^*)^* = P$, if P is a ball or halfspace, the converse holds, too: a proper ball with the origin on its boundary transforms to a halfspace not containing the origin. On the other hand, a ball $B = B(c, \rho)$ not containing the origin maps to a ball again, namely to $B^* = B(d, \sigma)$ where

$$d = \frac{c}{c^T c - \rho^2} \quad \text{and} \quad \sigma = \frac{\rho}{c^T c - \rho^2}. \tag{11}$$

B^* again does not contain the origin, and B^* is positive if and only if B is positive. All these facts are easily verified.[2]

The following lemma shows how the dominance relation in the 'primal' domain translates under inversion. For this, we say that a halfspace H of the form (9) *dominates* a ball $B = B(d, \sigma)$ if and only if

$$v^T d + \alpha \geq \sigma, \tag{12}$$

and we speak of *marginal dominance* in case of equality in (12).

As in the primal domain, the dominance relation has an interpretation in terms of containment and intersection: H dominates a positive ball B if and only if H contains B, and H dominates a negative ball B if and only if H intersects B. In both cases, marginal dominance corresponds to B being tangent to the hyperplane underlying H, in addition.

Lemma 5.1. *Let D be a positive ball through $\mathbf{0}$ and B a signed ball not containing $\mathbf{0}$. Then D dominates B if and only if the halfspace D^* dominates the ball B^*.*

Proof. We first show that D dominates B if and only if

$$\|c_D - c_B\|^2 \leq (\rho_D - \rho_B)^2. \tag{13}$$

The direction ($\Rightarrow$) is clear from the definition of dominance, and so is ($\Leftarrow$) under the assumption that $\rho_D - \rho_B \geq 0$. So suppose (13) holds with $\rho_D - \rho_B < 0$. Then

$$0 \leq \|c_D - c_B\| \leq \rho_B - \rho_D,$$

[n]The use of the closure operator guarantees that if P is a ball or halfspace containing $\mathbf{0}$, its image P^* is well-defined and has no 'holes;' we also set $\{\}^* := \{\mathbf{0}\}$ to have $(P^*)^* = P$.

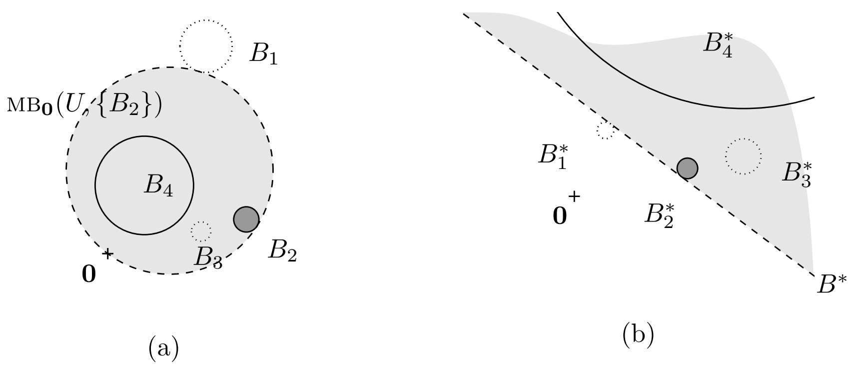

Fig. 14. An example illustrating our use of the inversion transform. (a) A configuration of four circles U, together with $B := \text{MB}_{\mathbf{0}}(U, \{B_2\})$ (dashed). (b) The balls from (a) after inversion: dominance carries over in the sense of Lemma 5.1, so B^* contains B_4^*, intersects B_1^* and B_3^*, and B_2^* is internally tangent to it.

from which we conclude that B is positive and dominates D. Thus, $\mathbf{0} \in D \subseteq B$, a contradiction to B not containing the origin.

It remains to show that Eq. (13) holds if and only if the halfspace D^* dominates the ball B^*. As $c_D^T c_D = \rho_D^2$, the former inequality is equivalent to

$$c_B^T c_B - \rho_B^2 \leq 2\,(c_D^T c_B - \rho_D \rho_B), \tag{14}$$

where the left hand side $\mu := c_B^T c_B - \rho_B^2$ is a strictly positive number.

Write the halfspace D^* in the form (9) with $\alpha < 0$, and observe from (10) and (11) that

$$c_D = -v/(2\alpha), \quad \rho_D = -1/(2\alpha), \quad d = c_B/\mu, \quad \sigma = \rho_B/\mu,$$

for the inverse $B(d, \sigma)$ of B. Using this, we obtain the equivalence of (12) and (14) by multiplying (14) with $\alpha/\mu < 0$. □

For $U \supseteq V$ two sets of balls, we define $\text{H}(U, V)$ to be the set of halfspaces not containing the origin that dominate the balls in U and marginally dominate the balls in V. The following is an immediate consequence of Lemma 5.1. Observe that any ball B satisfying $B \in \text{B}_{\mathbf{0}}(U, V)$ or $B^* \in \text{H}(U^*, V^*)$ is positive by definition.

Lemma 5.2. *Let $U \supseteq V$, $U \neq \emptyset$, be two sets of balls, no ball containing the origin. A ball B lies in $\text{B}_{\mathbf{0}}(U, V)$ if and only if the halfspace B^* lies in $\text{H}(U^*, V^*)$.*

We are interested in *smallest* balls in $\text{B}_{\mathbf{0}}(U, V)$. In order to obtain an interpretation for these in the dual, we use the fact that under inversion, the radius of a ball $B \in \text{B}_{\mathbf{0}}(U, V)$ is inversely proportional to the distance of the halfspace B^* to the origin, see (10). It follows that B is a smallest ball in $\text{B}_{\mathbf{0}}(U, V)$, i.e., $B \in \text{MB}_{\mathbf{0}}(U, V)$, if and only if the halfspace B^* has largest distance to the origin among all halfspaces in $\text{H}(U^*, V^*)$. We call such a halfspace B^* a *farthest* halfspace.

An example of four balls $U = \{B_1, \dots, B_4\}$ is shown in Fig. 14(a), together with the dashed ball $B := \mathrm{MB}_{\mathbf{0}}(U, \{B_2\})$. Positive balls are depicted with solid, negative balls with dotted boundary. Part (b) of the figure depicts the configuration after inversion w.r.t. the origin. The image B^* of B corresponds to the gray halfspace; it is the farthest among the halfspaces which avoid the origin, contain B_4^*, intersect B_1^* and B_3^*, and to which B_2^* is internally tangent.

The previous considerations imply that the following mathematical program searches for the halfspace(s) $\mathrm{MB}_{\mathbf{0}}(U, V)^*$ in the set $\mathrm{H}(U^*, V^*)$.

Corollary 5.3. *Let $U \supseteq V$, $U \neq \emptyset$, be two sets of balls, no ball containing the origin. Consider the program*

$$
\begin{array}{lll}
P_{\mathbf{0}}(U,V) & \textit{minimize} & \alpha \\
 & \textit{subject to} & v^T d_B + \alpha \geq \sigma_B, \quad B \in U \setminus V, \\
 & & v^T d_B + \alpha = \sigma_B, \quad B \in V, \\
 & & v^T v \qquad\quad = 1,
\end{array}
$$

where the d_B and σ_B are the centers and radii of the inverted balls U^*, see Eq. (11). Then $D \in \mathrm{MB}_{\mathbf{0}}(U, V)$ if and only if

$$D^* = \{x \in \mathbb{R}^d \mid \tilde{v}^T x + \tilde{\alpha} \geq 0\}$$

for some optimal solution $(\tilde{v}, \tilde{\alpha})$ to the above program satisfying $\tilde{\alpha} < 0$.

The assumption $U \neq \emptyset$ guarantees $D \neq \mathbf{0}$; if U is empty, program $P_{\mathbf{0}}(U, V)$ consists of a quadratic constraint only and is thus unbounded (the statement is still correct if we define the optimal solution to be $(\mathbf{0}, -\infty)$ in this case).

5.2. *The distance to the convex hull*

With the material from the previous subsection at hand, we can easily relate problems SEBB and $\mathrm{SEBB}_{\mathbf{0}}$ to the problem DHB of finding the distance from a point $p \in \mathbb{R}^d$ to the convex hull

$$\mathrm{conv}(U) := \mathrm{conv}(\textstyle\bigcup_{B \in U} B)$$

of a given set U of positive balls (Fig. 15). W.l.o.g. we may assume that $p = \mathbf{0}$, in which case the problem amounts to finding the minimum-norm point in $\mathrm{conv}(U)$. Also, we can make sure in linear time that no input ball contains the origin.

Observe that a point $q \neq \mathbf{0}$ (which we can always write as $q = -\alpha v$ with $v^T v = 1$ and $\alpha < 0$) is the minimum-norm point in $\mathrm{conv}(U)$ if and only if the halfspace (9) is the farthest halfspace containing the balls in U. By Lemma 2.6 and Corollary 5.3, if the latter halfspace exists, it is unique and can be computed by an algorithm for $\mathrm{SEBB}_{\mathbf{0}}$ over positive balls. (If the algorithm delivers $\mathrm{MB}_{\mathbf{0}}(U^*) = \emptyset$, we know that $\mathbf{0} \in \mathrm{conv}(U)$ and hence $q = \mathbf{0}$ is the minimum-norm point, see Lemma 2.7.) It is clear that we can also solve $\mathrm{SEBB}_{\mathbf{0}}$ for positive balls, given an algorithm for DHB. In this sense, both problems are equivalent.

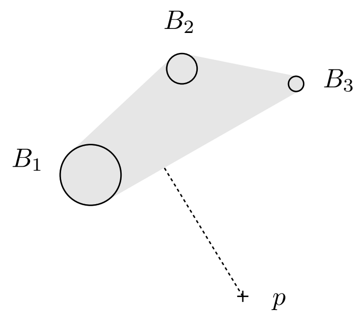

Fig. 15. The DHB problem: find the distance (dotted) from point $p \in \mathbb{R}^d$ to the nearest point in the convex hull conv(U) of the positive balls U (gray). In this example, $U = \{B_1, B_2, B_3\}$.

In addition, we can also solve SEBB for signed balls with an algorithm $\mathcal{D}$ for DHB.[o] To obtain $\text{MB}(U)$, we guess the smallest (possibly negative) ball $O \in U$ marginally dominated by $\text{MB}(U)$, find the set U' of positive shrunken balls,

$$U' = \{B \in s_O(U \setminus \{O\}) \mid B \text{ positive}\},$$

and compute $\text{MB}_{c_O}(U') = s_O(\text{MB}(U))$ from which the ball $\text{MB}(U)$ is easily reconstructed (Corollary 4.3). (If our guess was correct, the shrunken balls with negative radius come from balls that do not contribute to $\text{MB}(U)$ by Lemma 2.3.) As long as U is a small set, the at most $|U|$ guesses introduce a negligible polynomial overhead. For large input sets however, a direct application of the reduction leads to an unnecessarily slow algorithm, and thus it pays off to run algorithm msw and use the reduction for the small cases only (where $|U| \leq d+2$).

6. Small cases

We have shown in Sec. 4 that problem SEBB can be reduced to the problem $\text{SEBB}_{\mathbf{0}}$ of computing $\text{MB}_{\mathbf{0}}(T)$, for T some set of signed balls, $|T| \leq d+1$.

Using the fact that we now have the *origin* fixed on the boundary we can improve over the previous complete enumeration approach, by using inversion and the concept of *unique sink orientations*.[24]

In the sequel, we assume that T is a set of signed balls with linearly independent centers,[p] none of them containing the origin. The latter assumption is satisfied in our application, where $\text{MB}_{\mathbf{0}}(T)$ is needed only during the basis computation of algorithm msw (Fig. 9). The linear independence assumption is no loss of generality,

[o] An entirely different reduction from SEBP to DHP, the problem of finding the distance from a given point to the convex hull of a set of *points*, is well-known[7,22] but only works for small cases and does not generalize to SEBB.

[p] For this, we interpret the centers as vectors, which is quite natural because of the translation employed in the reduction from SEBB_p to $\text{SEBB}_{\mathbf{0}}$.

because we can embed the balls into $\mathbb{R}^{d+1}$ and symbolically perturb them; in fact, this is easy if T arises as the set V during the basis computation $\mathsf{basis}(V, B)$ of the algorithm msw.[11]

Our method for finding $\text{MB}_{\mathbf{0}}(T)$ computes as intermediate steps balls of the form

$$\text{MB}_{\mathbf{0}}(U, V) = \text{MB}(U \cup \{p\}, V \cup \{p\}),$$

for $V \subseteq U \subseteq T$. One obstacle we have to overcome for this is the possible nonexistence of $\text{MB}_{\mathbf{0}}(U, V)$: take for instance a positive ball B not containing the origin and place a positive ball B' into $\text{conv}(B \cup \{\mathbf{0}\})$, $U = \{B, B'\}$ and $V = \{B'\}$.[q] Our solution employs the inversion transform: it defines for all pairs (U, V) a 'generalized ball' $\text{GMB}_{\mathbf{0}}(U, V)$ which coincides with $\text{MB}_{\mathbf{0}}(U, V)$ if the latter exists.

Performing inversion as described in the previous section gives us $|T| \leq d$ balls T^* with centers d_B and radii σ_B, $B \in T$, as in (11). The latter equation also shows that the d_B are linearly independent. The following lemma is then an easy consequence of previous considerations.

Lemma 6.1. *For given $V \subseteq U \subseteq T$ with $U \neq \emptyset$, consider the following (nonconvex) optimization problem in the variables $v \in \mathbb{R}^d$, $\alpha \in \mathbb{R}$.*

$$\begin{array}{llll} \mathcal{P}_{\mathbf{0}}(U, V) & \textit{lexmin} & (v^T v, \alpha), & \\ & \textit{subject to} & v^T d_B + \alpha \geq \sigma_B, & B \in U \setminus V, \\ & & v^T d_B + \alpha = \sigma_B, & B \in V, \\ & & v^T v \quad\quad \geq 1. & \end{array}$$

(i) $\mathcal{P}_{\mathbf{0}}(U, V)$ has a unique optimal solution $(\tilde{v}, \tilde{\alpha})$.
(ii) $D \in \text{MB}_{\mathbf{0}}(U, V)$ if and only if

$$D^* = \{x \in \mathbb{R}^d \mid \tilde{v}^T x + \tilde{\alpha} \geq 0\}$$

for the optimal solution $(\tilde{v}, \tilde{\alpha})$ to the above program with $\tilde{v}^T \tilde{v} = 1$ and $\tilde{\alpha} < 0$.

In particular, part (i) implies that the set $\text{MB}_{\mathbf{0}}(U, V)$ contains *at most one ball* whenever the balls in U do not contain the origin and their centers are affinely independent.

Proof. (i) If we can show that $\mathcal{P}_{\mathbf{0}}(U, V)$ has a feasible solution then it also has an optimal solution, again using a compactness argument (this requires $U \neq \emptyset$). To construct a feasible solution, we first observe that by linear independence of the d_B, the system of equations

$$v^T d_B + \alpha = \sigma_B, \quad B \in U$$

has a solution v for any given α; moreover, if we choose α large enough, any corresponding v must satisfy $v^T v \geq 1$, in which case (v, α) is a feasible solution.

[q] Such a configuration may turn up in our application.

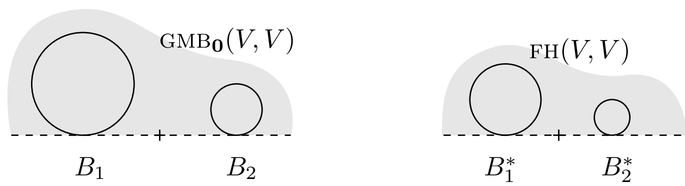

Fig. 16. Two balls $V = \{B_1, B_2\}$ (left) and their images under inversion (right). In this example, the value $(\tilde{v}, \tilde{\alpha})$ of (V,V) has $\tilde{\alpha} = 0$, in which case the 'generalized ball' $\textsc{gmb}_{\mathbf{0}}(V,V)$ is a halfspaces.

To prove uniqueness of the optimal solution, we again invoke linear independence of the d_B and derive the existence of a vector w (which we call an *unbounded direction*) such that

$$w^T d_B = 1, \quad B \in U. \tag{15}$$

Now assume that $\mathcal{P}_{\mathbf{0}}(U,V)$ has two distinct optimal solutions $(\tilde{v}_1, \tilde{\alpha})$, $(\tilde{v}_2, \tilde{\alpha})$ with $\tilde{v}_1^T \tilde{v}_1 = \tilde{v}_2^T \tilde{v}_2 = \delta \geq 1$. Consider any proper convex combination v of $\tilde{v}_1$ and $\tilde{v}_2$; v satisfies $v^T v < \delta$. Then there is a suitable positive constant Θ such that $(v + \Theta w)^T (v + \Theta w) = \delta$, and hence the pair $(v + \Theta w, \tilde{\alpha} - \Theta)$ is a feasible solution for $\mathcal{P}_{\mathbf{0}}(U,V)$, a contradiction to lexicographic minimality of the initial solutions.

(ii) Under $\tilde{v}^T \tilde{v} = 1$, this is equivalent to the statement of Corollary 5.3(ii). □

Even if $\textsc{mb}_{\mathbf{0}}(U,V) = \emptyset$, program $\mathcal{P}_{\mathbf{0}}(U,V)$ has a unique solution, and we call it the *value of* (U,V).

Definition 6.2. For $U \supseteq V$ with U nonempty, the *value of* (U,V), denoted by $\textsc{val}(U,V)$, is the unique solution $(\tilde{v}, \tilde{\alpha})$ of program $\mathcal{P}_{\mathbf{0}}(U,V)$, and we define $\textsc{val}(\emptyset,\emptyset) := (\mathbf{0}, -\infty)$. Moreover, we call the halfspace

$$\textsc{fh}(U,V) := \{x \in \mathbb{R}^d \mid \tilde{v}^T x + \tilde{\alpha} \geq 0\},$$

the *farthest (dual) halfspace* of (U,V). In particular, $\textsc{fh}(\emptyset,\emptyset) = \emptyset$.

The farthest halfspace of (U,V) has a meaningful geometric interpretation even if $\textsc{mb}_{\mathbf{0}}(U,V) = \emptyset$. If the value $(\tilde{v}, \tilde{\alpha})$ of (U,V) satisfies $\tilde{v}^T \tilde{v} = 1$, we already know that $\textsc{fh}(U,V)$ dominates the balls in U and marginally dominates the balls in V, see Eq. (12). If on the other hand $\tilde{v}^T \tilde{v} > 1$, it is easy to see that the halfspace $\textsc{fh}(U,V)$ dominates the *scaled* balls

$$B(d_B, \sigma_B/\sqrt{\tau}) \qquad \text{with} \qquad \tau := \tilde{v}^T \tilde{v}, \tag{16}$$

for $B \in U$, and marginally dominates the scaled versions of the balls in V^* (divide the linear constraints of program $\mathcal{P}_{\mathbf{0}}(U,V)$ by $\sqrt{\tau}$).

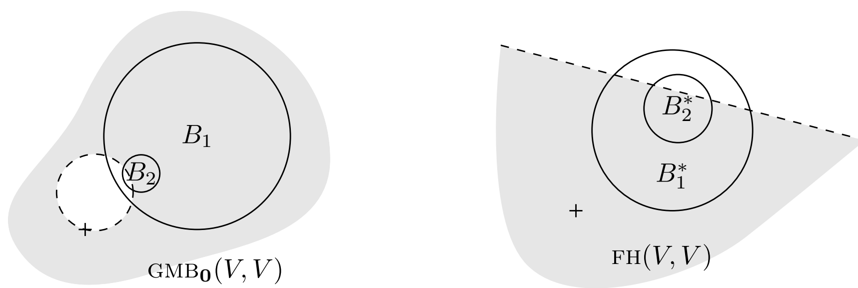

Fig. 17. Two positive balls $V = \{B_1, B_2\}$ (left) and their images under inversion (right). The value $(\tilde{v}, \tilde{\alpha})$ of (V, V) has $\tilde{v}^T\tilde{v} > 1$ and the 'generalized ball' $\text{GMB}_{\mathbf{0}}(V, V)$ is not tangent to the balls V^*.

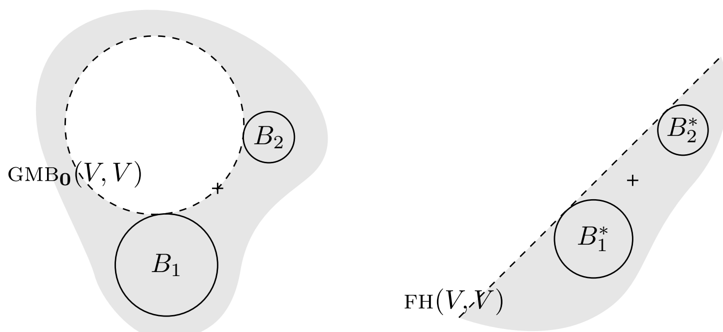

Fig. 18. Two positive balls $V = \{B_1, B_2\}$ (left) and their images under inversion (right). In this case, the value $(\tilde{v}, \tilde{\alpha})$ has $\tilde{v}^T\tilde{v} = 1$ but $\tilde{\alpha} > 0$, i.e., the balls do *not* admit a ball $\text{MB}_{\mathbf{0}}(V, V)$. Sill, all balls V are 'internally' tangent to the 'generalized ball' $\text{GMB}_{\mathbf{0}}(V, V)$.

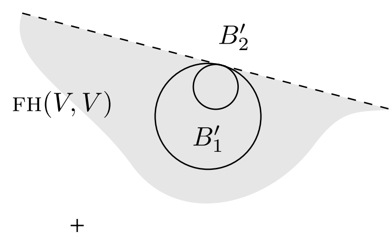

Fig. 19. The *scaled* balls $\{B_1', B_2'\}$, obtained from the balls $V^* = \{B_1^*, B_2^*\}$ in Fig. 17 by scaling their radii with $1/\sqrt{\tau}$, $\tau = \tilde{v}^T\tilde{v}$, are marginally dominated by $\text{FH}(V, V)$.

For an interpretation of $\text{FH}(U, V)$ in the primal, we associate to the pair (U, V) the 'generalized ball'

$$\text{GMB}_{\mathbf{0}}(U, V) := \text{FH}(U, V)^*,$$

which in general need *not* be a ball, as we will see. However, in the geometrically interesting case when the set $\text{MB}_{\mathbf{0}}(U, V)$ is nonempty, it follows from Lemma 6.1(ii) that $\text{GMB}_{\mathbf{0}}(U, V) = \text{MB}_{\mathbf{0}}(U, V)$. Recall that this occurs precisely if the value $(\tilde{v}, \tilde{\alpha})$ of the pair (U, V) fulfills $\tilde{v}^T \tilde{v} = 1$ and $\tilde{\alpha} < 0$.

In general, $\text{GMB}_{\mathbf{0}}(U, V)$ can be a ball, the complement of an open ball, or a halfspace. In case $\tilde{\alpha} > 0$, the halfspace $\text{FH}(U, V)$ contains the origin and hence $\text{GMB}_{\mathbf{0}}(U, V)$ is the complement of an open ball through the origin. If $\tilde{\alpha} = 0$ then $\text{FH}(U, V)$ goes through the origin, and inversion does not provide us with a ball $\text{GMB}_{\mathbf{0}}(U, V)$ but with a halfspace instead (Fig. 16). We remark that if $\tilde{v}^T \tilde{v} > 1$, $\text{GMB}_{\mathbf{0}}(U, V)$ will not even be tangent to the proper balls in V (Fig. 17).

In Fig. 18, the inverted balls V^* do not admit a halfspace that avoids the origin. Hence program $P_{\mathbf{0}}(V, V)$ has no solution, implying $\text{MB}_{\mathbf{0}}(V, V) = \emptyset$. In order to obtain $\text{GMB}_{\mathbf{0}}(V, V)$, we have to solve program $\mathcal{P}_{\mathbf{0}}(V, V)$. For this, we observe that the balls V^* admit two tangent hyperplanes, i.e., there are two halfspaces, parameterized by v and α, which satisfy the equality constraints of $\mathcal{P}_{\mathbf{0}}(V, V)$ with $v^T v = 1$. Since the program in this case *minimizes* the distance to the halfspace, $\text{FH}(V, V)$ is the enclosing halfspace corresponding to the 'upper' hyperplane in the figure (painted in gray). Since it contains the origin, $\text{GMB}_{\mathbf{0}}(V, V)$ is the complement of a ball. Finally, Fig. 19 depicts the scaled versions (16) of the balls V^* from Fig. 17. Indeed, $\text{FH}(V, V)$ marginally dominates these balls.[r]

We now investigate program $\mathcal{P}_{\mathbf{0}}(U, V)$ further. Although it is not a convex program, it turns out to be equivalent to one of two related convex programs. Program $\mathcal{C}'_{\mathbf{0}}(U, V)$ below finds the lowest point in a cylinder, subject to linear (in)equality constraints. In case it is infeasible (which will be the case if and only if $\text{MB}_{\mathbf{0}}(U, V) = \emptyset$), the other program $\mathcal{C}_{\mathbf{0}}(U, V)$ applies in which the cylinder is allowed to enlarge until the feasible region becomes non-empty.

Lemma 6.3. *Let $(\tilde{v}, \tilde{\alpha})$ be the optimal solution to program $\mathcal{P}_{\mathbf{0}}(U, V)$, for $U \neq \emptyset$, and let γ be the minimum value of the convex quadratic program*

$$\begin{array}{lll} \mathcal{C}_{\mathbf{0}}(U, V) & \textit{minimize} & v^T v \\ & \textit{subject to} & v^T d_B + \alpha \geq \sigma_B, \quad B \in U \setminus V, \\ & & v^T d_B + \alpha = \sigma_B, \quad B \in V. \end{array}$$

(i) Program $\mathcal{C}_{\mathbf{0}}(U, V)$ has a unique optimal solution, provided $V \neq \emptyset$.
(ii) If $\gamma \geq 1$ then $(\tilde{v}, \tilde{\alpha})$ is the unique optimal solution to $\mathcal{C}_{\mathbf{0}}(U, V)$.

[r]Since scaled balls do not invert to scaled balls in general—the centers may move—the situation is more complicated in the primal.

(iii) If $\gamma < 1$ then $\tilde{v}^T\tilde{v} = 1$ and $(\tilde{v}, \tilde{\alpha})$ is the unique optimal solution to the convex program

$$\begin{array}{lll} \mathcal{C}'_{\mathbf{0}}(U,V) & \textit{minimize} & \alpha \\ & \textit{subject to} & v^T d_B + \alpha \geq \sigma_B, \quad B \in U \setminus V, \\ & & v^T d_B + \alpha = \sigma_B, \quad B \in V, \\ & & v^T v \qquad\quad \leq 1. \end{array}$$

Also, $\mathcal{C}_{\mathbf{0}}(U,V)$ is strictly feasible (i.e., feasible values exist that satisfy all inequality constraints with strict inequality). If $\gamma < 1$, $\mathcal{C}'_{\mathbf{0}}(U,V)$ is strictly feasible, too.

Proof. (i) A compactness argument shows that some optimal solutions exists. Moreover, $\mathcal{C}_{\mathbf{0}}(U,V)$ has a unique optimal vector $\tilde{v}$ because any proper convex combination of two different optimal vectors would still be feasible with smaller objective function value. The optimal $\tilde{v}$ uniquely determines α because $\mathcal{C}_{\mathbf{0}}(U,V)$ has at least one equality constraint.

(ii) Under $\gamma \geq 1$, $(\tilde{v}, \tilde{\alpha})$ is an optimal solution to $\mathcal{C}_{\mathbf{0}}(U,V)$ and by (i) it is the unique one because $\gamma \geq 1$ implies $V \neq \emptyset$.

(iii) Under $\gamma < 1$, $\mathcal{C}'_{\mathbf{0}}(U,V)$ is feasible and a compactness argument shows that an optimal solution $(\tilde{v}', \tilde{\alpha}')$ exists. Using the unbounded direction (15) again, $\tilde{v}'^T\tilde{v}' = 1$ and the uniqueness of the optimal solution can be established. Because $(\tilde{v}', \tilde{\alpha}')$ is feasible for $\mathcal{P}_{\mathbf{0}}(U,V)$, we have $\tilde{v}^T\tilde{v} = 1$, and from lexicographic minimality of $(\tilde{v}, \tilde{\alpha})$, $(\tilde{v}, \tilde{\alpha}) = (\tilde{v}', \tilde{\alpha}')$ follows.

To see strict feasibility of $\mathcal{C}'_{\mathbf{0}}(U,V)$, first note that $\gamma < 1$ implies the existence of a feasible pair (v, α) for which $v^T v < 1$. Linear independence of the d_B yields a vector w such that

$$w^T d_B = \begin{cases} 1, & B \in U \setminus V, \\ 0, & B \in V \end{cases}.$$

For sufficiently small $\Theta > 0$, the pair $(v + \Theta w, \alpha)$ is strictly feasible for $\mathcal{C}'_{\mathbf{0}}(U,V)$. Strict feasibility of $\mathcal{C}_{\mathbf{0}}(U,V)$ follows by an even simpler proof along these lines. □

This shows that given the minimum value γ of program $\mathcal{C}_{\mathbf{0}}(U,V)$, the solution of program $\mathcal{P}_{\mathbf{0}}(U,V)$ can either be read off program $\mathcal{C}_{\mathbf{0}}(U,V)$ (in case $\gamma \geq 1$) or program $\mathcal{C}'_{\mathbf{0}}(U,V)$ (in case $\gamma \leq 1$).

The next step is to characterize the optimal solutions of the programs $\mathcal{C}_{\mathbf{0}}(U,V)$ and $\mathcal{C}'_{\mathbf{0}}(U,V)$, based on the *Karush-Kuhn-Tucker Theorem* for convex programming. Unifying both characterizations, we obtain necessary and sufficient optimality conditions for program $\mathcal{P}_{\mathbf{0}}(U,V)$. To this end, we invoke the following version of the Karush-Kuhn-Tucker Theorem which is a specialization of a general result (Theorem 5.3.1 and Theorem 4.3.8 with Slater's constraint qualification in the book by Bazaraa, Sherali & Shetty[1]).

Theorem 6.4. *Let $f, g_1, \ldots, g_m$ be differentiable convex functions, let $a_1, \ldots, a_\ell \in \mathbb{R}^n$ be linearly independent vectors, and let $\beta_1, \ldots, \beta_\ell$ be real numbers. Consider the optimization problem*

$$\begin{array}{lll} \textit{minimize} & f(x), & \\ \textit{subject to} & g_i(x) \leq 0, & i = 1, \ldots, m, \\ & a_i^T x \ = \beta_i, & i = 1, \ldots, \ell. \end{array} \tag{17}$$

(i) If $\tilde{x}$ is an optimal solution to (17) and if there exists a vector $\tilde{y}$ such that

$$g_i(\tilde{y}) < 0, \quad i = 1, \ldots, m,$$
$$a_i^T \tilde{y} = \beta_i, \quad i = 1, \ldots, \ell,$$

then there are real numbers $\mu_1, \ldots, \mu_m$ and $\lambda_1, \ldots, \lambda_\ell$ such that

$$\mu_i \geq 0, \quad i = 1, \ldots, m, \tag{18}$$
$$\mu_i g_i(\tilde{x}) = 0, \quad i = 1, \ldots, m, \tag{19}$$
$$\nabla f(\tilde{x}) + \sum_{i=1}^{m} \mu_i \nabla g_i(\tilde{x}) + \sum_{i=1}^{\ell} \lambda_i a_i^T = 0. \tag{20}$$

(ii) Conversely, if $\tilde{x}$ is a feasible solution to program (17) such that numbers satisfying (18), (19) and (20) exist then $\tilde{x}$ is an optimal solution to (17).

Applied to our two programs, we obtain the following optimality conditions.

Lemma 6.5. *Let $V \subseteq U \subseteq T$ with $U \neq \emptyset$.*

(i) A feasible solution $(\tilde{v}, \tilde{\alpha})$ for $\mathcal{C}_{\mathbf{0}}(U, V)$ is optimal if and only if there exist real numbers λ_B, $B \in U$, such that

$$\lambda_B \geq 0, \quad B \in U \setminus V$$
$$\lambda_B(\tilde{v}^T d_B + \tilde{\alpha} - \sigma_B) = 0, \quad B \in U \setminus V,$$
$$\textstyle\sum_{B \in U} \lambda_B d_B = \tilde{v}, \tag{21}$$
$$\textstyle\sum_{B \in U} \lambda_B = 0. \tag{22}$$

(ii) A feasible solution $(\tilde{v}, \tilde{\alpha})$ to $\mathcal{C}'_{\mathbf{0}}(U, V)$ satisfying $\tilde{v}^T \tilde{v} = 1$ is optimal if there exist real numbers λ_B, $B \in U$, such that

$$\lambda_B \geq 0, \quad B \in U \setminus V \tag{23}$$
$$\lambda_B(\tilde{v}^T d_B + \tilde{\alpha} - \sigma_B) = 0, \quad B \in U \setminus V, \tag{24}$$
$$\textstyle\sum_{B \in U} \lambda_B d_B = \tilde{v}, \tag{25}$$
$$\textstyle\sum_{B \in U} \lambda_B > 0. \tag{26}$$

Conversely, if $(\tilde{v}, \tilde{\alpha})$ is an optimal solution to $\mathcal{C}'_{\mathbf{0}}(U, V)$ and the minimum value γ of program $\mathcal{C}_{\mathbf{0}}(U, V)$ fulfills $\gamma < 1$ then there exist real numbers λ_B, $B \in U$, such that (23), (24), (25) and (26) hold.

In both cases, the λ_B are uniquely determined by $\tilde{v}$ via linear independence of the d_B.

From these sets of conditions for $\mathcal{C}_{\mathbf{0}}(U,V)$ and $\mathcal{C}'_{\mathbf{0}}(U,V)$ we can derive optimality conditions for the nonconvex program $\mathcal{P}_{\mathbf{0}}(U,V)$.

Theorem 6.6. *A feasible solution $(\tilde{v},\tilde{\alpha})$ for program $\mathcal{P}_{\mathbf{0}}(U,V)$ is optimal if and only if there exist real numbers λ_B, $B \in U$, with $\mu := \sum_{B\in U}\lambda_B$ such that*

$$\begin{aligned} \lambda_B &\geq 0, \quad B \in U \setminus V \\ \mu &\geq 0, \\ \lambda_B(\tilde{v}^T d_B + \tilde{\alpha} - \sigma_B) &= 0, \quad B \in U \setminus V, \\ \mu\,(\tilde{v}^T\tilde{v} - 1) &= 0, \\ \textstyle\sum_{B\in U}\lambda_B d_B &= \tilde{v}. \end{aligned} \tag{27}$$

Proof. The direction $(\Rightarrow)$ follows through Lemmata 6.3 and 6.5, so it remains to settle $(\Leftarrow)$. For this, we distinguish two cases, depending on the minimum value γ of program $\mathcal{C}_{\mathbf{0}}(U,V)$.

Consider the case $\gamma < 1$ first. If $\sum_{B\in U}\lambda_B = 0$ then Lemma 6.5 shows that $(\tilde{v},\tilde{\alpha})$, which is clearly feasible for $\mathcal{C}_{\mathbf{0}}(U,V)$, is optmial to $\mathcal{C}_{\mathbf{0}}(U,V)$; hence $\gamma = \tilde{v}^T\tilde{v} \geq 1$, a contradiction. Thus $\sum_{B\in U}\lambda_B > 0$, which by (27) implies $\tilde{v}^T\tilde{v} = 1$. So $(\tilde{v},\tilde{\alpha})$ is feasible and optimal to $\mathcal{C}'_{\mathbf{0}}(U,V)$, which together with Lemma 6.3(iii) establishes the claim.

In case $\gamma \geq 1$ the argument is as follows. If $\sum_{B\in U}\lambda_B = 0$ holds, the Lemmata 6.5 and 6.3(ii) certify that the solution $(\tilde{v},\tilde{\alpha})$ is optimal to $\mathcal{P}_{\mathbf{0}}(U,V)$. If on the other hand $\sum_{B\in U}\lambda_B > 0$ then $\tilde{v}^T\tilde{v} = 1$ by (27), from which we derive via an invocation of Lemma 6.5(ii) that $(\tilde{v},\tilde{\alpha})$ is optimal to program $\mathcal{C}'_{\mathbf{0}}(U,V)$. Moreover, $(\tilde{v},\tilde{\alpha})$ must also be optimal to $\mathcal{P}_{\mathbf{0}}(U,V)$ because the optimal solution to program $\mathcal{P}_{\mathbf{0}}(U,V)$ satisfies the constraint $v^Tv \geq 1$ with equality (recall $\tilde{v}^T\tilde{v} = 1$) and hence is feasible to program $\mathcal{C}'_{\mathbf{0}}(U,V)$. □

As promised, we can state a version of Welzl's Lemma.[25] We prepare this by presenting the statement in the dual space, i.e., in terms of values of pairs (U,V) and associated halfspaces $\textsc{fh}(U,V)$.

Lemma 6.7. *Let $V \subseteq U \subseteq T$ and $B \in U \setminus V$. Denote by $(\tilde{v},\tilde{\alpha})$ the value of the pair $(U \setminus \{B\}, V)$. Then*

$$\textsc{val}(U,V) = \begin{cases} \textsc{val}(U\setminus\{B\},V), & \text{if } \tilde{v}^T d_B + \tilde{\alpha} \geq \sigma_B, \\ \textsc{val}(U, V\cup\{B\}), & \text{otherwise.} \end{cases}$$

As the value of a pair uniquely determines its associated farthest halfspace, the lemma holds also for farthest halfspaces (i.e., if we replace 'VAL' by 'FH' in the lemma). In this case, we obtain the following geometric interpretation. The halfspace $\textsc{fh}(U,V)$ coincides with the halfspace $\textsc{fh}(U\setminus\{B\},V)$ if the latter dominates the scaled version (16) of ball B^*, and equals the halfspace $\textsc{fh}(U,V\cup\{B\})$ otherwise.

Proof. The case $U = \{B\}$ is easily checked directly, so assume $|U| > 1$. If $\tilde{v}^T d_B + \tilde{\alpha} \geq \sigma_B$ then $(\tilde{v}, \tilde{\alpha})$ is feasible and hence optimal to the more restricted problem $\mathcal{P}_{\mathbf{0}}(U, V)$, and $\text{VAL}(U, V) = \text{VAL}(U \setminus \{B\}, V)$ follows. Otherwise, the value $(\tilde{v}', \tilde{\alpha}')$ of (U, V) is different from $(\tilde{v}, \tilde{\alpha})$. Now consider the coefficient λ'_B resulting from the application of Theorem 6.6 to $(\tilde{v}', \tilde{\alpha}')$. We must have $\lambda'_B \neq 0$, because Theorem 6.6 would otherwise certify that $(\tilde{v}', \tilde{\alpha}') = \text{VAL}(U \setminus \{B\}, V)$. This, however, implies that

$$\tilde{v}'^T d_B + \tilde{\alpha}' = \sigma_B,$$

from which we conclude $\text{VAL}(U, V) = \text{VAL}(U, V \cup \{B\})$. □

Here is the fix for Dilemma 3.2 in case of affinely independent centers.

Lemma 6.8. *Let $V \subseteq U$, where U is any set of signed balls with affinely independent centers, and assume $\text{MB}(U, V) \neq \emptyset$. Then the sets $\text{MB}(U, V)$ and $\text{MB}(U \setminus \{B\}, V)$ are singletons, for any $B \in U \setminus V$. Moreover, if no ball in V is dominated by another ball in U, and if*

$$B \text{ is not dominated by } \text{MB}(U \setminus \{B\}, V), \tag{28}$$

for some $B \in U \setminus V$, then $\text{MB}(U, V) = \text{MB}(U, V \cup \{B\})$, and B is not dominated by another ball in U, either.

It easily follows by induction that Welzl's algorithm (Fig. 6 in Sec. 3, with the test '$B \not\subseteq D$' replaced by 'B not dominated by D') computes $\text{MB}(U)$ for a set of signed balls, provided the centers of the input balls are affinely independent (a perturbed embedding into $\mathbb{R}^{|U|-1}$ always accomplishes this). No other preconditions are required; in particular, balls can overlap in an arbitrary fashion.

Proof. For $V = \emptyset$, this is Lemma 2.3, with the obvious generalization to signed balls (refer to the discussion after Corollary 4.3). For all V, transitivity of the dominance relation shows that if B is not dominated by $\text{MB}(U \setminus \{B\}, V)$, it cannot be dominated by a ball in $U \setminus \{B\}$, either.

In case $V \neq \emptyset$, we fix any ball $O \in V$ and may assume—after a suitable translation and a shrinking step—that $O = \mathbf{0}$; Eq. (28) is not affected by this. Moreover, we can assume that O does not dominate any other (negative) ball in $U \setminus V$: such a ball can be removed from consideration (and added back later), without affecting the miniball (here, we again use transitivity of dominance).

Then, no ball in U contains $O = \mathbf{0}$, and the centers of the balls

$$U' = U \setminus \{O\}$$

are *linearly independent*. Under (28), we have $B \in U'$. Therefore, we can apply our previous machinery. Setting

$$V' = V \setminus \{O\},$$

Lemma 6.1 yields that the two sets $\text{MB}(U, V) = \text{MB}_{\mathbf{0}}(U', V')$ and $\text{MB}(U \setminus \{B\}, V) = \text{MB}_{\mathbf{0}}(U' \setminus \{B\}, V')$ contain at most one ball each. Also, the assumption $\text{MB}(U, V) \neq \emptyset$ implies $\text{MB}(U \setminus \{B\}, V) \neq \emptyset$ (this is easily verified using the program in Lemma 6.1). Consequently, the ball sets are singletons.

Now let $(\tilde{v}, \tilde{\alpha})$ be the value of pair $(U' \setminus \{B\}, V')$. $\text{MB}_{\mathbf{0}}(U' \setminus \{B\}, V') \neq \emptyset$ implies $\tilde{v}^T \tilde{v} = 1$ (Lemma 6.1). Then, Lemma 5.1 shows that B is not dominated by the ball $\text{MB}_{\mathbf{0}}(U' \setminus \{B\}, V')$ if and only if $\tilde{v}^T d_B + \tilde{\alpha} < \sigma_B$ holds, for d_B, σ_B being center and radius of the inverted ball B^*. Lemma 6.7 in turn implies

$$\text{VAL}(U', V') = \text{VAL}(U', V' \cup \{B\}), \tag{29}$$

and $\text{FH}(U', V') = \text{FH}(U', V' \cup \{B\})$ along with $\text{GMB}_{\mathbf{0}}(U', V') = \text{GMB}_{\mathbf{0}}(U', V' \cup \{B\})$ follows. By assumption, the former 'generalized ball' coincides with $\text{MB}_{\mathbf{0}}(U', V')$, from which it follows that the value $(\tilde{v}', \tilde{\alpha}')$ of (U', V') fulfills $\tilde{v}'^T \tilde{v}' = 1$ and $\tilde{\alpha}' < 0$ (Lemma 6.1). By (29), this shows that $\text{GMB}_{\mathbf{0}}(U', V' \cup \{B\}) = \text{MB}_{\mathbf{0}}(U', V' \cup \{B\})$, which establishes the lemma. □

6.1. *The unique sink orientation*

In this last part we want to use the results developed so far to reduce the problem of finding $\text{MB}_{\mathbf{0}}(T)$ to the problem of finding the sink in a unique sink orientation. To this end, we begin with a brief recapitulation of unique sink orientations and proceed with the presentation of our orientation.

As in the previous subsection, we consider a set T of $m \leq d$ balls such that the centers of T are linearly independent and such that no ball in T contains the origin. Consider the m-dimensional cube. Its vertices can be identified with the subsets $J \subseteq T$; faces of the cube then correspond to *intervals* $[V, U] := \{J \mid V \subseteq J \subseteq U\}$, where $V \subseteq U \subseteq T$. We consider the *cube graph*

$$\mathcal{G} = (2^T, \{\{J, J \oplus \{B\}\} \mid J \in 2^T, B \in T\}),$$

where $\oplus$ denotes symmetric difference. An orientation $\mathcal{O}$ of the edges of $\mathcal{G}$ is called a *unique sink orientation* (USO) if for any nonempty face $[V, U]$, the subgraph of $\mathcal{G}$ induced by the vertices of $[V, U]$ has a unique sink w.r.t. $\mathcal{O}$.[24]

As before, we write d_B and σ_B for the center and radius of the inverted balls $B^* \in T^*$, see (11). The following is the main result of this section.

Theorem 6.9. *Consider the orientation $\mathcal{O}$ of $\mathcal{G}$ defined by*

$$J \to J \cup \{B\} \quad :\Leftrightarrow \quad \text{VAL}(J, J) \neq \text{VAL}(J \cup \{B\}, J).$$

Then $\mathcal{O}$ is a USO, and the sink S of the cube is a basis of T, meaning that S is inclusion-minimal with $\text{VAL}(S, S) = \text{VAL}(T, \emptyset)$.

In terms of halfspaces $\text{FH}(U, V)$, we can interpret this as follows. The edge $\{J, J \cup \{B\}\}$ is directed towards the larger set if and only if the halfspace $\text{FH}(J, J)$ does not dominate the scaled version (16) of ball B^*.

Proof. Fix a face $[V, U]$ of the cube. If $U = \emptyset$, the face is a vertex and has obviously a unique sink. Otherwise let λ_B, $B \in U$, be the unique multipliers guaranteed by Theorem 6.6 for the value $(\tilde{v}, \tilde{\alpha})$ of (U, V). Let $S \in [V, U]$. We claim that S is a sink in $[V, U]$ if and only if S coincides with the set

$$V \cup \{B \in U \setminus V \mid \lambda_B > 0\}, \tag{30}$$

which implies uniqueness of the sink.

First assume that S equals the set in (30). Consider sets V', U' with $V \subseteq V' \subseteq S \subseteq U' \subseteq U$. By definition of S and Theorem 6.6,

$$\begin{aligned} \tilde{v}^T d_B + \tilde{\alpha} = \sigma_B, &\quad B \in V', \\ \lambda_B = 0, &\quad B \notin U', \end{aligned}$$

which—again using Theorem 6.6—implies that $(\tilde{v}, \tilde{\alpha})$ is both feasible and optimal for $\mathcal{P}_{\mathbf{0}}(U', V')$. In particular, we obtain

$$\text{VAL}(S, S) = \text{VAL}(S \cup \{B\}, S), \quad B \in U \setminus S, \tag{31}$$

and

$$\text{VAL}(S, S) = \text{VAL}(S, S \setminus \{B\}), \quad B \in S \setminus V,$$

where all values coincide with $(\tilde{v}, \tilde{\alpha})$. Because the λ_B are determined by $\tilde{v}$, and $\lambda_B \neq 0$ for $B \in S \setminus V$, we must have

$$\text{VAL}(S, S \setminus \{B\}) \neq \text{VAL}(S \setminus \{B\}, S \setminus \{B\}), \quad B \in S \setminus V. \tag{32}$$

Equations (31) and (32) together show that S is a sink in $[V, U]$ w.r.t. $\mathcal{O}$.

For the other direction, let S be any sink in $[V, U]$, let $(\tilde{v}', \tilde{\alpha}')$ be the value of (S, S), and let λ'_B, $B \in S$, be the multipliers proving this according to Theorem 6.6. We first show that $(\tilde{v}', \tilde{\alpha}') = \text{VAL}(U, V)$. S being a sink is equivalent to (31) and (32). From (31) it follows that $(\tilde{v}', \tilde{\alpha}')$ is feasible for $\mathcal{P}_{\mathbf{0}}(U, V)$, and (32) together with Lemma 6.7 yields

$$\text{VAL}(S, S) = \text{VAL}(S, S \setminus \{B\}). \tag{33}$$

This in turn implies $\lambda'_B \geq 0$, by Theorem 6.6, and via another invocation of the theorem, $(\tilde{v}', \tilde{\alpha}') = \text{VAL}(U, V)$ follows. In particular, $\lambda_B = \lambda'_B$ for $B \in S$ and $\lambda_B = 0$ for $B \in U \setminus S$.

To prove that S equals the set in (30) we need another ingredient: $\lambda'_B = \lambda_B > 0$ for $B \in S \setminus V$. With Theorem 6.6, this follows from $\lambda'_B \geq 0$ and (32). Consequently,

$$S = V \cup \{B \in U \setminus V \mid \lambda'_B > 0\},$$

as desired.

With the representation (30) for the sink S of $[V, U]$ it is easy to see that $S \in [V, U]$ is inclusion-minimal with $\text{VAL}(S, S) = \text{VAL}(U, V)$. A proper subset $S' \subset S$, $V \subseteq S'$, differs from S in at least one ball $B' \in S \setminus S'$, which must have $\lambda_{B'} \neq 0$ according to (30). It follows from Theorem 6.6 that $(\tilde{v}, \tilde{\alpha})$ disagrees with the value of (S', S'), so $\text{VAL}(S', S') \neq \text{VAL}(S, S) = \text{VAL}(U, V)$. □

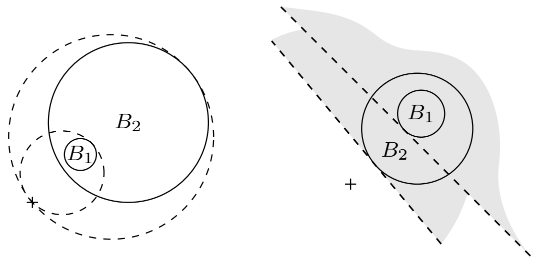

Fig. 20. The USO (c) from Theorem 6.9 for a set $T = \{B_1, B_2\}$ of two circles (a). A vertex $J \subseteq T$ of the cube corresponds to the solution $\text{VAL}(J, J)$ of program $\mathcal{P}_{\mathbf{0}}(J, J)$ and represents a halfspace $\text{FH}(J, J)$ in the dual (b) and the ball $\text{GMB}_{\mathbf{0}}(J, J)$ (gray) in the primal (a). Every edge $\{J, J \cup \{B\}\}$ of the cube is oriented towards $J \cup \{B\}$ if and only if the halfspace $\text{FH}(J, J)$ does not dominate the scaled version of B^* (which in this example is B^* itself). The global sink S in the resulting orientation corresponds to the inclusion-minimal subset S with $\text{VAL}(S, S) = \text{VAL}(T, \emptyset)$. For the definition of the USO, the halfspace $\text{FH}(T, T)$ is irrelevant, and since $\text{FH}(\emptyset, \emptyset)$ does not dominate any ball, all edges incident to $\emptyset$ are outgoing (therefore, the figure does not show these halfspaces).

Figure 20 illustrates Theorem 6.9 for a set T of two circles.

Specialized to the case of points, this result is already known;[10] however, our proof removes the general position assumption.

In order to apply USO-algorithms[24] to find the sink of our orientation $\mathcal{O}$, we have to evaluate the orientation of an edge $\{J, J \cup \{B\}\}$, i.e., we must check

$$\text{VAL}(J, J) \neq \text{VAL}(J \cup \{B\}, J). \tag{34}$$

If $J = \emptyset$, this condition is always satisfied. Otherwise, we first solve program $\mathcal{C}_{\mathbf{0}}(J, J)$, which is easy: by the Karush-Kuhn-Tucker conditions from Lemma 6.5(i), it suffices to solve the linear system consisting of the Eqs. (21), (22), and the feasibility constraints $v^T d_B + \alpha = \sigma_B$, $B \in J$. We know that this system is regular because the optimal solution is unique and uniquely determines the Karush-Kuhn-Tucker multipliers.

If the solution $(\tilde{v}, \tilde{\alpha})$ satisfies $\tilde{v}^T \tilde{v} \geq 1$, we have already found the value $(\tilde{v}, \tilde{\alpha})$ of (J, J) (Lemma 6.3(i)), and we simply check whether

$$\tilde{v}^T d_B + \tilde{\alpha} < \sigma_B, \tag{35}$$

a condition equivalent to (34). If $\tilde{v}^T \tilde{v} < 1$, we solve $\mathcal{C}'_{\mathbf{0}}(J, J)$: by the Karush-Kuhn-Tucker conditions from Lemma 6.6(ii), it suffices to solve the system consisting of the Eqs. (25), the feasibility constraints $v^T d_B + \alpha = \sigma_B$, $B \in J$, and $v^T v = 1$. For this, we first solve the subsystem of linear equations; like in the proof of Lemma 3.1 we can show that the solution space is one-dimensional. The additional constraint $v^T v = 1$ selects at most two candidate solutions[s], exactly one of which

[s]A continuous set of solutions would contradict the uniqueness of $(\tilde{v}, \tilde{\alpha})$ in Lemma 6.3(iii).

must satisfy (26). In this way we obtain the value of (J, J) and again evaluate (35).

Equation (34) gives an easy way to evaluate the orientation of the upward edge $\{J, J \cup \{B\}\}$, given the value of (J, J). We note that the orientation of the downward edge $\{J, J \setminus \{B\}\}$ can be read off the Karush-Kuhn-Tucker multiplier λ_B associated with $\textsc{val}(J, J)$: orient from $J \setminus \{B\}$ towards J if and only if $\lambda_B > 0$ (refer to the proof of Theorem 6.9). Moreover, the solution of the second program $\mathcal{C}'_{\mathbf{0}}(J, J)$ can be obtained by reusing intermediate results from the computations for the first program $\mathcal{C}_{\mathbf{0}}(J, J)$.[11]

With the currently best known USO algorithm we can find the sink of an m-dimensional USO with an expected number of $\mathcal{O}(c^m)$ *vertex evaluations*, where $c \approx 1.438$.[24] Since in our case a vertex evaluation (determine the orientations of all edges incident to some vertex) essentially requires to solve one system of linear equations, we obtain an expected running time of $\mathcal{O}(d^3 c^m)$ to solve problem $\text{SEBB}_{\mathbf{0}}$ for a set of $m \leq d$ signed balls. Plugging $f(d) = \mathcal{O}(d^3 c^d)$ into Theorem 4.4, we obtain the currently best known exact algorithm for SEBB.

It remains an open problem whether $\text{SEBB}_{\mathbf{0}}$ or SEBB can be solved in *sub-exponential* time as is the case for problem SEBP.[7]

7. Acknowledgment

We thank Menelaos Karavelas for pointing out the usefulness of the inversion transform in this context. We also thank two anonymous referees for providing additional references, and for valuable suggestions which helped to substantially improve the presentation. We are grateful to Isha Geigenfeind for his work on the practical aspects of the USO approach.

References

1. M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear programming: theory and applications*. Wiley, 1979.
2. M. Berger. *Geometry (vols. 1–2)*. Springer-Verlag, 1987.
3. B. Chazelle and J. Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *J. Algorithms*, 21:579–597, 1996.
4. V. Chvátal. *Linear programming*. W. H. Freeman, New York, NY, 1983.
5. M. E. Dyer. Linear time algorithms for two- and three-variable linear programs. *SIAM J. Comput.*, 13:31–45, 1984.
6. M. E. Dyer. A class of convex programs with applications to computational geometry. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 9–15, 1992.
7. B. Gärtner. A subexponential algorithm for abstract optimization problems. *SIAM J. Comput.*, 24(5):1018–1035, 1995.
8. B. Gärtner. Fast and robust smallest enclosing balls. In *Proc. 7th Annual European Symposium on Algorithms (ESA)*, volume 1643 of *Lecture Notes Comput. Sci.*, pages 325–338. Springer-Verlag, 1999.
9. B. Gärtner and E. Welzl. Linear programming—randomization and abstract frameworks. In *Proc. 13th Sympos. Theoret. Aspects Comput. Sci.*, volume 1046 of *Lecture Notes Comput. Sci.*, pages 669–687. Springer-Verlag, 1996.

10. B. Gärtner and E. Welzl. Explicit and implict enforcing—randomized optimization. In *Computational Discrete Mathematics*, volume 2122 of *LNCS*, pages 25–46. Springer-Verlag, 2001.
11. I. Geigenfeind. Implementation of Miniball-of-Balls using USO-algorithms. Semester thesis, ETH Zürich, 2004.
12. T. H. Hopp and C. P. Reeve. An algorithm for computing the minimum covering sphere in any dimension. Technical Report NISTIR 5831, National Institute of Standards and Technology, 1996.
13. V. Karamcheti, C. Li, I. Pechtchanski, and C. Yap. A core library for robust numeric and geometric computation. In *Proceedings of the fifteenth annual symposium on Computational geometry*, pages 351–359. ACM Press, 1999.
14. P. Kumar, J. S. B. Mitchell, and A. Yıldırım. Computing core-sets and approximate smallest enclosing hyperspheres in high dimensions. In *Algorithm Engineering and Experimentation (Proc. ALENEX '03)*, Lecture Notes Comput. Sci., pages 45–55. Springer-Verlag, 2003.
15. R. Kurniawati, J. S. Jin, and J. A. Shepherd. The SS^+-tree: an improved index structure for similarity searches in a high-dimensional feature space. In *Proc. 5th Storage and Retrieval for Image and Video Databases SPIE*, volume 3022, pages 110–120, 1997.
16. J. Matoušek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. *Algorithmica*, 16:498–516, 1996.
17. N. Megiddo. Linear-time algorithms for linear programming in $\mathbb{R}^3$ and related problems. *SIAM J. Comput.*, 12(4):759–776, 1983.
18. N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31:114–127, 1984.
19. N. Megiddo. On the ball spanned by balls. *Discrete Comput. Geom.*, 4:605–610, 1989.
20. K. Mehlhorn and S. Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, UK, 2000.
21. K. G. Murty. Note on Bard-type scheme for solving the complementarity problem. In *Operations Research*, volume 11, pages 123–130, 1974.
22. V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. In *Proc. 7th Annu. ACM Sympos. Comput. Geom.*, pages 357–363, 1991.
23. R. Seidel. Personal communication, 1997.
24. T. Szabó and E. Welzl. Unique sink orientations of cubes. In *Proc. 42nd annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 547–555, 2001.
25. E. Welzl. Smallest enclosing disks (balls and ellipsoids). In H. Maurer, editor, *New Results and New Trends in Computer Science*, volume 555 of *Lecture Notes Comput. Sci.*, pages 359–370. Springer-Verlag, 1991.
26. H.-M. Will. *Computation of additively weighted Voronoi cells for applications in molecular biology*. PhD thesis, ETH Zürich, 1998.
27. S. Xu, R. Freund, and J. Sun. Solution methodologies for the smallest enclosing circle problem. *Computational Optimization and Applications*, 25:283–292, 2003.
28. G. Zhou, K.-C. Toh, and J. Sun. Efficient algorithms for the smallest enclosing ball problem. http://www.bschool.nus.edu.sg/depart/ds/sunjiehomepage/, 2002.