

LLL reducing with the most significant bits

Saruchi

Indian Institute of Technology New Delhi
saruchigoel@gmail.com

Damien Stehlé

U. Lyon, CNRS, ENS de Lyon, INRIA, UCBL
Laboratoire LIP
damien.stehle@gmail.com

Ivan Morel

U. Lyon, CNRS, ENS de Lyon, INRIA, UCBL
U. Sydney, Laboratoire LIP
ivan.morel@ens-lyon.org

Gilles Villard

U. Lyon, CNRS, ENS de Lyon, INRIA, UCBL
Laboratoire LIP
gilles.villard@ens-lyon.fr

ABSTRACT

Let B be a basis of a Euclidean lattice, and \tilde{B} an approximation thereof. We give a sufficient condition on the closeness between \tilde{B} and B so that an LLL-reducing transformation U for \tilde{B} remains valid for B . Further, we analyse an efficient reduction algorithm when B is itself a small deformation of an LLL-reduced basis. Applications include speeding-up reduction by keeping only the most significant bits of B , reducing a basis that is only approximately known, and efficiently batching LLL reductions for closely related inputs.

Categories and Subject Descriptors

I.1.2 [Symbolic and Algebraic Manipulation]: Algorithms; F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems; G.1.3 [Numerical Linear Algebra]: Conditioning

General Terms

Algorithms

Keywords

Lattice basis reduction, LLL, floating point arithmetic

1. INTRODUCTION

We consider the problem of computing a reducing transformation for a Euclidean lattice basis B using an approximation \tilde{B} to $B \in \mathbb{R}^{m \times n}$. The necessary background on basis reduction is given in Section 2.2. Our notion of reduction is a variation, robust to perturbations, of the LLL reduction introduced by Lenstra, Lenstra and Lovász in [10]. We assume that the basis vectors are given by the columns of a full column rank matrix B . We establish a bound on the closeness between \tilde{B} and B so that if U is a reducing transformation for \tilde{B} , i.e., the matrix $\tilde{B}U$ is LLL-reduced, then BU also is reduced.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ISSAC'14, July 23-25, 2014, Kobe, Japan.
Copyright 2014 ACM 978-1-4503-2501-1/14/07 ...\$15.00.
<http://dx.doi.org/10.1145/2608628.2608645>.

A main application is to compute a reducing transformation using only a limited number of bits of the input basis, hence possibly at a lower cost than using the entire initial basis. An approach for LLL-reducing a basis B may then be: 1) use the bound and compute an appropriate rounding precision, and deduce an approximation \tilde{B} sufficiently close to B ; 2) compute U by reducing \tilde{B} ; 3) output BU . We follow this general approach for designing Algorithm 1.

We then develop a column scaling strategy for handling cases where the input basis vectors have unbalanced magnitudes. Indeed, in addition to mastering the bit-size of the inputs, dealing with homogeneous magnitudes is often better for lowering the computational cost: as far as we are aware of, no known LLL-reduction algorithm preserves the small bit-size of a “floating-point” basis during the execution, if the magnitudes of the columns differ. A scaling is a pre-processing of the basis B that provides with a “more balanced” matrix B' that is also appropriate for computing a reducing transformation. In this purpose we design Algorithm 2 that may be used as Step 2') in the general LLL-reducing scheme above. We note that after both the approximation and the scaling processes, “almost any” LLL-reduction algorithm could be used at Step 2) or 2') for computing the reducing transformation.

An important contribution in the study of approximation conditions for preserving reducing transformations has been made by Buchmann in [2]. In the rest of the paper the matrix norm induced by the Euclidean norm is denoted by $\|B\|$. All our vectors will be column vectors, and will be denoted in bold. Buchmann considers an approximation $\tilde{B} = B + \Delta B$ of absolute precision p to B such that $\|\Delta B\| \leq 2^{-p}$. He provides (see [2, Cor. 4] and its proof) a sufficient condition bound on p to guarantee that if U is such that $(B + \Delta B)U$ is LLL-reduced then BU also has small norm vectors, i.e., within a factor $2^{O(n)}$ of the successive minima of the lattice spanned by B (see Section 2.2 for the definition of the lattice minima). In his analysis, Buchmann relies on the orthogonality defect $\text{od}(B) = \frac{\prod_i \|\mathbf{b}_i\|}{\sqrt{\det(B^t B)}}$, and requires a precision p that is logarithmic in the dimension and the orthogonality defect $\text{od}(B)$, and involves the sizes of the successive minima. In terms of relative precision, i.e., for ΔB such that $\frac{\|\Delta B\|}{\|B\|} \leq 2^{-p}$, the bound is at least logarithmic in the orthogonality defect.

We work with a wider class of approximations: we consider columnwise perturbations of matrices such that $B +$

ΔB satisfies $\max_{i \leq n} \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq \varepsilon$, where $\varepsilon \geq 0$ is small. Hence $B + \Delta B$ is an approximation to B with small columnwise relative precision. In the case of approximate computations, the error bound ε is of the order of 2^{-p} , where p is a working precision. Our type of approximations models matrix truncation with small error relatively to each column magnitude. This choice is appropriate for taking into account the backward rounding errors of standard numerical QR-factorization algorithms such as Householder's, that are at the heart of fast reducing algorithms. It also preserves LLL-reducedness, as shown in [4]. To introduce our results, we need to define the QR-factorization. Let $B \in \mathbb{R}^{m \times n}$ be full column rank. There exists a unique pair (Q, R) such that

$$B = Q \cdot R, \quad Q \in \mathbb{R}^{m \times n}, \quad R \in \mathbb{R}^{n \times n},$$

the columns of Q are orthonormal and the matrix R is upper-triangular with positive diagonal coefficients. The matrices Q and R are respectively called the *Q-factor* and *R-factor* of B . For insights into perturbation analysis we may refer the reader to Higham [6], and to [3, 4] in the context of QR-factorization.

Our first result is an improved bound that characterizes which columnwise approximations $B + \Delta B$ to B are allowed such that a reducing transformation U for $B + \Delta B$ remains valid for B . Instead of using the defect and the lattice minima, we relate the sufficient precision $1/\varepsilon$ to the quantity $\text{cond}(R) = \| |R| \cdot |R^{-1}| \|$. This quantity is defined for any square invertible matrix, and may be viewed as a condition number for the problem of computing R . (See, e.g., Zha [19] and Higham [6, Ch. 19].)

In Theorem 1, we show that as soon as

$$\max_{i \leq n} \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq \frac{1}{cm^6 \beta^n \text{cond}^2(R)}$$

for a constant c that may be made explicit and a constant β that can be chosen arbitrarily close to 2, if U reduces $B + \Delta B$, then U also reduces B . Because of the approximation, the matrix BU is LLL-reduced for parameters slightly weaker than those for which $(B + \Delta B)U$ is reduced. However, one may ensure that the parameters are degraded by an arbitrarily small amount (by increasing c), and the main relevant properties of LLL-reduced bases are preserved, such as those reminded in Lemma 4.

The bound above indicates that taking $p = 2 \log \text{cond}(R) + n(1 + \varepsilon) + O(\log m)$ bits of precision on the input basis B suffices for an arbitrarily small ε . We will see, with Lemma 11, that $\text{cond}(R)$ is a more accurate measure than the orthogonality defect $\text{od}(B)$. Indeed, the condition we propose is never more restrictive (up to a $O(\log n)$ additive term) and may be much less so than Buchmann's. In particular, we exhibit a family of bases for which we divide the number of required bits by $\Omega(n)$.

A direct application is an interesting situation considered in [1]. It concerns matrices B whose R-factor satisfies $|r_{ij}| \leq \eta r_{ii}$ for $\eta \geq 0$ (see Definition 1, with $\theta = 0$), and for which the ratio $h = \frac{\max r_{ii}}{\min r_{ii}}$ is bounded. In Lemma 12, we show that in this case $\log \text{cond}(B) \lesssim n \log(1 + \eta) + \log h$. It follows that $\approx n(1 + 2 \log(1 + \eta)) + 2 \log h + O(\log m)$ bits of columnwise precision suffice for computing a reducing matrix. The bound is especially interesting when $h \ll \log \|B\|$, such as in [1, Se. 3.2]. The authors of [1] use a constant factor fewer bits than we do (they indeed work with about $\log h$ bits). However, their study is restricted to the first vector

of the output basis, which is shown to be no more than a factor $2^{O(n)}$ longer than that of an LLL-reduced basis, compared to the n th root of the lattice determinant (see Section 2.2 for the definition of the lattice determinant). This interesting result is an example of transfer between the precision and the quality of the reduction. It also demonstrates the accuracy of our all-purpose bound.

It is essential to note that the ratio h given by the diagonal of R may not be relevant in general for indicating allowed truncations. Indeed, a large h may not imply a large $\text{cond}(\cdot)$, as shown by LLL-reduced bases, which can have arbitrarily large ratios h but always satisfy $\text{cond}(R) = 2^{O(n)}$ (see Lemma 9).

Our second contribution is an LLL-reduction specifically designed for (floating-point) matrices of the form $B = M_B \cdot E_B$, where $M_B \in \mathbb{Z}^{m \times n}$ is full column rank and $E_B = \text{diag}(2^{e_i})$, with $e_i \in \mathbb{Z}$ for all i . Such matrices include the ones obtained by columnwise rounding. If the norms of the columns are unbalanced (i.e., the e_i 's have different orders of magnitude), the compactness of the representation may be lost when applying an LLL-reduction algorithm to B , as the columns get mixed. To circumvent this issue, we propose an algorithm (Algorithm 2) that applies a column scaling D^{-1} to $M_B E_B$ before calling an LLL-reduction algorithm on $B' = M_B E_B D^{-1}$. The obtained transformation U for B' is then mapped to a transformation $D^{-1}UD$ for B .

Concerning the correctness of the approach, the main difficulty is to find a scaling D such that $D^{-1}UD$ is unimodular and indeed reduces B . This is solved by numerically estimating the diagonal of the R-factor R of B and identifying blocks of consecutive vectors such that $r_{ii} \ll r_{jj}$ if j belongs to a block subsequent to the one of i . These blocks are the main source of unbalancedness between the norms of the columns of B , and the computed scaling annihilates it. Our algorithm then relies on a "well-behaved" LLL-reduction algorithm that does not destroy the block structure (most known LLL-reduction algorithms are well-behaved, as explained in Section 5.1).

To measure the efficiency of the approach, the relevant quantity is the bit-size of the entries of B' once converted to an integer matrix. To estimate it, we view $B = M_B E_B$ as a distortion $B = \Sigma C$ of an LLL-reduced matrix C , where $\Sigma \in \mathbb{R}^{m \times m}$ is non-singular. We then prove, in Theorem 3, that if $\sigma_1 \geq \dots \geq \sigma_m$ are the singular values of Σ , then the bit-size of B' is as $O\left(n + \log \|M_B\| + \log \prod_{i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}}\right)$. Several Σ 's may exist so that $\Sigma^{-1}B$ is LLL-reduced, and one may optimize the choice of Σ to lower $\prod_{i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}}$.

A direct application is to LLL-reduce for strong LLL parameters, a basis that is already reduced for some weak LLL parameters. This is a tempting approach in practice, as LLL-reducing for weak parameters is typically much faster. This strategy is mentioned, e.g., in [5, Se. 2.6.1] and [9, pp. 70–72].¹ Here Σ is the identity and $\prod_{i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}} = 1$. Another particular case is $\Sigma = \text{diag}(2^\ell, 1, \dots, 1)$, for which we have $\prod_{i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}} = 2^\ell$. This is used in the polynomial factoring algorithm of [7] as well as at the bottom of the recursion in the \tilde{L}^1 algorithm [15].

Finally, our algorithm may be used to batch reductions of

¹In [9], the idea is attributed to He, but we could not find the corresponding work. Preliminary results, for strengthening the reducedness, were presented as a poster [11].

closely related lattices $(L_k)_k$ described by bases $(B_k)_k$ such that $L_{k+1} = \Sigma_k L_k$ for all k , where the Σ_k 's have balanced singular values. One may then LLL-reduce B_1 , and, for $k \geq 1$, use the transformation matrix U_k computed for B_k , before calling our algorithm on $B_{k+1}U_k = \Sigma_k(B_kU_k)$. This strategy is used, e.g., in communications theory and crypt-analysis applications of LLL [13, 1]. Our algorithm could prove useful to accelerate and analyse these applications.

Remark. The present work generalizes several results previously investigated for the design of the \tilde{L}^1 algorithm [15]. Perturbation analyses and approximations indeed play a key role since \tilde{L}^1 heavily rely on well-chosen truncations. The results in [15] are essentially focused on LLL-reduced bases, and on specific deformations of such bases. The generalization here is a study with no restrictive assumptions on the initial basis B .

Notations. If B is a real-valued matrix, then $|B|$ (resp. $\lfloor B \rfloor$) is the matrix obtained by replacing each entry of B by its absolute value (resp. the largest integer no greater than it). If B and B' are two matrices of identical dimensions, the relation $B \leq B'$ must be understood as a componentwise bound. The notations $\|B\|_F$ and $\|B\|_1$ respectively refer to the Frobenius and induced Manhattan norms of B . If B is square and non-singular we define $\kappa(B) = \|B\| \cdot \|B^{-1}\|$. Clearly, we have $\text{cond}(B) \leq \kappa(B)$. If $(x_i)_i$ is a sequence of cardinality n , we let $\text{diag}_i(x_i)$ denote the $n \times n$ diagonal matrix whose diagonal entries are the x_i 's. The computational costs are given in terms of bit operations. We let $\mathcal{M}(t)$ the cost of multiplying two t -bit long integers. Finally, all logarithms are in base 2.

2. QR AND BASIS REDUCTION

We extensively rely on roundings and perturbations. We say that \tilde{B} is an approximation to B of *columnwise relative precision* $p \geq 0$ if $B - \tilde{B} = \Delta B$ satisfies $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 2^{-p}$.

Approximating real matrices by floating-point ones fits into this context. A precision- p floating-point number is of the shape $m2^e$ with m, e integers with $|m| \leq 2^p - 1$. For any $x \in \mathbb{R}$, there exists m_x, e_x integers with $|m_x| \leq 2^p - 1$ such that $|m_x 2^{e_x} - x| \leq 2^{-p}|x|$. We call $m_x 2^{e_x}$ a precision p approximation to x . If x is a non-zero integer with known bit-length, such an approximation may be computed in time $O(p + \log(1 + \log|x|))$, and e_x has bit-length $O(\log \log|x|)$.

An interesting matrix perturbation is *columnwise rounding*. Let $B = (\mathbf{b}_i)_i \in \mathbb{R}^{m \times n}$ be full column rank, and p be a non-negative integer. For each $i \leq n$, let e_i be an integer such that $\frac{2^{e_i} - \|\mathbf{b}_i\|}{\|\mathbf{b}_i\|} < 3/4$. Let $E_B = 2^{-p} \text{diag}_i(2^{e_i})$ and $M_B = \lfloor B \cdot E_B^{-1} \rfloor$. Then $\Delta B = B - M_B E_B$ satisfies $\max_{i \leq n} \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq \frac{\sqrt{m}}{2^{p-2}}$. We may therefore view $M_B E_B = M_B \text{diag}_i(2^{e_i-p})$ as a columnwise floating-point approximation to B . Each entry of the mantissa matrix M_B is an integer of magnitude smaller than 2^{p+1} , and the matrix E_B that may be represented on $O(n \log \log \|B\| + \log p)$ bits, gives column exponents.

2.1 Numerical aspects of QR-factorization

The numerical aspects of QR-factorization have been extensively studied, and we refer the reader to [6, Ch. 19] for a comprehensive entry point to the topic. The following is an explicit variant of classical results.

LEMMA 1 ([4, SE. 6]). *Let $p \geq 0$ and $B \in \mathbb{R}^{m \times n}$ be non-singular with R-factor R . Let \hat{R} be the R-factor computed by Householder's algorithm with floating-point precision p . If $80mn2^{-p} < 1$, then there exists an orthogonal \hat{Q} such that $\hat{Q}\hat{R} = B + \Delta B$ with $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 80mn2^{-p}$.*

Given a matrix B , the number of bit operations consumed by Householder's algorithm for computing an approximation to the R-factor of B is $O(mn^2(\mathcal{M}(p) + \log \log \|B\|))$.

The backward stability lemma above is often combined with a sensitivity result, such as the one below, in order to obtain forward error bounds on the computed quantities.

LEMMA 2 (ADAPTED FROM [4, TH. 2.3]). *Let B be full column rank in $\mathbb{R}^{m \times n}$, and let R denote its R-factor. Let $\Delta B \in \mathbb{R}^{m \times n}$. If $\max_{i \leq n} \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} < 1/(12m\sqrt{n}\text{cond}(R))$, then $B + \Delta B$ is full column rank and its R-factor $R + \Delta R$ satisfies $\|\Delta R \cdot R^{-1}\|_F \leq 6m\sqrt{n}\text{cond}(R) \max_{i \leq n} \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|}$.*

PROOF. The assertion on the rank follows from [4, Le. 2.2]. From the end of the proof of [4, Th. 2.3] with $D = I$, we have $\|\Delta R \cdot R^{-1}\|_F \leq (\sqrt{6} + \sqrt{3})\sqrt{2m}\sqrt{n}\text{cond}(R) \max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|}$. \square

We will also use the following result on the effect on the R-factor of a matrix B of applying a distortion to B .

LEMMA 3. *Let $B \in \mathbb{R}^{m \times n}$ be full column rank and let R denote its R-factor. Let $\Sigma \in \mathbb{R}^{m \times m}$ be non-singular, and let R' denote the R-factor of ΣB . Then, for all i , we have $\|\Sigma^{-1}\|^{-1} \leq \frac{r'_{ii}}{r_{ii}} \leq \|\Sigma\|$.*

PROOF. The proof is adapted from the proof of [7, Le. 4]. Let $V_i(B) = \{\mathbf{b}_i - \sum_{j < i} y_j \mathbf{b}_j : y_1, \dots, y_{i-1} \in \mathbb{R}\}$. Then r_{ii} is the norm of the shortest vector \mathbf{b} in $V_i(B)$. Now, the vector $\Sigma \mathbf{b}$ belongs to $V_i(\Sigma B)$. As a result, we have $r'_{ii} \leq \|\Sigma \mathbf{b}\| \leq \|\Sigma\| r_{ii}$. The proof that $r'_{ii} \geq \frac{r_{ii}}{\|\Sigma^{-1}\|}$ is analogous. \square

2.2 Lattices and LLL basis reduction

A lattice L is the set of integer combinations of linearly independent vectors in a euclidean space \mathbb{R}^n : any lattice may be written as $L = L(B) = B\mathbb{Z}^n$, for some full column rank matrix $B \in \mathbb{R}^{m \times n}$. The columns of B are said to form a basis of L . If the lattice dimension satisfies $n \geq 2$, the lattice admits infinitely many lattice bases, related by unimodular matrices (i.e., square integer matrices of determinant ± 1): for two full column rank matrices $B, C \in \mathbb{R}^{m \times n}$, we have $B\mathbb{Z}^n = C\mathbb{Z}^n$ if and only if there exists a unimodular matrix U such that $C = B \cdot U$.

The sparsity of a lattice L may be quantified by its successive minima, defined as $\lambda_i(L) = \inf\{r : \dim(\text{span}L \cap \mathcal{B}(\mathbf{0}, r)) \geq i\}$, for all $i \leq n$. It may also be quantified with the lattice determinant $\det L = \prod_i r_{ii}$, where R is the R-factor of any basis of L .

In 1982, Lenstra *et al.* [10] introduced the notion of LLL-reduction of a lattice basis and the LLL-algorithm. If a basis is LLL-reduced, then it is short with respect to the minima of the spanned lattice, and, further, such a basis can be efficiently found using the LLL algorithm. Here we use a variation of LLL-reduction that is more suited to numerical computations.

DEFINITION 1 ([4, DEF. 5.3]). *Let $\Xi = (\delta, \eta, \theta)$ with $\eta \in (1/2, 1)$, $\theta \in (0, 1]$ and $\delta \in (\eta^2, 1)$. Let $B \in \mathbb{R}^{m \times n}$ be non-singular with QR factorization $B = Q \cdot R$. The matrix B is Ξ -LLL-reduced if:*

- for all $i < j$, we have $|r_{i,j}| \leq \eta r_{i,i} + \theta r_{j,j}$ (B is said size-reduced);
- for all i , we have $\delta \cdot r_{i,i}^2 \leq r_{i,i+1}^2 + r_{i+1,i+1}^2$ (B is said to satisfy Lovász' conditions).

Let $\Xi = (\delta, \eta, \theta)$ and $\Xi_w = (\delta_w, \eta_w, \theta_w)$ be valid LLL-parameters. We say that Ξ_w is weaker than Ξ and write $\Xi > \Xi_w$ if $\delta > \delta_w$, $\eta < \eta_w$ and $\theta < \theta_w$. If a basis is Ξ -LLL-reduced and if $\Xi > \Xi_w$, then it is also Ξ_w -LLL-reduced. This LLL-reduction variant is as powerful as the classical definition.

LEMMA 4 ([4, TH. 5.4]). *Let $B \in \mathbb{R}^{m \times n}$ be (δ, η, θ) -LLL-reduced with R-factor R , for valid parameters (δ, η, θ) . Let $\alpha = (\eta\theta + \sqrt{(1+\theta^2)\delta - \eta^2})/(\delta - \eta^2)$. Then, for all i , $r_{i,i} \leq \alpha \cdot r_{i+1,i+1}$ and $r_{i,i} \leq \|\mathbf{b}_i\| \leq \alpha^i \cdot r_{i,i}$. This implies that $\|\mathbf{b}_1\| \leq \alpha^{\frac{n-1}{2}} |\det B|^{1/n}$ and $\alpha^{i-n} r_{i,i} \leq \lambda_i(L(B)) \leq \alpha^i r_{i,i}$.*

The use of Ξ -LLL-reduction rather than the classical definition of LLL-reduction (corresponding to taking $\theta = 0$) is motivated by the following result. It says that a sufficiently precise approximation to a Ξ -LLL-reduced is Ξ_w -LLL-reduced for some $\Xi_w < \Xi$. This result is incorrect if one imposes $\theta_w = 0$.

LEMMA 5 (ADAPTED FROM [4, COR. 5.7]). *For any valid sets of parameters $\Xi = (\delta, \eta, \theta)$ and $\Xi_w = (\delta_w, \eta_w, \theta_w)$ with $\Xi_w < \Xi$, there exists a constant $c > 0$ (that may be made explicit) such that the following holds. For any Ξ -LLL-reduced $B \in \mathbb{R}^{m \times n}$ and any ΔB satisfying $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 1/(cm^2(1+\eta+\theta)^n \alpha^n)$ where α is as in Lemma 4, the basis $B + \Delta B$ is Ξ_w -LLL-reduced.*

The L^3 algorithm from [10] allows one to compute an LLL-reduced basis of the lattice spanned by a given $B \in \mathbb{Z}^{m \times n}$ in time $O\left(mn^4 \log^2 \|B\| \frac{\mathcal{M}(n+\log \|B\|)}{n+\log \|B\|}\right)$ (see [8]). The L^2 and H-LLL algorithms from [14, 12] achieve it within $O(m\mathcal{M}(n)n^3(n+\log \|B\|)\log \|B\|)$ bit operations, while the \tilde{L}^1 from [15] runs in time $\tilde{O}(mn^4(n+\log \|B\|))$.

Finally, we will use the following generalization of [15, Le. 5] to arbitrary bases, which provides a bound on the size of the unimodular matrix between any basis of a lattice and an LLL-reduced basis of the same lattice.

LEMMA 6. *Let $B \in \mathbb{R}^{m \times n}$ be full column rank. Let Ξ be a valid LLL-parameter, α as in Lemma 4, and U such that $C = BU$ is Ξ -reduced. We have:*

$$\forall i, j : |u_{ij}| \leq m^3 \alpha^n \text{cond}(R) \cdot \frac{r'_{jj}}{r_{ii}},$$

where R and R' respectively denote the R-factors of B and C .

PROOF. Let $B = QR$, $C = Q'R'$ be the QR-factorizations of B and C , respectively. Then

$$U = R^{-1}Q^t Q' R' = \text{diag}_i(r_{ii}^{-1}) \overline{R}^{-1} Q^t Q' \overline{R}' \text{diag}_i(r'_{ii}),$$

with $\overline{R} = R \cdot \text{diag}(1/r_{ii})$ and $\overline{R}' = R' \cdot \text{diag}(1/r'_{ii})$. We have $|\overline{R}^{-1}| \leq |R| |R^{-1}| \leq \text{cond}(R) \cdot T$, where $t_{ij} = 1$ if $i \leq j$ and $t_{ij} = 0$ otherwise. By Lemma 4, we have $|\overline{R}'| \leq \alpha^n T$. We also have $|Q|, |Q'| \leq M$, where $m_{ij} = 1$ for all i, j . Using the triangular inequality, we obtain:

$$\begin{aligned} |U| &\leq \text{cond}(R) \alpha^n \cdot \text{diag}(r_{ii}^{-1}) T M^t M T \text{diag}(r'_{ii}) \\ &\leq mn^2 \alpha^n \text{cond}(R) \cdot \text{diag}(r_{ii}^{-1}) N \text{diag}(r'_{ii}), \end{aligned}$$

where N is the all-1 matrix with appropriate dimensions. \square

3. WELL-CONDITIONED MATRICES

As we have seen, the quantity $\text{cond}(\cdot)$ plays a role both for the sensitivity of the R-factor under columnwise perturbations (Lemma 2) and for the size of the unimodular transformation between a lattice basis and an LLL-reduced basis of the same lattice (Lemma 6). It is therefore interesting to investigate sufficient conditions to ensure a small value of $\text{cond}(\cdot)$.

LEMMA 7. *Let $B \in \mathbb{R}^{m \times n}$ full column rank and $\Sigma \in \mathbb{R}^{m \times m}$ non-singular. Let R and R' respectively denote the R-factors of B and ΣB . Then $\text{cond}(R') \leq mn \kappa(\Sigma) \text{cond}(R)$. Further, if B and B' are square and upper-triangular, then $\text{cond}(R') \leq \kappa(\Sigma) \text{cond}(R)$.*

PROOF. Write $B = QR$ (resp. $\Sigma B = Q'R'$), where the columns of Q (resp. Q') are orthogonal. The first result follows from the equation $|R'| |(R')^{-1}| = |((Q')^T \Sigma Q) R| \cdot |R^{-1} ((Q')^T \Sigma Q)^{-1}|$ and the triangular inequality. The proof simplifies when Q and Q' are square, as in that case we have $\|(Q')^T \Sigma Q\| = \|\Sigma\|$ and $\|((Q')^T \Sigma Q)^{-1}\| = \|\Sigma^{-1}\|$. \square

As a corollary, we obtain the fact that a small columnwise perturbation cannot increase $\text{cond}(\cdot)$ by much.

LEMMA 8. *Let $B \in \mathbb{R}^{m \times n}$ be full column rank with R-factor R , and $\Delta B \in \mathbb{R}^{m \times n}$. If we have $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} < 1/(12m\sqrt{n}\text{cond}(R))$, then $B + \Delta B$ is full column rank and its R-factor $R + \Delta R$ satisfies $\text{cond}(R + \Delta R) \leq 4 \text{cond}(R)$.*

PROOF. We write $R + \Delta R = \Sigma R$, with $\Sigma = I + \Delta R \cdot R^{-1}$. By Lemma 2, we know that $\|\Delta R \cdot R^{-1}\| \leq \|\Delta R \cdot R^{-1}\|_F \leq 6m\sqrt{n}\text{cond}(R) \max_i \|\Delta \mathbf{b}_i\|/\|\mathbf{b}_i\|$. Thanks to the assumption on $\max_i \|\Delta \mathbf{b}_i\|/\|\mathbf{b}_i\|$, we obtain that Σ is non-singular. Further $\|\Sigma\| \leq 2$ and $\|\Sigma^{-1}\| \leq \|I + \sum_{k \geq 1} (\Delta R \cdot R^{-1})^k\| \leq 2$. As a result, we obtain $\kappa(\Sigma) \leq 4$. Lemma 7 provides the result. \square

The following result shows that any LLL-reduced basis has a small $\text{cond}(\cdot)$.

LEMMA 9 ([4, LE. 5.5]). *Let $\Xi = (\delta, \eta, \theta)$ be any valid set of LLL-parameters. If $B \in \mathbb{R}^{m \times n}$ is Ξ -LLL-reduced and R is its R-factor, then $\text{cond}(R) \leq \frac{1-\eta-\theta}{1+\eta+\theta} \alpha^{n-1} ((1+\eta+\theta)\alpha)^n$, with α as in Lemma 4.*

In fact, LLL-reducedness is a much stronger assumption than needed, for $\text{cond}(R)$ to be bounded as $2^{O(n)}$. A weaker assumption is used in the following result. Note that the assumption is satisfied for LLL-reduced bases, by Lemma 4 applied to square diagonal sub-blocks or the R-factor.

LEMMA 10. *Let $B \in \mathbb{R}^{m \times n}$ be full column rank with R-factor R . Assume that there exists $\alpha > 1$ such that for all $i \leq j$, we have $|r_{ij}| \leq \alpha^{j-i+1} r_{jj}$. Then $\text{cond}(R) \leq \frac{\alpha^2}{(\alpha^2-1)\sqrt{4\alpha^4-1}} (2\alpha^3)^n$.*

PROOF. Let $\overline{R} = R \cdot \text{diag}(r_{ii}^{-1})$. We have $\text{cond}(R) = \text{cond}(\overline{R}) \leq \kappa(\overline{R})$. A direct computation shows that $\|\overline{R}\| \leq (\sum_{i=1}^n \sum_{j=1}^i \alpha^{2j})^{1/2} \leq (\sum_{i=1}^n \alpha^{2i} \cdot \frac{\alpha^2}{\alpha^2-1})^{1/2} \leq \frac{\alpha^{n+2}}{\alpha^2-1}$. It now suffices to bound $\|\overline{R}^{-1}\|$ from above.

Write $\overline{R} = I + M$, where M is the matrix having same elements as \overline{R} but with zeroed diagonal coefficients. We have

$\bar{R}^{-1} = (I + M)^{-1} = \sum_{0 \leq k < n} (-M)^k$. Using the triangle inequality, we obtain that $|\bar{R}^{-1}| \leq \sum_{0 \leq k < n} |M|^k$. Let J denote the $n \times n$ matrix such that $J_{ij} = 1$ if $i - j = 1$, and $J_{ij} = 0$ otherwise. By assumption, we have $|M| \leq \alpha \sum_{1 \leq k < n} (\alpha J)^k = \alpha^2 J(I - \alpha J)^{-1}$. As a consequence:

$$\begin{aligned} |\bar{R}^{-1}| &\leq \sum_{0 \leq k < n} (\alpha^2 J(I - \alpha J)^{-1})^k \\ &= (I - \alpha^2 J(I - \alpha J)^{-1})^{-1} \\ &= (I - \alpha J)(I - (\alpha + \alpha^2)J)^{-1} \\ &\leq \sum_{0 \leq k < n} (\alpha + \alpha^2)^k J^k \leq \sum_{0 \leq k < n} (2\alpha^2)^k J^k. \end{aligned}$$

We derive that $\|\bar{R}^{-1}\| \leq (\sum_{k=0}^{n-1} (2\alpha^2)^{2k})^{1/2} \leq \frac{(2\alpha^2)^n}{\sqrt{4\alpha^4 - 1}}$, which leads to the result. \square

As discussed in the introduction, Buchmann provides in [2] a sufficient bound on the input precision to guarantee the correctness of the algorithm of the next section. The bound is at least logarithmic in the *orthogonality defect* $\text{od}(B) = \prod_i \frac{\|\mathbf{b}_i\|}{r_{ii}}$ of the full column rank matrix $B \in \mathbb{R}^{m \times n}$ with R-factor R . Our sufficient condition involves a precision logarithmic in $\text{cond}(B)$. The following lemma reveals the relationship between $\text{cond}(B)$ and $\text{od}(B)$.

LEMMA 11. *Let $B \in \mathbb{R}^{m \times n}$ be full column rank and R be its R-factor. Then for all i , we have $\|\mathbf{b}_i\| \leq r_{ii} \text{cond}(R)$, implying that $\text{od}(B) \leq \text{cond}(R)^n$. Oppositely, we have $\text{cond}(R) \cdot n^{-3/2} \leq \text{od}(B)$. Finally there exists a sequence of full column rank matrices $B \in \mathbb{R}^{n \times n}$ of growing dimension n such that $\text{od}(B) = \text{cond}(R)^{\Theta(n)}$.*

PROOF. For each $j \geq i$, the coefficient (j, i) of $|R| \cdot |R^{-1}|$ is bounded from below by $|r_{ji}|/r_{ii}$ (when doing the inner product of the j th row of $|R|$ with the i th column of $|R^{-1}|$), the coefficient $|r_{ji}|$ is multiplied with $1/r_{ii}$. This implies that $\frac{\|\mathbf{r}_i\|}{r_{ii}} \leq \text{cond}(R)$ holds for all i . Multiplying over varying i gives the first statement.

We now prove the second statement. The coefficient (i, j) of $|R^{-1}|$ is bounded from above by $\frac{1}{r_{ii}} \cdot \prod_{i < k \leq j} \frac{\|\mathbf{b}_k\|}{r_{kk}}$ (this can be obtained by using the cofactors of R to compute $|R^{-1}|$ and then applying Hadamard's bound). As a result we obtain that the coefficient (i, j) of $|R| \cdot |R^{-1}|$ is bounded from above by $\sum_{i \leq k \leq j} \prod_{k \leq \ell \leq j} \frac{\|\mathbf{b}_\ell\|}{r_{\ell\ell}}$, which is itself $\leq n \cdot \text{od}(B)$.

Finally, consider the n -dimensional upper triangular matrix R defined by $r_{ij} = \alpha^{j-i+1}$, for $\alpha > 1$ arbitrary. By considering only the first row of R , we obtain that $\text{od}(R) \geq \alpha^{n(n-1)/2}$. Lemma 10 allows us to complete the proof. \square

Another class of bases with relatively small $\text{cond}(\cdot)$ is given by upper triangular matrices B whose diagonal entries have balanced magnitudes, and which are size-reduced with $\theta = 0$ in Definition 1. If the largest ratio h between two diagonal entries is small, then as shown by next lemma, the quantity $\text{cond}(B) \leq h2^{O(n)}$ may be thought as small. (A geometric interpretation is given in [4, Se. 3.3].)

LEMMA 12. *Let $B \in \mathbb{R}^{n \times n}$ be an upper triangular, invertible matrix with the property that for all $i < j$ we have, $|b_{ij}| \leq \eta |b_{ii}|$, for some $\eta \geq 0$. Then we have $\text{cond}(B) \leq 2n(1 + \eta)^{n-1} \max_{i,j} \frac{|b_{ii}|}{|b_{jj}|}$.*

PROOF. Let $\bar{B} = \text{diag}(b_{ii}^{-1}) \cdot B$. For all $i < j$ we have $|\bar{b}_{ij}| \leq \eta \bar{b}_{ii} = \eta$. Therefore, we have $|\bar{B}^{-1}| \leq T^{-1}$, where $T \in \mathbb{R}^{n \times n}$ is upper triangular with $t_{ii} = 1$ and $t_{ij} = -\eta$ for $i < j$ (see, e.g., [6, Th. 8.12]). Since $S = T^{-1}$ satisfies $s_{ii} = 1$ and $s_{ij} = \eta(1 + \eta)^{j-i-1}$ for $i < j$ (see, e.g., [6, Eq. (8.4)]), we obtain $v_{ij} = 2\eta(1 + \eta)^{j-i-1}$, for $i < j$, where $V = |T|T^{-1}$. It follows that $|V| \leq 2(1 + \eta)^{n-1}M$ where $m_{ij} = 1$ for all $i \leq j$, and $m_{ij} = 0$ otherwise. Since $|\bar{B}||\bar{B}^{-1}| \leq V$, we may now write $\text{cond}(B) = \|\text{diag}(b_{ii}) \cdot V \cdot \text{diag}(b_{jj}^{-1})\| \leq 2(1 + \eta)^{n-1} \left\| \left(\frac{b_{ii}}{b_{jj}} \right)_{1 \leq i, j \leq n} \right\|$, which shows the assertion. \square

Given an invertible $B \in \mathbb{R}^{n \times n}$ with R-factor R , one may estimate $\text{cond}(R)$ in the following way. By [17], we have:

$$\begin{aligned} \text{cond}(R) &\leq n \text{cond}(B) \leq n^{3/2} \|B\| \cdot |B^{-1}|_1 \\ &\leq n^{3/2} \|BD^{-1}\|_1 \|DB^{-1}\|_1 \\ &\leq n^{5/2} \|BD^{-1}\| \|DB^{-1}\|, \end{aligned}$$

where $D = \text{diag}_i(\|\mathbf{b}_i\|_1)$. Therefore, it suffices to find estimates of $\|BD^{-1}\|$ and $\|DB^{-1}\|$. We refer the reader to [6, Ch. 15] for a presentation of classical approaches for estimating a matrix norm $\|A\|$, such as through a random sampling of vectors \mathbf{x}_i for measuring $\max_i \frac{\|A\mathbf{x}_i\|}{\|\mathbf{x}_i\|}$. If B is an integer matrix, this results in an algorithm of bit-complexity $\tilde{O}(n^3 \log \|B\|)$ using [18].

4. REDUCING BY ROUNDING

Our first main result is Theorem 1. We analyse the effect of applying to a lattice basis B a transformation matrix U reducing a perturbation $B + \Delta B$ of B . We rely on Lemma 13, which shows that a reducing transformation U for a given basis B remains a reducing transformation for any basis sufficiently close to B . This result, with a backward stability flavor, is then applied to $B + \Delta B$ for establishing the reducedness of B .

LEMMA 13. *For any valid sets of LLL-parameters $\Xi = (\delta, \eta, \theta)$ and $\Xi_w = (\delta_w, \eta_w, \theta_w)$ with $\Xi_w < \Xi$, there exists a constant $c > 0$ (that may be made explicit) such that the following holds. Let $B \in \mathbb{R}^{m \times n}$ full column rank, R its R-factor, and U such that BU is Ξ -LLL-reduced. Assume that $\Delta B \in \mathbb{R}^{m \times n}$ satisfies $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 1/(cm^6 \beta^n \text{cond}^2(R))$ with $\beta = (1 + \eta + \theta)\alpha^2$ and α as in Lemma 4. Then $(B + \Delta B)U$ is Ξ_w -reduced.*

PROOF. By Lemma 6, we have $|u_{ji}| \leq m^3 \alpha^n \text{cond}(R) \cdot \frac{r'_{ii}}{r_{jj}}$ for all i, j , where R (resp. R') is the R-factor of B (resp. $C = BU$). Let $C + \Delta C = (B + \Delta B)U$. We obtain that $\Delta \mathbf{c}_i = \sum_j u_{ji} \Delta \mathbf{b}_j$ satisfies (for all i):

$$\|\Delta \mathbf{c}_i\| \leq \left(m^3 \alpha^n \text{cond}(R) \max_j \frac{\|\Delta \mathbf{b}_j\|}{\|\mathbf{b}_j\|} \right) \cdot \sum_j \frac{r'_{ii}}{r_{jj}} \|\mathbf{b}_j\|.$$

Now, by Lemma 11, we have that $\|\mathbf{b}_j\|/r_{jj} \leq \text{cond}(R)$ holds for all j . By using the fact that $r'_{ii} \leq \|\mathbf{c}_i\|$, we derive that:

$$\max_i \frac{\|\Delta \mathbf{c}_i\|}{\|\mathbf{c}_i\|} \leq m^4 \alpha^n \text{cond}^2(R) \max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|}.$$

Applying Lemma 5 to C and $C + \Delta C$ provides the result. \square

The following result extends [15, Le 7], to any full column rank matrix B .

THEOREM 1. For any valid sets of parameters $\Xi_w < \Xi$, there exists $c > 0$ (that may be made explicit) such that the following holds. Let $B \in \mathbb{R}^{m \times n}$ full column rank, R its R-factor, and ΔB satisfying $\max_i \frac{\|\Delta \mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 1/(cm^6 \beta^n \text{cond}^2(R))$ with β as in Lemma 13. Then if U is such that $(B + \Delta B)U$ is Ξ -LLL-reduced, then BU is Ξ_w -LLL-reduced.

PROOF. By Lemma 8, we have $\text{cond}(R + \Delta R) \leq 4\text{cond}(R)$, where $R + \Delta R$ is the R-factor of $B + \Delta B$. We conclude by using Lemma 13 on $B + \Delta B$ with perturbation $-\Delta B$, to establish the reducedness of $B = (B + \Delta B) - \Delta B$. \square

As a corollary of the theorem just above, Algorithm 1 is correct. Note that an upper bound of $\text{cond}(R)$ is required as part of the input, where R denotes the R-factor of B . Such a bound may be derived from a priori information on B (e.g., using Lemmata 7, 8 and 9), or may be estimated, as explained at the end of Section 3. At Step 7, any LLL-reducing algorithm may be used. In the next section, we describe and analyze an LLL-reducing algorithm specifically designed for floating-point lattice bases $M_B E_B$, when they are themselves small distortions of LLL-reduced bases.

Input: $B \in \mathbb{R}^{m \times n}$ full column rank;
valid LLL-parameters Ξ_w ;
 $\chi \geq \text{cond}(R)$, where R is the R-factor of B .

Output: A Ξ_w -reduced basis of the lattice spanned by B .

- 1 Choose valid LLL-parameters $\Xi > \Xi_w$.
- 2 Compute the constants c and β of Theorem 1.
- 3 Set $p := \lceil \log(4cm^7 \beta^n \chi^2) \rceil$.
- 4 For each $i \leq n$, find $e_i \in \mathbb{Z}$ such that $\frac{|2^{e_i} - \|\mathbf{b}_i\|}{\|\mathbf{b}_i\|} \leq 3/4$.
- 5 Set $E_B := 2^{-p} \text{diag}_i(2^{e_i})$.
- 6 Set $M_B := \lfloor B \cdot E_B^{-1} \rfloor$.
- 7 Compute U such that $(M_B E_B) \cdot U$ is Ξ -LLL-reduced.
- 8 Return $B \cdot U$.

Algorithm 1: LLL-reduction of B using a columnwise floating-point approximation of B .

5. REDUCING BY SCALING

We now describe and analyze an algorithm for efficiently LLL-reducing floating-point lattice bases $M_B E_B$, such as the one involved at Step 7 of Algorithm 1. To LLL-reduce the floating-point matrix $M_B E_B$, we may interpret it as an integer matrix, and LLL-reduce that integer matrix. However, if the exponents are very unbalanced, the bit-size of $M_B E_B$ as an integer matrix (and hence the cost of the LLL-reduction) may be much higher than the bit-size of $M_B E_B$ as a floating-point matrix. Our algorithm scales the columns of $M_B E_B$, to obtain a matrix $M_B E_B D^{-1}$, so that the conversion to an integer matrix essentially preserves the small bit-size of the floating-point representation. The main difficulty to establish the correctness of the algorithm is to ensure that the transformation matrix U when LLL-reducing $M_B E_B D^{-1}$ is relevant for LLL-reducing $M_B E_B$ (note that the spanned lattices are different).

Input: $M_B \in \mathbb{Z}^{m \times n}$ full column rank;
 $E_B = 2^{-p} \text{diag}_i(e_i)$ with $e_i \in \mathbb{Z}$ for all i ;
valid LLL-parameters $\Xi = (\delta, \eta, \theta)$;
 $\chi \geq \text{cond}(R)$, where R is the R-factor of $M_B E_B$.

Output: A matrix pair (U, D) such that $D^{-1}UD$ is unimodular, $(M_B E_B)(D^{-1}UD)$ is Ξ -LLL-reduced and $D = \text{diag}(2^{d_i})$ with $d_i \in \mathbb{Z}$ for all i .

- 1 Set $p := 10 + \lceil \log(m^{3.5} \chi) \rceil$.
- 2 Call Householder's algorithm on $M_B E_B$ with precision p ; let \hat{R} be its output.
- 3 Set $i_0 := 1$ and $k := 1$.
- 4 For $i \leq n$, do: If $(\min_{j \geq i} \hat{r}_{jj}) > (8/\theta) \cdot \max_{j < i} \hat{r}_{jj}$, then increment k and set $i_k := i$.
- 5 For all $1 \leq \ell < k$,
set $g_\ell := (\min_{i_\ell \leq i < i_{\ell+1}} \hat{r}_{ii}) / (\max_{i_{\ell-1} \leq i < i_\ell} \hat{r}_{ii})$.
- 6 For all $1 \leq \ell < k$ and all $i_\ell \leq i < i_{\ell+1}$,
set $d_i := e_1 + \sum_{\ell' < \ell} \lceil \log(g_{\ell'}/4) \rceil$.
- 7 Set $D := \text{diag}(2^{d_i})$. /* Column scaling */
- 8 Set $\Xi' = (\delta, \eta, \theta/2)$.
- 9 Compute U s.t. $(M_B E_B D^{-1}) \cdot U$ is Ξ' -LLL-reduced.
- 10 Return (U, D) .

Algorithm 2: LLL-reduction of a floating-point matrix $M_B E_B$ using column scaling.

Algorithm 2 can be divided into four main parts:

- Finding approximations of the diagonal coefficients of the R-factor of the input basis $M_B E_B$ (Steps 1-2) for determining the scaling.
- Finding blocks, delimited by the i_ℓ 's, of consecutive vectors in $M_B E_B$, such that typical LLL-reduction algorithms do not swap vectors between these blocks, because the r_{ii} 's increase (Steps 3-4). Appropriate gaps between blocks allow to preserve the block structure after the scaling, which is a key ingredient for ensuring that U is block upper-triangular, and $D^{-1}UD$ is unimodular.
- Scaling the columns of $M_B E_B$, to shrink the eventual magnitude gaps between the r_{ii} 's of different blocks (Steps 5-7).
- LLL-reducing the scaled matrix (Steps 8-10).

5.1 Correctness

The following lemma ensures that the \hat{r}_{ii} 's are good approximations of the r_{ii} 's.

LEMMA 14. The matrix \hat{R} computed at Step 2 of Algorithm 2 satisfies $\max_i |\hat{r}_{ii} - r_{ii}|/r_{ii} \leq 1/2$.

PROOF. Thanks to the choice of p , Lemmata 1 and 2 ensure that $\|\Delta R \cdot R^{-1}\|_F \leq 2^9 m^{3.5} \chi 2^{-p} \leq 1/2$, where $\Delta R = \hat{R} - R$. Looking at the diagonal coefficients of $\Delta R \cdot R^{-1}$ leads to the result. \square

The next part of the algorithm aims at determining the column scalings to be applied to $M_B E_B$. The scalings are computed by grouping the columns of $M_B E_B$ according to the magnitudes of the \hat{r}_{ii} 's. Columns with indices in $I_\ell = [i_\ell, i_{\ell+1})$ belong to the same block. By construction, the index $i_{\ell+1}$ is the smallest $i > i_\ell$ such that $\min_{j \geq i} \hat{r}_{jj} > (8/\theta) \cdot \max_{j < i} \hat{r}_{jj}$. Let the amplitude gap between two consecutive blocks $I_{\ell-1}$ and I_ℓ be $g_\ell = (\min_{i \in I_\ell} \hat{r}_{ii}) / (\max_{i \in I_{\ell-1}} \hat{r}_{ii})$. By construction of the blocks, and $\theta \leq 1$ (see Definition 1), we have $g_\ell \geq (8/\theta) \geq 8$ for all ℓ .

At Step 7, the column scaling is set to $D = \text{diag}_i(2^{d_i})$, for each ℓ and each $i \in I_\ell$. By choice of the d_i 's, the block structure of $M_B E_B$ is preserved for $M_B E_B D^{-1}$, but the gap between two blocks gets shrunk to at most a constant. The following result is a direct consequence of Lemma 14 and of the choice of the d_i 's.

LEMMA 15. *Let R' denote the R-factor of $M_B E_B D^{-1}$. Then, for all ℓ , we have $4/3 \leq \frac{\min_{i \geq i_\ell} r'_{ii}}{\max_{i < i_\ell} r'_{ii}} \leq 32$.*

At Step 9 of Algorithm 2, an LLL-reduction algorithm is called. It is required that this algorithm does not interfere with the block structure. In most LLL-reduction algorithms, the only operations performed on the current lattice basis A are of two types: size-reductions of vectors (an integer linear combination of basis vectors \mathbf{a}_j with $j < i$ is subtracted from the basis vector \mathbf{a}_i), and swaps (two consecutive basis vectors \mathbf{a}_{i-1} and \mathbf{a}_i are exchanged). We require that swaps occur only when $r_{i,i} < r_{i-1,i-1}$. This is the case for most known LLL-reduction algorithms, including [10, 16, 14, 12, 15]. We say that these LLL-reduction algorithms are *well-behaved*. Further, if the used LLL-reduction algorithm handles only integer matrices, we may multiply matrix $M_B E_B D^{-1}$ by a power of 2 to make it integral, and reduce the scaled matrix: the computed transformation U will also be a valid LLL-reducing matrix for $M_B E_B D^{-1}$ as LLL-reducedness is invariant under basis scaling.

THEOREM 2. *Assuming the LLL-reducing algorithm used at Step 9 is well-behaved (as defined just above), Algorithm 2 is correct. In particular, the matrix $D^{-1}UD$ is unimodular and the matrix $(M_B E_B)(D^{-1}UD)$ is Ξ -LLL-reduced.*

PROOF. Using Lemma 15 and the assumption on the LLL-reducing algorithm used at Step 9, we obtain that the computed matrix U is block-upper triangular, in the following sense. For any ℓ, ℓ' , we define $U_{\ell\ell'} = (u_{ij})_{i \in I_\ell, j \in I_{\ell'}}$. Then for any $\ell > \ell'$, we have $U_{\ell\ell'} = 0$. Now, the diagonal coefficients of D are non-decreasing powers of 2, and $d_i = d_j$ when i, j belong to the same I_ℓ . As a result, the matrix $D^{-1}UD$ is integral: for $\ell \leq \ell'$, submatrix $U_{\ell\ell'}$ becomes $2^{d_{\ell'} - d_\ell} \cdot U_{\ell\ell'}$. Further, since $1 = |\det U| = \prod_\ell |\det U_{\ell\ell}|$, we obtain that all $U_{\ell\ell}$'s are unimodular. This implies that $D^{-1}UD$ is unimodular. It remains to show that $(M_B E_B)(D^{-1}UD)$ is Ξ -LLL-reduced. Let R' and R'' respectively denote the R-factors of $(M_B E_B D^{-1})U$ and $(M_B E_B)(D^{-1}UD)$. We have $r''_{ij} = r'_{ij} 2^{d_j}$, for all i, j . By $(\delta, \eta, \theta/2)$ -reducedness of $(M_B E_B D^{-1})U$, we have, for any $i \leq j$:

$$\begin{aligned} |r''_{ij}| &= |r'_{ij}| 2^{d_j} \leq (\eta 2^{d_j}) \cdot r'_{ii} + \left(\frac{\theta}{2} 2^{d_j}\right) \cdot r'_{jj} \\ &= (\eta 2^{d_j - d_i}) \cdot r'_{ii} + \frac{\theta}{2} \cdot r'_{jj}. \end{aligned} \quad (1)$$

If i and j belong to the same I_ℓ , then $d_j = d_i$ and the size-reduction condition of Definition 1 is satisfied. Otherwise, we have $i \in I_{\ell_i}$ and $j \in I_{\ell_j}$ for some $\ell_i < \ell_j$. Thanks to the assumption on the LLL-reducing algorithm (and noting that the R-factor of $M_B E_B D^{-1}$ is RD^{-1}), we have:

$$\begin{aligned} r'_{jj} &\geq \min_{t \in I_{\ell_j}} r'_{tt} \geq \min_{t \in I_{\ell_j}} (r_{tt} 2^{-d_j}) \\ &\geq \frac{2}{\theta} \max_{t \in I_{\ell_i}} (r_{tt} 2^{-d_i}) \geq \frac{2}{\theta} \max_{t \in I_{\ell_i}} r'_{tt} \geq \frac{2}{\theta} r'_{ii}. \end{aligned}$$

For the second inequality, we used the fact that for a well-behaved LLL-reduction algorithm, the minimum of the R-factor diagonal factors in a block cannot decrease. Similarly, in the fourth inequality, we used the fact that, the maximum

of the R-factor diagonal factors in a block cannot increase. For the third inequality, we used the definition of the blocks and the lower bound on the gap between two blocks, and Lemma 14. As a result, we have $r'_{jj} \geq 2^{d_j - d_i} \frac{2}{\theta} r'_{ii}$, and, by (1), we obtain that $|r''_{ij}| \leq \theta r'_{jj}$. The output basis satisfies the size-reduction condition of Definition 1.

Similarly, by reducedness of $(M_B E_B D^{-1})U$, we have:

$$\begin{aligned} \forall i : \delta (r''_{i,i})^2 &= \delta 2^{2d_i} (r'_{i,i})^2 \leq 2^{2d_i} ((r'_{i,i+1})^2 + (r'_{i+1,i+1})^2) \\ &\leq 2^{2(d_i - d_{i+1})} ((r'_{i,i+1})^2 + (r'_{i+1,i+1})^2) \\ &\leq (r'_{i,i+1})^2 + (r'_{i+1,i+1})^2, \end{aligned}$$

where we used the fact that $d_{i+1} \geq d_i$. The output basis satisfies the Lovász' conditions of Definition 1. This completes the proof of the theorem. \square

5.2 Complexity analysis

So far, we have shown that Algorithm 2 is correct. We now turn to estimating its run-time. Unless the exponents in E_B are uncommonly huge, the dominating component of the cost is the LLL-reduction of Step 9. Our aim here is to bound the bit-size of the coefficients of the matrix $M_B E_B D^{-1}$, when this matrix is viewed as an integer matrix. The algorithm takes any floating-point lattice basis as input, but the run-time bound will depend on how close is $M_B E_B$ is to be LLL-reduced. More precisely, we consider a non-singular matrix Σ and a set Ξ' of valid LLL-parameters such that $\Sigma^{-1} M_B E_B$ is Ξ' -LLL-reduced. Such a Σ always exists (take Σ such that $\Sigma^{-1} M_B E_B$ is orthonormal), but the bit-size bound to be proven will depend on the singular values of Σ . More precisely, for all ℓ , we define E_ℓ as the $|I_\ell|$ -dimensional subvector space of \mathbb{R}^m that is spanned by the columns of $M_B E_B$ with indices in I_ℓ . We define F_ℓ as the projection of E_ℓ orthogonally to $F_1 + \dots + F_{\ell-1}$, so that the column span of $M_B E_B$ is the orthogonal union of the F_ℓ 's. Now, by orthogonality, the distortion Σ acts independently on any of the F_ℓ 's. We let Σ_ℓ denote the corresponding $|I_\ell|$ -dimensional non-singular linear map. The bit-size bound of the integer matrix $M_B E_B D^{-1}$ to be given as input to an LLL-reduction algorithm at Step 9 will involve the quantity $\log \prod \kappa(\Sigma_\ell)$. By orthogonality of the F_ℓ 's, the latter is bounded from above by $\log \prod_{1 \leq i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}}$, where $\sigma_1 \geq \dots \geq \sigma_m$ are the singular values of Σ . The following lemma provides a bound on the amplitude of the \hat{r}_{ii} 's within a block.

LEMMA 16. *With α' as in Lemma 4 (for Ξ'), for any ℓ , we have $\frac{\max_{i \in I_\ell} \hat{r}_{ii}}{\min_{i \in I_\ell} \hat{r}_{ii}} \leq 3(8\alpha'/\theta)^{|I_\ell|} \kappa(\Sigma_\ell)$.*

PROOF. We prove that for any $i, j \in I_\ell$, we have $\hat{r}_{jj}/\hat{r}_{ii} \leq 3(8\alpha'/\theta)^{|I_\ell|} \kappa(\Sigma_\ell)$. Suppose first that $j \leq i$. Then, by Lemma 4, we have $r'_{jj}/r'_{ii} \leq (\alpha')^{i-j}$, where R' denote the R-factor of C . By Lemma 3, we obtain that $r_{jj}/r_{ii} \leq (\alpha')^{i-j} \kappa(\Sigma_\ell)$. Finally, by Lemma 14, we obtain that $\hat{r}_{jj}/\hat{r}_{ii} \leq 3(\alpha')^{i-j} \kappa(\Sigma_\ell)$. Suppose now that $j > i$. If $\hat{r}_{ii} = \max_{t \geq i} \hat{r}_{tt}$, then the bound holds since the right hand side is ≥ 1 . Otherwise, from the definition of blocks, there exists some $i' \in I_\ell$ with $i' > i$, such that $\hat{r}_{i'i'} \leq (8/\theta) \cdot \hat{r}_{ii}$. Applying the same idea to i' we get $\hat{r}_{i''i''} \leq (8/\theta)^t \cdot \hat{r}_{ii}$, with $i'' = i_{\ell+1} - 1$, where $t \leq |I_\ell|$ is the number of times this recursion is applied. Since $j \leq i''$, we conclude that $\hat{r}_{jj} \leq 3(\alpha')^{i''-j} \kappa(\Sigma_\ell) \hat{r}_{i''i''} \leq 3(8\alpha'/\theta)^{|I_\ell|} \kappa(\Sigma_\ell) \hat{r}_{ii}$ (using the first part of the proof). \square

From here, we can derive a bound on the amplitude of the diagonal coefficients of the R-factor of $M_B E_B D^{-1}$.

LEMMA 17. Let R' denote the R-factor of $M_B E_B D^{-1}$. Then $\frac{\max_i r'_{ii}}{\min_i r_{ii}} \leq c^n \cdot \prod \kappa(\Sigma_\ell)$, for some c depending only on Ξ' .

PROOF. Thanks to Lemma 15, we have that for all ℓ , $\min_{i \in I_\ell} r'_{ii} \leq 32 \cdot \max_{i \in I_{\ell-1}} r'_{ii}$.

Using Lemma 14 we then translate the bound of Lemma 16 for $I_{\ell-1}$ in terms of the r'_{ii} 's rather than the \hat{r}_{ii} 's, hence get: $\min_{i \in I_\ell} r'_{ii} \leq 100(8\alpha'/\theta)^{|I_\ell|} \kappa(\Sigma_\ell) \cdot \min_{i \in I_{\ell-1}} r'_{ii}$. Taking the product over all values of ℓ leads to the result. \square

We can now prove the following bit-size bound for the input to the LLL-reduction algorithm.

THEOREM 3. Assume that the input matrix M_B is integral with $\|M_B\| \leq 2^p$. Let $\Sigma \in \mathbb{R}^{m \times m}$ be non-singular, such that $\Sigma^{-1} M_B E_B$ is Ξ' -LLL-reduced, for some valid set Ξ' of LLL-parameters. Let $C = M_B E_B D^{-1}$ be the matrix given as input to an LLL-reduction algorithm at Step 9 of Algorithm 2. Then there exists a constant c such that $C \in 2^{-k} \mathbb{Z}^{m \times n}$ for some k satisfying:

$$\log \|C\| - k \leq c \cdot n + 2p + 2 \log \prod_{1 \leq i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i+1}},$$

where $\sigma_1 \geq \dots \geq \sigma_m$ denote the singular values of Σ .

PROOF. Lemma 11 gives: $\max \|\mathbf{c}_i\| \leq \text{cond}(C) \max r'_{ii}$, where R' denotes the R-factor of C . Further, by Lemmata 7 and 9, there exists a constant c_1 such that $\text{cond}(C) \leq \kappa(\Sigma) \cdot c_1^n$. Using Lemma 17 and the fact that $\min \|\mathbf{c}_i\| \geq \min r'_{ii}$, we obtain that

$$\frac{\max \|\mathbf{c}_i\|}{\min \|\mathbf{c}_i\|} \leq c_2^n \kappa(\Sigma) \prod_{\ell} \kappa(\Sigma_\ell) \leq c_2^n \prod_{1 \leq i \leq \lfloor n/2 \rfloor} \frac{\sigma_i^2}{\sigma_{m-i+1}^2},$$

for some constant c_2 . Now, using the assumption that M_B is integral of norm $\leq 2^p$, we have

$$\frac{\max 2^{e_i - d_i}}{\min 2^{e_i - d_i}} \leq 2^p \cdot \frac{\max \|\mathbf{c}_i\|}{\min \|\mathbf{c}_i\|} \leq 2^p c_2^n \prod_{1 \leq i \leq \lfloor n/2 \rfloor} \frac{\sigma_i^2}{\sigma_{m-i+1}^2}.$$

To complete the proof, we note that the entries of C are sums of powers of 2, and we use that $2^p \frac{\max 2^{e_i - d_i}}{\min 2^{e_i - d_i}}$ bounds from above the ratio between the smallest and the largest entries. \square

As our bound applies to any Σ such that $\Sigma^{-1} M_B E_B$ is LLL-reduced, we are interested in the existence of such a Σ with a small $\log \prod_{i \leq \lfloor n/2 \rfloor} \frac{\sigma_i}{\sigma_{m-i}}$.

6. PRACTICAL CONSIDERATIONS

Several important points deserve further investigations, especially from a practical point of view. The sharpness of the sufficient bound $2 \log \text{cond}(R) + n(1 + \epsilon) + O(\log m)$ on the input precision is unclear. Studying heuristic values for the different quantities to choose, e.g. for the constant c of Lemma 13, remains to be done. In the same vein, understanding the impact of the chosen precision and the input basis structure on the output parameters $\Xi = (\delta, \eta, \theta)$ is an interesting problem. For an idea of practical accelerations that can be obtained thanks to the scaling, we may refer to [1, Se. 5]. A extensive experimental study should be made.

Acknowledgements. We thank Mark Watkins for pointing out the reference [5, Se. 2.6.1], and Cong Ling for valuable discussions. Part of this work was undergone while the second and third authors were visiting the University of Sydney, whose hospitality is gratefully acknowledged. This work has been supported in part by ERC Starting Grant ERC-2013-StG-335086-LATTAC, and by the French National Research Agency Grant ANR-11-BS02-013-HPAC.

7. REFERENCES

- [1] J. Bi, J.-S. Coron, J.-C. Faugère, P. Nguyen, G. Renault, and R. Zeitoun. Rounding and chaining LLL: Finding faster small roots of univariate polynomial congruences. In *Proc. PKC'14, Buenos Aires, Argentina, LNCS 8383, 160–168*. Springer, 2014.
- [2] J. Buchmann. Reducing Lattice Bases by Means of Approximations. In *Proc. ANTS, LNCS 877, 160–168*. Springer, 1994.
- [3] X.-W. Chang and C.C. Paige. Componentwise perturbation analyses for the QR factorization. *Numer. Math.*, 88:319–345, 2001.
- [4] X.-W. Chang, D. Stehlé, and G. Villard. Perturbation analysis of the QR factor R in the context of LLL lattice basis reduction. *Math. Comp.*, 81(279):1487–1511, 2012.
- [5] H. Cohen. *A Course in Computational Algebraic Number Theory*, volume 138 of *Graduate Texts in Mathematics*. Springer, Berlin, 1996.
- [6] N. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 2nd edition, 2002.
- [7] M. van Hoeij and A. Novocin. Gradual sub-lattice reduction and a new complexity for factoring polynomials. *Algorithmica*, 63(3):616–633, 2012. Preliminary version: *Proc. LATIN, 539–553, 2010*.
- [8] E. Kaltofen. On the complexity of finding short vectors in integer lattices computer algebra. In *Proc. EUROCAL, LNCS 162, 236–244*. Springer, 1983.
- [9] B. A. Lamacchia. Basis reduction algorithms and subset sum problems. Technical report, SM thesis, Massachusetts Inst. Technol, 1991.
- [10] A. Lenstra, H. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.
- [11] I. Morel, D. Stehlé, and G. Villard. From an LLL-reduced basis to another. *ACM Commun. Comput. Algebra*, 42(3):142–143, February 2009. ISSAC'08 poster.
- [12] I. Morel, D. Stehlé, and G. Villard. H-LLL: Using Householder inside LLL. In *Proc. ISSAC, Seoul, Republic of Korea, 271–278*. ACM, 2009.
- [13] H. Najafi, M. Jafari, and M.-O. Damen. On adaptive lattice reduction of correlated fading channels. *IEEE Trans. Commun.*, 59(5):1224–1227, 2011.
- [14] P. Nguyen and D. Stehlé. An LLL algorithm with quadratic complexity. *SIAM J. Comput.*, 39(3):874–903, 2009. Preliminary version: *Proc. EUROCRYPT, LNCS 3494, 215–233, 2005*.
- [15] A. Novocin, D. Stehlé, and G. Villard. An LLL-reduction algorithm with quasi-linear time complexity: extended abstract. In *Proc. STOC, San Jose, USA, 403–412*. ACM, 2011.
- [16] A. Schönhage. Factorization of univariate integer polynomials by Diophantine approximation and an improved basis reduction algorithm. In *Proceedings of ICALP, LNCS 172, 436–447*. Springer, 1984.
- [17] A. van der Sluis. Condition numbers and equilibration of matrices. *Numer. Math.*, 14:14–23, 1969.
- [18] A. Storjohann. The shifted number system for fast linear algebra on integer matrices. *J. Compl.*, 21(4):609–650, 2005.
- [19] H. Zha. A componentwise perturbation analysis of the QR decomposition. *SIAM J. Matrix Anal. Appl.*, 14(4):1124–1131, 1993.