



THÈSE
en vue de l'obtention du grade de Docteur, délivrée par
l'ÉCOLE NORMALE SUPERIEURE DE LYON

École Doctorale N°512
INFOMATHS

Discipline : Informatique

Soutenue publiquement le 26/11/2024, par :

Joël Felderhoff

**Hardness of Structured Lattice Problems
for Post-Quantum Cryptography**

**Difficultés de Problèmes de Réseaux Structurés
pour la Cryptographie Post-Quantique**

Devant le jury composé de :

CASTRYCK, Wouter	Research Expert Katholieke Universiteit Leuven	Rapporteur
THOMÉ, Emmanuel	Directeur de Recherche Inria Nancy	Rapporteur
BOUDGOUST, Katharina	Chargée de Recherche LIRMM, Montpellier	Examinatrice
KIRSHANOVA, Elena	Chercheuse Technology Innovation Institute	Examinatrice
SALVY, Bruno	Directeur de Recherche Inria Lyon	Directeur de thèse
STÉHLE, Damien	Chercheur CryptoLab	Co-encadrant
HANROT, Guillaume	Chercheur CryptoLab	Co-encadrant

À Mamie Dolly qui aurait sans doute été catastrophée par mon accent et mes fautes d'accords lors de la soutenance mais - j'espère - quand même été un peu fière.

À mes parents, parce que je sais tout ce que je leur dois.

À toutes les personnes dont j'ai la chance et le bonheur d'être aimé.

Foreword

My official PhD advisor is Bruno Salvy, but since he is thematically far from what this manuscript is about, I need to comment on my supervisory context.

I started my PhD in late 2021 under the supervision of Damien Stehlé. He left the ENS de Lyon for CryptoLab in early 2023, a little more than a year after. At this stage, Guillaume Hanrot took over my supervision. When Guillaume also left for CryptoLab in late 2023, the official PhD advisor became Bruno Salvy.

Bruno Salvy did provide an administrative supervision (and precious advices) but all the scientific one was done by Damien and Guillaume even after they left for Cryptolab. In the scientific sense, my PhD advisors are Damien Stehlé and Guillaume Hanrot, even if they are not noted as such on the front page of this manuscript for administrative reasons.

Acknowledgements/Remerciements

My PhD position was funded by the Direction Générale de l'Armement (Pôle de Recherche CYBER).

Je souhaite pour commencer remercier mes encadrants. Merci Damien de m'avoir recruté à l'issue de mon stage de PLR (et d'avoir permis que ma thèse commence en septembre malgré les difficultés administratives), merci pour ton encadrement, tes conseils et nos discussions plus informelles, et aussi d'avoir été mon tuteur dès mon entrée à l'ENS, ton aide et tes conseils à cette époque ont grandement contribué à mes choix de sujets et de carrière.

Merci Guillaume pour ton encadrement à partir de ma 2^e année. Ça a été un plaisir d'ajouter des aspects plus « Théorie Analytique des Nombres » à mon travail, et de manière générale d'échanger (et de digresser) régulièrement avec toi. Je veux aussi te remercier pour les conseils que tu as pu me donner toi aussi dès mes premières années à l'ENS, où mon profil de « vrai informaticien/faux matheux » (ou l'inverse) me faisaient me poser beaucoup de questions. Merci à tous les deux d'être restés disponible pour répondre à mes questions et m'encadrer scientifiquement même après votre départ du LIP.

Pour finir, et même si je sais que tu ne te considères pas comme mon encadrant, merci Bruno de m'avoir accepté comme doctorant lors du départ de Guillaume et Damien. Merci pour tes conseils et tes retours lors de la préparation de la soutenance, et de manière générale d'avoir été là pour la gestion administrative de la fin de thèse.

Now, I want to thank Emmanuel Thomé and Wouter Castryck for reviewing this manuscript and for their helpful comments. I also want to thank a lot Katharina Boudgoust and Elena Kirshanova for accepting to be part of my jury.

Écrire des remerciements pour un travail qui a duré 3.5 ans est une tâche condamnée à la non-exhaustivité. J'ai essayé de mettre le plus de monde possible et j'ai évidemment échoué, toutes mes excuses.

Je voudrais commencer par remercier Alice Pellet--Mary. J'ai été très heureux de travailler avec toi pendant ces 3.5 ans, et ton oreille attentive et tes conseils durant la thèse m'ont énormément aidé. J'espère vraiment pouvoir continuer à travailler avec toi dans la suite de ma carrière.

Merci beaucoup à Alid, Arthur, Danaé, Mahshid, Maman, Pouria et Xavier qui ont participé à la relecture de l'introduction de ce manuscrit (et à la chasse aux fautes).

Merci à l'équipe des MALIPs pour toute l'aide qu'elles m'ont apporté durant mes années de thèse. Le LIP a énormément de chance de vous avoir. En particulier un énorme merci à Chiraz. Merci à mes co-doctorant-es (qui, pour certain-es, sont maintenant docteur-es!) de la « sous-équipe crypto » d'AriC : Arthur (si, si, t'es toujours AriC dans mon cœur), Calvin, Emily (yeah,

you don't do crypto, but you have long been adopted), Julien, Mahshid et Pouria¹. Je retiendrai les voyages en conférence, les moresques, le support quand ça n'allait pas bien, les cafés, la procrastination et les batailles de nerf aléatoires au bureau. Ça a été génial de faire ma thèse avec vous. Et puisque certain-es choisissent de travailler sur autre chose que de la cryptographie (étrange, mais je vais pas juger), merci énormément à Adrien, Amélie, Alaa, Esther, Johann, Léo, Louis, Meriem, Thaïs ainsi qu'à tous-tes les autres avec qui j'ai eu le plaisir de papoter, râler, débattre et rigoler au coin café.

Un autre grand merci aux autres collègues avec qui j'ai eu la chance de travailler, d'enseigner, ou simplement de discuter pendant ces quelques années au LIP, et notamment à Benjamin, Claude-Pierre, Cyril, Daniel, Fabrice, Gilles, Jean-Michel, Michaël, Nathalie, Nicolas et Vincent.

A bit further from Lyon, I want to thank Léo, Koen and Yael. I am very happy to have had the chance to work with you, and I really hope we continue working together in the future!

En m'éloignant du bureau pour mes remerciements, je me dois de remercier chaleureusement l'équipe de Maths en Jeans, et en tout particulier Stéphanie pour la bouffée d'air frais que cette activité a représenté durant la thèse. Merci aussi à tous-tes les élèves de Jean Perrin et de la Tourette que j'ai eu le plaisir d'encadrer. N'arrêtez jamais d'être aussi formidables.

Maintenant que je suis sorti du bureau, il faut que je parle des potes. J'ai la chance incroyable d'avoir été très entouré dans ma thèse et d'avoir trop de bons souvenirs avec chacun-es d'entre vous pour pouvoir tout dire sans dépasser ma limite de page. Merci George, Janelle, Malo, Simon, Xavier et les chatons, vous savez pourquoi. Merci à Aaren, Adrien, Alban, Alexandra, Alix, Alyd, Antoine, Arnaud, Avril, Bertrand, Charline, Corentin, Élodie, Florian, Florine, Gabriel, Gabrielle, Garance, Guillaume, Henry, Juliette, Lambert, Laureline, Marie, Mista, Morgan, Nattes, Octave, Solène, Thomas, Yohann et tous-tes les autres pour votre affection et votre soutien constant pendant ces années de thèse et avant.

Merci aux potes de la CRF69, en particulier à Ambre, Aurore, Cecile, Margot, Maya, Mickaël, Roxane, Sarah, ainsi qu'à tous-tes les potes de la formation CI. Pouvoir travailler sur des choses plus pratiques que *la difficulté des problèmes de réseaux structurés* pendant 3 ans m'a aidé à garder un minimum les pieds sur terre et a été un vrai plaisir.

Je dois aussi remercier les pauvres hères qui m'ont supporté au quotidien pendant la thèse : merci Dana, Marlysa et Youssef. Vous êtes les meilleurs colocs du monde et vivre avec vous a été fantastique.

Évidemment, merci à ma famille pour tout leur soutien. Merci Papa, Maman, Noé, Laura (j'espère que vous apprécierez ma petite référence lors de la soutenance), Mamette, Papet, Papi, Tatie, Tonton, Léna et tous les autres. Merci pour votre écoute et vos encouragements, et toutes mes excuses pour les descriptions incompréhensibles de mon travail aux repas de famille².

Finalement, merci Danaé. Ce n'est pas vraiment possible pour moi de mettre en mots à quel point ta présence, ton soutien, nos discussions, et de manière générale notre vie commune ont été importantes pour moi pendant ces années de thèse. Tu es fantastique et j'ai hâte de voir ce que l'avenir nous réserve. Je t'aime.

Merci également à Claire Martinod et à Guillaume Labeille.

Je voudrais pour finir, remercier, les personnels de ménage, de cuisine et de gardiennage de l'ENS, qui m'ont permis de travailler dans de bonnes conditions, les artistes (notamment musicaux) dont j'ai profité des créations, ainsi que les auteurs et autrices des travaux que je cite dans ce manuscrit.

¹You know *exactly* what dwells beneath Zürichsee...

²Non pas que je prévois de m'arrêter.

Contents

Contents	7
I Introduction	9
I.1 Introduction (Français)	9
I.2 Introduction (English)	24
II Preliminaries	39
II.1 Lattices	39
II.2 Number Theory	44
II.3 Modules	49
II.4 Computational Problems	51
II.5 Probabilities	54
III Counting Small Ideals	57
III.1 Preliminaries	57
III.2 Bounds on the Dedekind's Zeta Function of K	59
III.3 Bounds on the integral	60
III.4 Bounding the ideal-counting function	63
IV Ideal-SVP is Hard for Small-Norm Uniform Prime Ideals	65
IV.1 Introduction	65
IV.2 Preliminaries	68
IV.3 Self-Reducibility of id-HSVP to Inverses	70
IV.4 The Sampling Set	77
IV.5 Wrapping Up	85
IV.6 NTRU with Polynomial Modulus	87
V On Module Unique-SVP and NTRU	91
V.1 Introduction	91
V.2 Preliminaries	96
V.3 New Tools on Module Lattices	97
V.4 From mod-uSVP ₂ to NTRU	99
V.5 Randomization of Rank-2 Modules with Gaps	103
V.6 Random Self-Reducibility of Module uSVP	110
VI Conclusion and Perspectives	113
VI.1 Summary of Contributions	113
VI.2 Perspective and open problems	114

Bibliography	117
A Appendices of Chapter II	123
A.1 Missing Proofs	123
B Appendices of Chapter III	125
B.1 Analysis proofs	125
B.2 Proof of Theorem III.1.8	128
B.3 Bounds for $ \ln \zeta_K $	128
C Appendices of Chapter IV	129
C.1 Proof of Lemma IV.2.3	129
C.2 Proof of Theorem IV.2.4	132
D Appendices of Chapter V	139
D.1 Properties of the Rényi Divergence	139
D.2 Missing proofs from Section V.2	140
D.3 Missing Proofs from Section V.3	145
D.4 Missing Proofs from Section V.4	147
D.5 Removing $\zeta_K(2)$ from Theorem V.4.1	154
D.6 Missing Proofs from Section V.5	162
D.7 Missing Proofs from Section V.6	169

Chapter I

Introduction

I.1 Introduction (Français)

« Du coup... T'es mathématicien
ou informaticien ? »

J.M. Felderhoff (mon père), 2022

La notion de « communication sécurisée » recouvre un grand nombre de problématiques, comme l'authentification des messages (garantir la provenance d'un message reçu) ou le chiffrement de ceux-ci (rendre le contenu d'un message inintelligible pour d'autres que son destinataire), et on en trouve des proto-exemples datant de l'Antiquité. De nos jours, les ordinateurs modernes rendent caduques les techniques naïves de sécurisation (telles que les chiffrements mono-alphabétiques, où on remplace simplement une lettre par une autre).

La démocratisation de l'outil informatique et d'internet s'est accompagnée de la mise en place de nombreux protocoles de chiffrement et de signatures numériques, tels que le protocole TLS (standard, utilisé notamment dans le HTTPS), le standard OpenPGP (utilisé pour signer des emails) ou encore les protocoles de chiffrement bout-en-bout (utilisés dans WhatsApp ou Signal). Dans ce contexte d'utilisation de grande ampleur, il est nécessaire de trouver des moyens de garantir la sécurité des communications utilisant des protocoles cryptographiques.

Nous nous placerons dans ce manuscrit dans le cadre de la cryptographie à clé publique, qui permet de communiquer de manière sécurisée lorsque les deux parties ne peuvent pas partager une clé au préalable (ce qui est régulièrement le cas des échanges sur internet). Dans ce paradigme, on distingue clé secrète (qui n'est à disposition que d'une seule partie) et clé publique (à disposition de tout le monde). La sécurité des protocoles à clés publiques repose sur la difficulté de deviner une clé secrète en ayant connaissance des données publiques (clé publique et messages transitant sur le réseau par exemple). De nos jours, garantir la sécurité de protocoles de cryptographie à clé publique se fait au moyen de preuves de sécurité.

I.1.1 Garantir la sécurité d'un protocole

Prenons l'exemple d'un protocole de chiffrement permettant à deux parties d'échanger des messages de manière incompréhensible pour une tierce personne. Prouver la sécurité d'un tel protocole se fait en trois étapes.

La première est la définition de l’adversaire, c’est-à-dire l’entité abstraite contre qui on veut garantir la sécurité de notre système, par exemple une agence de renseignement ou une entreprise (légal ou non) voulant revendre des données... Les questions à se poser sont typiquement :

- Quel serait l’objectif de l’adversaire pour « casser » mon protocole ? Pour notre exemple, cela pourrait être de déchiffrer un message, ou de distinguer un chiffré d’une chaîne de bits aléatoires.
- À quelle puissance de calcul l’adversaire a-t-elle accès ? Pendant combien de temps veut-on lui résister ?
- Peut-elle interagir avec le système cryptographique ? Par exemple, peut-elle envoyer de faux messages et observer le comportement de son interlocutrice ? A-t-elle accès à une partie de la clé secrète ?

La seconde étape est de définir une ou plusieurs « hypothèses de sécurité ». Ce sont des énoncés mathématiques de la forme : « il est impossible de résoudre tel problème en un temps raisonnable ». Un exemple classique est la factorisation : « Étant donné un grand nombre N , il est impossible en un temps raisonnable de trouver p et q différents de 1 tel que $N = p \cdot q$ ». ¹

La troisième étape est de faire le lien entre les deux. En pratique, cela consiste à prouver un énoncé mathématique de la forme : « Supposons qu’il existe un adversaire cassant notre protocole, alors il existe aussi forcément un algorithme qui casse l’hypothèse de sécurité ». Ce genre d’énoncé est appelé une réduction de sécurité. Il faut l’interpréter comme la formalisation mathématique du fait que « tant que l’hypothèse de sécurité est vérifiée, le protocole est sécurisé contre ce type d’adversaires ». On dit qu’étudier la sécurité du protocole contre cet adversaire se *réduit* à étudier la validité de l’hypothèse de sécurité.

I.1.2 L’adversaire quantique et la cryptographie Post-Quantique

Dans ce manuscrit, nous nous plaçons sous l’hypothèse d’une adversaire ayant accès à un calculateur quantique (contrairement à nous). L’informatique quantique s’intéresse au calcul par des ordinateurs quantiques (par opposition à nos ordinateurs « classiques », fonctionnant avec des transistors). Il n’est pas question dans cette introduction d’expliquer précisément le fonctionnement de possibles ordinateurs quantiques. Il suffit de dire qu’ils effectuent des calculs d’une manière différente d’un ordinateur classique dans le même sens qu’une calculatrice mécanique fait des calculs différemment qu’une calculatrice numérique. Cette différence fait que certains problèmes qui étaient supposés difficiles pour des ordinateurs classiques sont résolubles efficacement avec des ordinateurs quantiques.

En particulier, l’algorithme quantique de Shor [Sho94] permet de résoudre avec une quantité raisonnable de ressources le problème de la factorisation et celui du logarithme discret. Les protocoles de sécurité basés sur la difficulté de ces problèmes sont alors caducs dans le cas où l’adversaire dispose d’un calculateur quantique. Étant donné que les protocoles les plus largement déployés (par exemple TLS et OpenPGP) dépendent fortement de la difficulté à résoudre ces problèmes, la possibilité (largement débattue...) de l’apparition d’un ordinateur quantique dans les années à venir a poussé les instituts de standardisation, les autorités étatiques et les industriels à augmenter l’effort de recherche autour d’hypothèses de sécurité résistantes aux ordinateurs quantiques.

¹Bien sûr, la notion de « grand » et de « temps raisonnable » doivent être définies plus précisément, voir Section II.4.




FIGURE I.1 : Couche de graphène au niveau atomique [Ale09] et réseau euclidien associé.

On peut citer notamment le processus de standardisation, sous forme d'une compétition, de cryptographie post-quantique du NIST (l'institut de standardisation et des technologies aux États-Unis) commencé en 2016 et terminé en 2022 [NIST]. À l'issue de cette compétition, il apparaît quatre grandes familles de protocoles de chiffrements semblant résister aux ordinateurs quantiques (les hypothèses sur lesquels ils sont basés sont dites *post-quantiques*). Les protocoles basés sur les systèmes polynomiaux, sur les codes correcteurs d'erreurs, sur les isogénies entre courbes elliptiques et sur les réseaux euclidiens. Le travail effectué durant cette thèse concerne les hypothèses de sécurité reliées aux réseaux euclidiens.

I.1.3 Réseaux euclidiens

Le problème du plus court vecteur.

Informellement, un réseau euclidien (on parlera souvent simplement de réseau) est un ensemble infini de points de l'espace répartis de manière régulière. On peut par exemple s'en servir en 2 ou 3 dimensions pour représenter la répartition des atomes dans un cristal (voir Figure I.1). En cryptographie, on utilise des réseaux de haute dimension ($n \approx 500$) pour leurs propriétés algorithmiques.

Soit $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{n \times n}$ une matrice (invertible), qu'on appelle une base. Le réseau euclidien défini par cette matrice, noté $\mathcal{L}(\mathbf{B})$ est l'ensemble des combinaisons entières des vecteurs colonnes de $(\mathbf{b}_i)_{1 \leq i \leq n}$. Mathématiquement, on écrira

$$\mathcal{L}(\mathbf{B}) = \mathbf{B} \cdot \mathbb{Z}^n = \left\{ \sum_{i=1}^n \mathbf{b}_i \cdot x_i, (x_i)_{1 \leq i \leq n} \in \mathbb{Z}^n \right\},$$

Un exemple de réseau euclidien L est donné en Figure I.2. Ce réseau particulier est généré par les vecteurs en bleus sur la figure, soit la matrice

$$\begin{pmatrix} 1.1 & -0.1 \\ -0.1 & 1 \end{pmatrix}.$$

Il est à noter que les vecteurs bleus ne sont pas les seuls à générer L , c'est aussi le cas des vecteurs rouges : un réseau possède plusieurs bases.

De nombreux problèmes algorithmiques sont associés aux réseaux euclidiens. Certains sont utilisables pour construire de la cryptographie (une présentation des principaux ainsi que de leurs




FIGURE I.2 : Un exemple de réseau euclidien.

relations peut être trouvée dans une revue de littérature par Peikert [Pei16]). Dans ce manuscrit, nous étudierons le problème consistant à trouver un ou plusieurs vecteurs les plus courts possibles dans un réseau euclidien, étant donné une base de celui-ci. On définit le problème SVP_γ (Shortest Vector Problem) comme suit :

Definition I.1.1 (SVP_γ). Pour $\gamma \geq 1$, étant donné une matrice inversible $\mathbf{B} \in \mathbb{Z}^{n \times n}$, le problème SVP_γ demande de trouver $\mathbf{v} \in \mathcal{L}(\mathbf{B}) \setminus \{\mathbf{0}\}$ tel que

$$\|\mathbf{v}\| \leq \gamma \cdot \lambda_1(\mathcal{L}(\mathbf{B})),$$

où $\lambda_1(\mathcal{L}(\mathbf{B}))$ est la taille du plus petit vecteur non-nul du réseau euclidien $\mathcal{L}(\mathbf{B})$.

En particulier, SVP_1 (aussi nommé « exact-SVP ») demande de trouver un des vecteurs les plus courts de $\mathcal{L}(\mathbf{B})$. Dans le cas du réseau de la Figure I.2, le vecteur vert serait une solution à SVP_1 . Ce problème est difficile à résoudre. En effet, les meilleurs algorithmes connus (que ce soit classiques ou quantiques) résolvant SVP_1 en dimension n nécessitent soit de faire de l'ordre de $n^{O(n)}$ opérations (algorithmes d'énumération, voir par exemple [FP85, Kan87, HS07]), soit de l'ordre de $2^{O(n)}$ opérations, au prix d'une quantité de mémoire de $2^{O(n)}$ bits (algorithmes de crible [AKS01] ou basé sur les cellules de Voronoï [MV13]). Cet état de l'art rend SVP_1 insoluble avec des ressources raisonnables (disons moins de 10^{10} ans en utilisant l'ensemble des ordinateurs du monde) dès que la dimension n devient plus grande que quelques centaines (le plus grand record enregistré sur [Nam] au moment de la rédaction de ce manuscrit est $n = 190$).

Comparer les problèmes algorithmiques. La notion centrale pour comparer la difficulté de deux problèmes algorithmiques est la notion de réduction. En théorie de la complexité, on dit qu'un problème \mathcal{B} est aussi ou plus difficile qu'un problème \mathcal{A} si, étant donné un algorithme efficace résolvant le problème \mathcal{B} , on peut écrire un algorithme efficace² (appelé réduction) pour résoudre le problème \mathcal{A} . On dira que le problème \mathcal{A} se réduit au problème \mathcal{B} . Deux problèmes \mathcal{A} et \mathcal{B} sont dits équivalents si \mathcal{A} se réduit à \mathcal{B} et \mathcal{B} se réduit à \mathcal{A} .

²S'exécutant en temps polynomial en la taille de son entrée.




FIGURE I.3 : Compromis temps/approximation pour SVP avec BKZ en dimension n . En vert les facteurs d'approximation utilisés en cryptographie.

Plus γ est grand, plus SVP_γ est facile à résoudre, au sens où une solution de SVP_γ est également une solution de $\text{SVP}_{\gamma'}$ si $\gamma' \geq \gamma$. Il a été démontré que SVP_1 est NP-difficile³, ce qui implique qu'il est plausible qu'un algorithme efficace pour le résoudre en toute généralité n'existe pas, y compris avec un calculateur quantique⁴. Sa difficulté en fonction de γ peut être décrite plus précisément comme suit. Si n est la dimension du réseau et β un entier entre 1 et n , il existe un algorithme (l'algorithme BKZ [SE94]) résolvant SVP_γ pour $\gamma = 2^{\tilde{O}(n/\beta)}$ en temps proportionnel à 2^β (le résultat est volontairement simplifié dans cette introduction : pour les conditions sur β et une valeur précise de γ , voir le Lemme II.1.14). Cet algorithme nous permet de donner un gradient de difficulté pour SVP_γ en fonction de γ , représenté en Figure I.3.

Cela dit, il n'a pas été prouvé que SVP_γ est NP-difficile pour les facteurs d'approximation γ utilisés dans les constructions cryptographiques. Il est même peu probable que ce soit le cas, car il a été prouvé [AR05] que pour $\gamma = \sqrt{n}$, SVP_γ appartient à la classe de complexité $\text{NP} \cap \text{coNP}$. La construction de cryptographie basée sur la difficulté d'un problème NP-difficile est un problème ouvert à l'heure actuelle.

Un autre problème central dans la cryptographie à base de réseaux euclidiens est le problème SIVP_γ , qui demande, grossièrement, de trouver une famille de rang plein de petits vecteurs dans un réseau donné (leur taille est contrôlée par γ comme pour SVP_γ). Il a été prouvé que SIVP_1 et SVP_1 sont équivalents [GMSS99, Mic08] et que $\text{SIVP}_{\sqrt{n} \cdot \gamma}$ se réduit à SVP_γ [Ste15].

Problèmes cas-moyen

Un protocole cryptographique à clé publique basé sur les réseaux euclidiens fonctionne typiquement comme suit : on tire un réseau L au hasard (selon une certaine distribution de probabilité) avec de petits vecteurs de celui-ci ; on publie ensuite une base du réseau comme clé publique et

³Pour les réductions randomisées [Ajt98], ou pour la norme ℓ_∞ [Emd].

⁴La NP-difficulté et sa relation avec le calcul quantique est un large sujet que nous n'avons pas le temps d'introduire en détails ici. La lectrice intéressée est redirigée vers [AB09].

on utilise les petits vecteurs comme clé secrète. La sécurité du système repose alors sur l'impossibilité pour l'adversaire de trouver des petits vecteurs du réseau étant donné sa base, c'est-à-dire sur la difficulté de SVP_γ pour le réseau qui a été tiré. Il faut noter que dans ce cas, la sécurité du protocole ne repose pas exactement sur la difficulté de SVP_γ , qui consiste à trouver un vecteur court dans *n'importe quel réseau*, mais sur la difficulté de SVP_γ sur un réseau tiré au hasard selon cette distribution de probabilité. Résoudre SVP_γ sur un réseau au hasard est un problème dit *cas-moyen* (ou moyen-cas). Il est plus facile que résoudre SVP_γ sur tous les réseaux, que l'on qualifie de problème *pire-cas*.

Un des intérêts de la cryptographie à base de réseaux euclidiens est que la difficulté de SVP_γ sur certaines de ces distributions peut être reliée à la difficulté de SVP_γ dans le pire-cas. Nous en décrivons deux dans cette introduction.

Short Integer Solution (SIS). Le premier exemple de problème cas-moyen que nous présentons est le problème Short Integer Solutions, introduit en 1996 par Miklós Ajtai [Ajt96]. Il consiste à résoudre un système linéaire avec une condition sur la taille de la solution. Dans tout ce qui suit, \mathbb{Z}_q désignera l'anneau $\mathbb{Z}/q\mathbb{Z}$.

Definition I.1.2 ($\text{SIS}_{q,n,m,\beta}$). Soit $q \geq 2$, $n \geq 1$, $m \geq n \log(q)$, et $\beta \geq \sqrt{n \log(q)}$. On définit le problème $\text{SIS}_{q,n,m,\beta}$ comme suit. Étant donné une matrice \mathbf{A} une matrice uniforme dans $\mathbb{Z}_q^{n \times m}$, trouver un vecteur $\mathbf{x} \in \mathbb{Z}^m \setminus \{\mathbf{0}\}$ tel que $\|\mathbf{x}\| \leq \beta$ tel que

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}.$$

On peut noter que l'entrée de ce problème est une matrice tirée uniformément sur $\mathbb{Z}_q^{n \times m}$: c'est un problème cas-moyen. On peut relier SIS à un problème de réseaux euclidiens en notant que si $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ une matrice, alors l'ensemble des solutions possibles de SIS sur l'entrée \mathbf{A} est l'ensemble

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m, \mathbf{A} \cdot \mathbf{x} = 0 \pmod{q}\},$$

qui est un réseau euclidien. Trouver un petit vecteur dans $\Lambda_q^\perp(\mathbf{A})$ est équivalent à trouver une solution de SIS. On peut donc reformuler le problème SIS comme « résoudre SVP dans un réseau $\Lambda_q^\perp(\mathbf{A})$ pour \mathbf{A} uniforme ».

Lors de l'introduction [Ajt96] de SIS, sa difficulté est reliée à SVP_{n^c} pour un certain $c > 1$. Des résultats ultérieurs [MR04, GPV08, MP13] ont précisé cette dépendance : si $m = \text{poly}(n)$ et $q \geq \beta \cdot n^\varepsilon$ pour $\varepsilon > 0$, alors SVP_γ en dimension n se réduit à $\text{SIS}_{q,n,m,\beta}$ pour $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$.

Learning with Errors (LWE). Un autre problème cas-moyen central dans la cryptographie à base de réseaux moderne est le problème Learning With Errors⁵, introduit par Oded Regev en 2005 [Reg05].

Definition I.1.3. Soient $1 \leq n \leq m$ et $q \geq 2$ trois entiers, et $\alpha \in [0, 1]$ un paramètre réel. Le problème $\text{LWE}_{n,m,q,\alpha}$ demande de distinguer avec une probabilité $\geq 2/3$ entre les distributions

$$(\mathbf{A}, \mathbf{u}) \text{ et } (\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}),$$

où \mathbf{A} est une matrice uniforme dans $\mathbb{Z}_q^{m \times n}$, \mathbf{u} est un vecteur uniforme de \mathbb{Z}_q^m , \mathbf{s} est un vecteur uniforme de \mathbb{Z}_q^n et \mathbf{e} est un vecteur gaussien de paramètre $\alpha \cdot q$ dans \mathbb{Z}_q^m .

⁵Nous présentons ici le problème LWE avec secret uniforme et erreur Gaussienne, d'autres variantes sont considérées dans la littérature, notamment avec des restrictions sur \mathbf{s} [Mic18] et d'autres distributions sur \mathbf{e} .

Le problème LWE peut être résumé au fait de distinguer entre un vecteur uniforme et un vecteur proche d'un réseau, comme représenté en Figure I.4. LWE est également présent dans la littérature sous une version « Recherche », où seul le couple $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ est donné, et il est demandé de retrouver \mathbf{s} .




FIGURE I.4 : Les deux distributions de LWE.

Comme pour SIS, on peut relier la difficulté de LWE au problème pire-cas SIVP. Si $q > 2\sqrt{n}/\alpha$ et $m = \text{poly}(n)$, alors SIVP_γ se réduit quantiquement⁶ [Reg05] à $\text{LWE}_{n,m,q,\alpha}$ pour $\gamma = \tilde{O}(n/\alpha)$.

Nous résumons dans la Figure I.5 les relations de difficultés des problèmes présentés précédemment. La partie pire-cas du diagramme est extraite de la revue de littérature de Noah Stephens-Davidowitz [Ste15, Page 1] et de [Mic08]. Une flèche du problème A au problème B indique que A se réduit à B (donc « A est au mieux aussi difficile que B »). Une flèche en tirets désigne une réduction utilisant un ordinateur quantique.




FIGURE I.5 : Difficulté relative de quelques problèmes de réseaux.

⁶La réduction nécessite un ordinateur quantique.

I.1.4 Les réseaux structurés

Polynômes et réseaux euclidiens

Dans les problèmes décrits précédemment, les réseaux sont représentés par leurs bases, sous la forme d'une matrice. Pour un réseau en dimension n , cela fait alors n^2 entiers à manipuler pour faire nos opérations cryptographiques (chiffrement, déchiffrement, signatures...), ce qui rend le temps d'exécution des protocoles important quand n devient grand. Une méthode trouvée pour diminuer ces temps d'exécution est d'utiliser des matrices présentant une structure.

Définissons une version du problème SIS sur l'anneau polynomial $\mathbb{Z}[X]/(X^n + 1)$.

Definition I.1.4 (Ring-SIS $_{q,n,m,\beta}$ [Mic02, LM06, PR06]). *Soit $q \geq 2$, $n \geq 1$, $m \geq n \log(q)$, et $\beta \geq \sqrt{n \log(q)}$. Le problème Ring-SIS $_{q,n,m,\beta}$ est défini comme suit. Étant donné des polynômes uniformes $P_1, \dots, P_k \in \mathbb{Z}_q[X]/(X^n + 1)$, trouver des polynômes $Q_1, \dots, Q_k \in \mathbb{Z}[X]/(X^n + 1)$ non tous nuls avec⁷ $\|(Q_i)_i\| \leq \beta$ tel que*

$$P_1(X) \cdot Q_1(X) + \dots + P_k(X) \cdot Q_k(X) = 0 \pmod{(X^n + 1, q)}. \quad (\text{I.1})$$

Si on choisit de représenter les polynômes comme des vecteurs, l'équation (I.1) peut se réécrire

$$[\text{nrot}(P_1) | \dots | \text{nrot}(P_k)] \cdot [Q_1, \dots, Q_k]^T = \mathbf{0} \pmod{q},$$

où $\text{nrot}(P)$ désigne la matrice néga-circulante associée au polynôme $P(X) = p_0 + p_1 \cdot X + \dots + p_{n-1} \cdot X^{n-1}$:

$$\text{nrot}(P) := \begin{pmatrix} p_0 & -p_{n-1} & -p_{n-2} & \cdots & -p_1 \\ p_1 & p_0 & -p_{n-1} & \cdots & -p_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ p_{n-2} & p_{n-3} & p_{n-4} & \cdots & -p_{n-1} \\ p_{n-1} & p_{n-2} & p_{n-3} & \cdots & p_0 \end{pmatrix}.$$

Avec cette précision, on voit que le problème Ring-SIS est exactement le problème SIS restreint aux matrices de la forme $[\text{nrot}(P_1) | \dots | \text{nrot}(P_k)]$ avec les P_i uniformes dans $\mathbb{Z}[X]/(q, X^n + 1)$. L'arithmétique dans l'anneau $\mathbb{Z}_q[X]/(X^n + 1)$ étant rendue plus rapide par l'utilisation de la transformée de Fourier discrète (qui permet une multiplication de deux polynômes en temps quasi linéaire), les systèmes cryptographiques construits à partir de Ring-SIS seront plus efficaces que ceux construits à partir de SIS. Il est également possible de définir le problème Ring-LWE [SSTX09, LPR10, PRS17] de manière similaire.

Cela étant dit, la plus grande efficacité des protocoles utilisant des réseaux structurés vient avec des hypothèses de sécurités plus fortes, qu'il faut étudier spécifiquement.

Théorie algébrique des nombres

Une manière naturelle de définir des hypothèses de sécurités plus génériques serait d'étendre nos problèmes calculatoires structurés aux anneaux de polynômes génériques de la forme $\mathbb{Z}[X]/P(X)$ pour P un polynôme entier quelconque. Cette approche peut cependant mener, selon le choix du polynôme P , à des problèmes de faible difficulté (par exemple, sur l'anneau $\mathbb{Z}[X]/(X^n - 1)$ le problème Ring-SIS est résoluble en temps raisonnable avec bonne probabilité sous certaines conditions sur n [PR06]). Pour construire des anneaux de polynômes avec de bonnes propriétés,

⁷Il y a plusieurs façons de définir les normes sur $\mathbb{Z}[X]/(X^n + 1)$, pour une présentation de la norme que nous utilisons dans ce manuscrit, se référer à la Section II.2.1. Si $n = 2^k$, cette norme est identique - à un facteur multiplicatif près - à la norme euclidienne du vecteur des coefficients du polynôme.




FIGURE I.6 : Visualisation des matrices.

Non structurées ($n = 9$) ; Module de rang 1 ($d = 9, k = 1$) ; Module de rang 3 ($d = 3, k = 3$).

nous allons nous reposer sur la théorie des nombres. Si $P \in \mathbb{Z}[X]$ est un polynôme irréductible de degré d , l'anneau $K = \mathbb{Q}[X]/P(X)$ est un corps, dit « corps de nombres de degré d » et on peut définir son anneau d'entiers \mathcal{O}_K comme l'ensemble de ses éléments vérifiant une équation polynomiale unitaire entière. Cet anneau d'entiers est toujours un sous-anneau d'un anneau de polynômes et possède un certain nombre de propriétés, notamment d'être de Dedekind et d'être le seul ordre maximal de K . Nous ne rentrerons pas dans les détails sur ce que ces propriétés signifient, pour nous, elles sont une garantie que \mathcal{O}_K n'a pas de propriété algébrique « trop différente de \mathbb{Z} ». En particulier, si K est le corps cyclotomique de degré $d = 2^n$, l'anneau \mathcal{O}_K est alors exactement l'anneau de polynôme $\mathbb{Z}[X]/(X^{2^n} + 1)$. Dans toute la suite du manuscrit, la lectrice plus habituée aux anneaux polynomiaux qu'à la théorie des nombres pourra remplacer toutes les occurrences de \mathcal{O}_K par $\mathbb{Z}[X]/(X^d + 1)$ pour $d = 2^n$.

Dans tout le reste de cette introduction, K est un corps de nombres de degré $d \geq 2$ et d'anneau d'entiers \mathcal{O}_K . Un réseau module sur K de rang $k \geq 1$ est sous-ensemble de K^k , stable par addition, soustraction et multiplication par un élément de \mathcal{O}_K ⁸. Un réseau module $M \subset K^k$ est associé à un réseau euclidien $\Phi(M) \subset \mathbb{R}^{d \cdot k}$ via $\Phi(\cdot)$, le plongement canonique du corps⁹, on peut donc restreindre tous les problèmes de réseaux euclidiens présentés précédemment aux réseaux modules, avec l'avantage que la structure de K permet d'effectuer les opérations sur les réseaux modules de manière plus efficace que sur des réseaux non-modules de même dimension.

Le corps le plus utilisé en pratique est le corps cyclotomique de conducteur une puissance de 2, c'est-à-dire le corps $K = \mathbb{Q}(\zeta_{2^n})$. Avec ce choix de corps, les réseaux modules libres¹⁰ de rang k sont exactement¹¹ les réseaux euclidiens générés par les matrices par blocs de la forme

$$\mathbf{B} = \begin{pmatrix} \text{nrot}(P_{1,1}) & \dots & \text{nrot}(P_{1,k}) \\ \vdots & & \vdots \\ \text{nrot}(P_{k,1}) & \dots & \text{nrot}(P_{k,k}) \end{pmatrix}$$

pour $(P_{i,j})_{1 \leq i,j \leq k} \in \mathbb{Q}[X]/(X^{2^{n-1}} + 1)^{k \times k}$ (voir Figure I.6). La structure de ce corps a notamment été utilisée pour les algorithmes de signature (CRYSTALS-Dilithium [LDK⁺20]) et d'échange de clés (CRYSTALS-Kyber [ABD⁺19]) sélectionnés par la compétition du NIST [NIST].

Certaines particularités des corps cyclotomiques (en particulier son grand nombre de sous-corps) ont mené certain-es auteur-ices à proposer d'utiliser des corps différents pour les protocoles

⁸Dans ce manuscrit, on ne parlera que de \mathcal{O}_K -modules sans torsion

⁹Ou le plongement par coefficients.

¹⁰Tous les modules ne sont pas libre, mais nous nous restreignons aux modules libres pour cette introduction.

¹¹À rotation et homothétie près.

cryptographiques, de manière à « diminuer la surface d'attaque » [BCLV17] (le corps proposé dans [BCLV17] est le corps défini par le polynôme $X^p - X - 1$, pour p un nombre premier). Nous nous sommes efforcés d'être agnostique sur le choix du corps dans ce manuscrit. Lorsque des résultats plus précis sont disponibles sur les cyclotomiques, nous le mentionnerons.

On peut alors, comme pour les réseaux euclidiens non structurés, définir le problème SVP_γ restreint aux réseaux modules. Si K est un corps de nombres, $\gamma \geq 1$ un facteur d'approximation et $k \geq 1$ un rang, le problème $\text{modSVP}_{k,\gamma}^K$ est le problème SVP_γ restreint aux réseaux modules de rang k sur le corps K . Comme dit précédemment, $\text{modSVP}_{k,\gamma}^K$ est un sous-problème de SVP_γ en dimension $d \cdot k$.

Il faut noter que contrairement au cas des réseaux non structurés, il y a ici deux variables à ajuster pour augmenter la difficulté du problème : le rang k du réseau et le degré d du corps. Contrairement aux réseaux non structurés, ici, c'est le degré du corps que nous ferons augmenter pour accroître la difficulté des problèmes. En particulier, nous nous intéressons aux réseaux de rang 1 et 2 dans des corps de plus en plus grands. Les variables desquelles dépendront le temps d'exécution de nos algorithmes et nos facteurs d'approximation seront le degré d du corps, qui correspondra à la dimension des réseaux, et le discriminant-racine $\Delta_K^{1/(2d)}$ du corps, qui - grossièrement - représentera la taille du corps de nombres.

Le problème id-HSVP

Le cas le plus élémentaire de réseau module est le cas du rang 1, qui correspond à trouver un vecteur court dans un idéal fractionnaire de K . Ce problème est donc nommé id-HSVP_γ ¹² La difficulté de ce problème n'est pas encore précisément comprise. Il semble que pour certains paramètres γ grands, id-HSVP_γ soit plus facile que SVP_γ : il existe des algorithmes pour le résoudre en temps polynomial quantique pour des facteurs d'approximation $\gamma \geq 2^{\tilde{O}(\sqrt{d})}$ dans le cas des corps cyclotomiques [CDPR16, CDW17] ou sur n'importe quel corps K si des pré-calculs ont été effectués [PHS19]. Pour $\gamma = \text{poly}(d)$ les meilleurs algorithmes connus actuellement sont toutefois ceux utilisés sur les réseaux non structurés (voir la Figure I.7).

Les idéaux sont l'exemple le plus simple de réseaux modules, ce qui met en avant l'importance de la compréhension de la difficulté de id-HSVP . La structure algébrique des réseaux idéaux a également permis de développer des réductions pire-cas vers cas-moyen les concernant.

Relation entre id-HSVP et les autres problèmes de réseaux structurés. Les problèmes moyens-cas Ring-SIS et Ring-LWE ont été reliés à la recherche de petits vecteurs dans des réseaux idéaux. Il a été prouvé que id-HSVP_γ se réduisait au problème Ring-LWE [SSTX09, LPR10] et au problème Ring-SIS [PR06, LM06] pour des facteurs d'approximation polynomiaux. Ces résultats doivent être vus comme des bornes inférieures de difficulté sur Ring-SIS et Ring-LWE : ces problèmes n'étant pas définis sur des réseaux idéaux directement, les attaques sur id-HSVP ne s'étendent pas à Ring-SIS ou Ring-LWE.

Auto-réduction pire-cas vers cas-moyen pour id-HSVP. Dans [Gen09], Gentry prouve que (avec un oracle de factorisation), id-HSVP dans le pire-cas se réduit à résoudre id-HSVP avec bonne probabilité sur l'inverse d'un petit idéal premier tiré au hasard.

¹²Le H - pour "Hermite" Short Vector Problem - vient du fait que pour les réseaux idéaux, il est plus naturel de chercher des vecteurs petits par rapport à la norme algébrique de l'idéal et non par rapport au λ_1 . On peut montrer que ces deux approches sont équivalentes, voir le Lemme II.4.9.




FIGURE I.7 : Compromis temps/approximation pour id-HSVP pour $\deg(K) = d$ avec pré-calcul exponentiel en d [PHS19, Fig. 2].

Un résultat similaire a été prouvé en 2022 par de Boer et al. [BDPW20], où id-HSVP dans le pire-cas est réduit à id-HSVP pour l'arrondi gaussien d'un idéal uniforme dans l'ensemble des idéaux replets¹³ de norme 1.

Le problème NTRU

Le problème NTRU¹⁴ introduit en 1998 [HPS98] (dans une version légèrement différente de celle présentée ici) est un autre problème d'équation polynomiale modulaire, où, étant donné un polynôme $h \in \mathbb{Z}[X]/(X^n + 1)$, il faut trouver une écriture $h = g/f \bmod q$ avec f et g de petits polynômes (avec la promesse que de tels polynômes existent) :

Definition I.1.5 $((\gamma, \gamma', q)$ -NTRU). Soit $q \geq 2$, $\gamma \geq \gamma' > 0$. Une instance (γ, q) -NTRU est un polynôme $h \in \mathbb{Z}[X]/(X^n + 1)$ tel que $h = g/f \bmod q$ avec $\|f\|, \|g\| \leq \sqrt{q}/\gamma$. Le problème (γ, γ', q) -NTRU demande, étant donné une instance (γ, q) -NTRU h , de trouver \tilde{f}, \tilde{g} de norme $\leq \sqrt{q}/\gamma'$ vérifiant $h = \tilde{g}/\tilde{f} \bmod q$.

Le problème NTRU peut être généralisé en prenant $f, g, h \in \mathcal{O}_K$ pour un corps de nombres K . On a remarqué très tôt que le problème NTRU peut être interprété en termes de réseaux euclidiens [HPS98, CS97]. En effet, l'ensemble

$$L_h := \begin{bmatrix} 1 \\ h \end{bmatrix} \mathcal{O}_K + \begin{bmatrix} 0 \\ q \end{bmatrix} \mathcal{O}_K = \left\{ (\tilde{f}, \tilde{g})^T \in \mathcal{O}_K^2, \quad h \cdot \tilde{f} = \tilde{g} \bmod q \right\}$$

est un réseau module de rang 2. Ce module est défini par h , à partir duquel une base peut être calculée, et possède une particularité : il contient un vecteur non nul inhabituellement court (f, g) . En effet, pour la plupart des h , on a $\det L_h = \Delta_K \cdot q^d$, (où Δ_K désigne le discriminant du corps).

¹³ « Replete ideal » en anglais, la dénomination française n'est pas fixe.

¹⁴La signification du nom semble avoir été perdue.




FIGURE I.8 : À gauche un réseau quelconque, à droite un réseau possédant un vecteur particulièrement court.

En conséquence, on s'attendrait¹⁵ à ce que les vecteurs non nuls les plus courts aient une norme autour de $q^{1/2}$, à quelques facteurs près en fonction de Δ_K et d . Cependant, $(f, g)^T$ est par hypothèse beaucoup plus court. Nous avons donc un module de rang 2 sur \mathcal{O}_K avec la promesse qu'il contient un vecteur non nul inhabituellement court, c'est-à-dire un sous-module de rang 1 inhabituellement dense. Nous appelons mod-uSVP_2 le problème qui consiste à trouver un vecteur non nul court dans un module de rang 2 contenant un vecteur court (voir Figure I.8).

Difficulté de NTRU. Les problèmes NTRU et mod-uSVP_2 existent en fait sous deux formes. La plus naturelle, décrite ci-dessus, demande de récupérer un vecteur court du module de rang 2 correspondant. C'est la variante que nous considérons implicitement dans cette introduction lorsque nous discutons de NTRU et de mod-uSVP_2 . D'autres versions existent et sont considérées dans ce manuscrit, elles demandent de trouver une base du sous-module le plus dense (donc généré par le vecteur inhabituellement court) plutôt que de trouver le vecteur court directement et seront désignées avec un exposant mod : NTRU^{mod} et $\text{mod-uSVP}_2^{\text{mod}}$. Les deux versions de ce problème sont équivalentes si un oracle à id-HSVP est donné. Comme on l'a vu plus haut, le problème NTRU peut être considéré comme un cas particulier du problème de réseaux modules mod-uSVP_2 , cependant, il n'est pas clair si ses instances sont représentatives de toutes les instances de mod-uSVP_2 . Dans [Pei16, Section 4.4.4], Peikert esquisse une réduction d'une version décisionnelle du problème NTRU au problème Ring-LWE [SSTX09, LPR10]; cette réduction peut être adaptée au problème NTRU de recherche que nous considérons ici. Il convient de noter que sous certaines contraintes de paramètres, le problème Ring-LWE est équivalent à mod-SIVP_2 [LS15, AD17], qui est la restriction aux réseaux modules de rang 2 du problème SIVP décrit précédemment.

Liens entre NTRU et id-HSVP

Dans l'autre direction, Pellet-Mary et Stehlé [PS21] ont présenté une réduction du problème du vecteur le plus court pour les réseaux correspondant à des idéaux de \mathcal{O}_K (id-HSVP) vers NTRU.

¹⁵C'est une conséquence de l'heuristique Gaussienne (voir par exemple [GNR10]), qui décrit à quoi un réseau « typique » ressemble.

Dans l'ensemble, on voit que NTRU se situe entre id-SVP et mod-SIVP. Comme noté précédemment, id-HSVP admet des algorithmes plus performants que les algorithmes de réduction de réseaux génériques [LLL82, Sch87] pour certaines plages de paramètres [CDW21, PHS19]. Comme un tel phénomène est inconnu dans le cas du mod-SIVP, il pourrait y avoir un grand saut de difficulté entre id-HSVP et mod-SIVP. Il n'est pas clair actuellement lequel de ces problèmes capture la véritable difficulté de NTRU, ou si NTRU se situe quelque part entre les deux.




FIGURE I.9 : Difficulté relative de quelques problèmes de réseaux modules (facteurs d'approximation omis).

Une présentation partielle de la difficulté de certains problèmes de réseaux structurés est disponible en Figure I.9. Tous les problèmes décrits le sont pour un même corps K . Pour des raisons de lisibilité, nous ne représentons ni les pertes de facteurs d'approximation, ni les conditions d'applications. Comme précédemment, une flèche de A vers B indique que A se réduit à B . Les flèches en tiret désignent les réductions quantiques, les flèches sans citation sont les réductions triviales ou folklores.

I.1.5 Contribution de cette thèse

Publications

Ce manuscrit est basé sur les deux publications réalisées pendant mon doctorat :

- [FPS22] **On Module Unique-SVP and NTRU**. Joël Felderhoff, Alice Pellet-Mary and Damien Stehlé. ASIACRYPT 2022.
- [FPSW23] **Ideal-SVP is Hard for Small-Norm Uniform Prime Ideals**. Joël Felderhoff, Alice Pellet-Mary, Damien Stehlé and Benjamin Wesolowski. TCC 2023.

Difficulté de id-HSVP pour des idéaux entiers aléatoires

Les deux distributions moyen-cas d'idéaux pour lesquelles une réduction pire-cas vers moyen-cas existe dans la littérature ne sont pas satisfaisantes. La première [Gen09] est une distribution d'idéaux inverses ce qui, en plus d'être peu naturel d'un point de vue algébrique, ne permet pas d'utiliser la réduction de [PS21] (valide uniquement pour les idéaux entiers) pour décrire une distribution NTRU dont la difficulté est basée sur id-HSVP dans le pire-cas. La seconde [BDPW20] est bien une distribution d'idéaux entiers, mais le processus d'arrondi au cœur de la réduction ne permet d'espérer que des idéaux de très grandes normes algébriques (de l'ordre de $2^{\mathcal{O}(d^2)}$). Utiliser des idéaux de cette taille pour faire de la cryptographie mènerait à des coûts de calculs tels que l'utilisation des réseaux structurés perdrait tout intérêt.

Nous proposons dans le Chapitre IV (tiré de [FPSW23]) une nouvelle réduction moyen-cas vers moyen-cas (avec oracle de factorisation) de id-HSVP. Soit \mathcal{W} un ensemble quelconque d'idéaux et \mathcal{W}^{-1} l'ensemble des inverses des éléments de \mathcal{W} . Nous prouvons dans le Théorème IV.5.1 que résoudre id-HSVP sur un idéal uniforme de \mathcal{W}^{-1} se réduit à résoudre id-HSVP sur un petit idéal entier uniforme et sur un idéal uniforme de \mathcal{W} . En particulier, appliquer la réduction de [PS21] à cette distribution donne des instances de NTRU avec module exponentiel.

En spécialisant ce résultat avec \mathcal{W} l'ensemble des idéaux premiers de petites normes (de l'ordre de $\text{poly}(d, \Delta_K^{1/d})^d$), nous démontrons dans le Corollaire IV.5.4 que la difficulté de id-HSVP sur la distribution décrite dans [Gen09] n'est pas différente de celle de id-HSVP sur la distribution uniforme sur les petits idéaux premiers, et donc que id-HSVP dans le pire-cas se réduit à id-HSVP sur un petit idéal premier uniforme.

Cette nouvelle distribution nous permet, grâce à la réduction de [PS21], de définir une distribution sur les instances NTRU de module polynomial basée sur la difficulté de id-HSVP dans le pire-cas, c'est le Corollaire IV.6.2. Cette distribution nécessite cependant un oracle de factorisation pour pouvoir être utilisée dans un contexte cryptographique.

Liens entre NTRU et mod-uSVP₂ et distribution cas-moyen pour mod-uSVP₂

Comme indiqué précédemment, le problème NTRU est un cas particulier du problème de modules mod-uSVP₂, où le module est de la forme $(1, h)^T \cdot \mathcal{O}_K + (0, q)^T \cdot \mathcal{O}_K$. Dans le Chapitre V, nous prouvons que ces modules NTRU sont en fait représentatifs des réseaux mod-uSVP₂ généraux, au sens où tout module mod-uSVP₂ peut être transformé en temps polynomial en un réseau NTRU de géométrie similaire, c'est le Théorème V.4.1. Nous accompagnons ce résultat avec des réductions pour les variantes pire-cas et cas-moyen de NTRU^{mod} et de NTRU^{vec}.

Nous proposons ensuite dans le Théorème V.6.1 une réduction pire-cas vers cas-moyen pour les réseaux mod-uSVP₂, qui tourne en temps polynomial avec appel à un oracle de factorisation (donc en temps polynomial quantique). Ce résultat permet de donner une autre distribution moyen-cas NTRU dont la difficulté est basée sur un problème pire-cas en combinant les Théorèmes V.4.1 et V.6.1. Cette distribution nécessite cependant un oracle de factorisation pour pouvoir être utilisée et semble assez artificielle.

Apports de ce manuscrit par rapport à [FPS22, FPSW23]

Comme mentionné précédemment, ce manuscrit est en grande partie issu des deux articles publiés pendant ma thèse : [FPS22, FPSW23]. Nous avons néanmoins rajouté un résultat additionnel au Chapitre III qui nous permet de préciser certains théorèmes de [FPS22] et [FPSW23].

Soit K un corps de nombres de degré ≥ 3 . Suivant un résultat classique de théorie analytique des nombres, lorsque X tend vers l'infini, le nombre d'idéaux de norme algébrique inférieure à X , noté $N_K(X)$, est équivalent à $\rho_K \cdot X$, où ρ_K est le résidu en $s = 1$ de la fonction zêta de

Dedekind associé à K . Dans certain de nos résultats, nous avons eu besoin de borner le terme d'erreur dans cette approximation, mais les bornes de la littérature ne faisaient pas apparaître explicitement la dépendance en K . Nous proposons au Chapitre III une explicitation, nouvelle à notre connaissance, de la borne sur cette erreur où la dépendance dans le corps est précisée, c'est le Théorème III.1.2.

Dans [FPSW23], le temps d'exécution de la réduction de id-HSVP dans le pire-cas vers id-HSVP pour un idéal premier uniforme [FPSW23, Corollaire 5.4] dépend d'un paramètre ad hoc $\tilde{\rho}_A$, qui est la proportion inverse d'idéaux premiers du corps de norme inférieure à A . Il est connu que $\tilde{\rho}_A$ se comporte asymptotiquement comme $\rho_K \cdot \ln(A)$, mais les bornes d'erreurs sur cette approximation dans la littérature ne nous permettaient pas d'appliquer ce résultat pour les valeurs de A qui nous intéressent. Le Corollaire III.1.3 nous permet de donner une réduction dont le temps d'exécution ne dépend que de ρ_K , et non plus de $\tilde{\rho}_A$. Cela mène à une taille d'idéaux et des facteurs d'approximation plus importants, qui sont présentés dans le Corollaire IV.5.5. Nous utilisons également notre résultat concernant $N_K(\cdot)$ dans le Chapitre V. La réduction de mod-uSVP₂ à NTRU (Théorème V.4.1, ou [FPS22, Théorème 4.1]) avait un temps d'exécution dépendant de $\zeta_K(2)$ qui, en fonction du corps de nombres considéré, peut être exponentiel en le degré. Nous utilisons les résultats du Chapitre III pour proposer une nouvelle version de ce théorème où le temps d'exécution ne dépend plus de $\zeta_K(2)$ (au prix d'une perte de facteur d'approximation plus grande). Le théorème en question est le Théorème V.4.2, dont la preuve est en Section D.5.

Une deuxième modification a été apportée au Théorème V.4.1 par rapport à [FPS22, Théorème 4.1]. Dans [FPS22, Théorème 4.1], une condition sur le corps de nombres K est présente : on demande que $\zeta_K(2) = 2^{o(d)}$. Cette condition est non triviale, car elle a des implications sur l'arithmétique de K (on peut l'interpréter comme le fait que « \mathcal{O}_K n'ait pas trop de petits idéaux »). Cette condition était issue de la preuve du Lemme 4.3 de [FPS22], et était purement technique. Nous avons modifié la preuve de ce lemme (dans ce manuscrit, le Lemme V.4.4) pour la retirer.

I.2 Introduction (English)

“So... Are you a mathematician
or a computer scientist?”

J.M. Felderhoff (my father), 2022

The notion of “secure communication” covers a wide range of problematics, such as message authentication (guaranteeing the origin of a received message) or message encryption (making the content of a message unintelligible to anyone other than the recipient), with proto-examples dating back to Antiquity. Nowadays, modern computers render naïve security techniques (such as mono-alphabetic encryption, where one letter is simply replaced by another) obsolete.

The democratization of computers and the Internet has been accompanied by the introduction of numerous encryption and digital signature protocols, such as TLS (the standard used in HTTPS), OpenPGP (used to sign emails) and end-to-end encryption protocols (used in WhatsApp, Telegram and Signal). In this context of widespread use, we need to find ways of guaranteeing the security of cryptographic protocols.

In this manuscript, we are working within the framework of public-key cryptography, which enables secure communication when both parties are unable to share a key beforehand (which is regularly the case in Internet exchanges). In this paradigm, a distinction is made between a secret key (available to only one party) and a public key (available to everyone). The security of public-key protocols is based on the difficulty of guessing the secret key with knowledge of public data, such as public key and messages transiting the network. Nowadays, the security of public-key cryptographic protocols is guaranteed by means of security proofs.

I.2.1 Guaranteeing the security of a protocol

Let us take the example of an encryption protocol that enables two parties to exchange messages in a way that is unintelligible to a third party. Proving the security of such a protocol involves three steps.

The first is to define the adversary, i.e. the abstract entity against which we want to guarantee the security of our system, such as an intelligence agency or a company (legal or not) wishing to resell data... The questions to ask are typically:

- What would be the adversary’s goal in breaking my protocol? For our example, it could be to decrypt a message, or to distinguish an encrypted message from a random bit string.
- How much computing power does the adversary have access to? How long do we want to resist them?
- Can they interact with the cryptographic system? For example, can they send false messages and observe their interlocutor’s behavior? Does they have access to part of the secret key?

The second step is to define one or more security assumptions. These are mathematical statements of the form: “it is impossible to solve such and such a problem in a reasonable time”. A classic example is factorization: “Given a large number N , it is impossible in a reasonable time to find p and q different from 1 such that $N = p \cdot q$ ”.¹⁶

¹⁶Of course, the notion of “large” and “reasonable time” need to be defined more precisely, see Section II.4.

The third step is to make the connection between the two. In practice, this means proving a mathematical statement of the form : “Suppose there is an adversary that breaks our protocol, then there must also be an algorithm that breaks the security assumption”. This kind of statement is called a security reduction. It should be interpreted as the mathematical formalization of the fact that, as long as the security assumption holds, the protocol is secure against this type of adversary. We say that studying the security of the protocol against this adversary *reduces to* studying the validity of the security assumption.

I.2.2 Quantum adversary and post-quantum cryptography

In this manuscript, we assume that our adversary has access to a quantum computer (as opposed to our classical transistor-based computers) and that we do not. It is beyond the scope of this introduction to precisely explain how a possible quantum computer would work. Suffice it to say that they perform calculations differently from a classical computer, in the same way that a mechanical calculator performs calculations differently from a digital calculator. This difference in particular implies that certain problems that were thought to be difficult for classical computers can be solved efficiently with quantum computers.

In particular, Shor’s quantum algorithm [Sho94] allows the factorization and discrete logarithm problems to be solved with a reasonable amount of resources. Security protocols based on the difficulty of these problems are therefore obsolete if the adversary has access to a quantum computer. Given that the most widely deployed protocols (e.g., TLS and OpenPGP) depend heavily on the difficulty of solving these problems, the (widely debated...) possibility of a quantum computer appearing in the next few years has prompted standardization institutes, national authorities and industry to intensify research into quantum-proof security assumptions (we talk about *post-quantum* security assumptions).

In particular, the NIST (the U.S. Institute of Standards and Technology) post-quantum cryptography standardization process, in the form of a competition, began in 2016 and has been completed in 2022 [NIST]. As a result of this competition, four main families of encryption protocols appear to be resistant to quantum computers. These are protocols based on polynomial systems, error-correcting codes, isogenies between elliptic curves and lattices. The work carried out in this thesis deals with security assumptions related to *lattices*.

I.2.3 Lattices

The Shortest Vector Problem

Informally, a lattice is an infinite set of regularly distributed points in space. It can be used in 2 or 3 dimensions, for example, to represent the distribution of atoms in a crystal (see Figure I.10). In cryptography, high-dimensional lattices ($n \approx 500$) are used for their algorithmic properties.

Let $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{n \times n}$ be an (invertible) matrix, which we call a basis. The lattice spanned by this matrix, denoted by $\mathcal{L}(\mathbf{B})$, is the set of integer combinations of the column vectors of $(\mathbf{b}_i)_{1 \leq i \leq n}$. Mathematically, we write

$$\mathcal{L}(\mathbf{B}) = \mathbf{B} \cdot \mathbb{Z}^n = \left\{ \sum_{i=1}^n \mathbf{b}_i \cdot x_i, (x_i)_{1 \leq i \leq n} \in \mathbb{Z}^n \right\},$$

An example of lattice L is given in Fig. I.11. This particular lattice is spanned by the blue vectors in the figure, i.e., the matrix

$$\begin{pmatrix} 1.1 & -0.1 \\ -0.1 & 1 \end{pmatrix}.$$




Figure I.10: Graphene layer at atomic level [Ale09] and corresponding lattice.

Note that not only blue vectors, but also red vectors generate L : a lattice has multiple bases.




Figure I.11: An example of lattice.

Numerous computational problems are associated with lattices, some of which can be used to construct cryptography (a presentation of the main ones and their relationships can be found in a literature review by Peikert [Pei16]). In this manuscript, we study the problem of finding one or many shortest vectors in a lattice, given a basis of it. We define the Shortest Vector Problem as follows:

Definition I.2.1 (SVP_γ). Let $\gamma \geq 1$. The problem SVP_γ asks, given $\mathbf{B} \in \mathbb{Z}^{n \times n}$, to find $\mathbf{v} \in \mathcal{L}(\mathbf{B}) \setminus \{\mathbf{0}\}$ satisfying

$$\|\mathbf{v}\| \leq \gamma \cdot \lambda_1(\mathcal{L}(\mathbf{B})),$$

where $\lambda_1(\mathcal{L}(\mathbf{B}))$ is the Euclidean norm of a shortest non-zero vector of the lattice $\mathcal{L}(\mathbf{B})$.

In particular, SVP_1 (also called “exact-SVP”) asks to find one of the shortest vectors of $\mathcal{L}(\mathbf{B})$. In the case of the lattice shown in Figure I.11, the green vector would be an answer to SVP_1 . This problem is difficult to solve. In fact, the best known algorithms (whether classical or

quantum) solving SVP_1 in dimension n either require $n^{O(n)}$ operations (enumeration algorithms, see for example [FP85, Kan87, HS07]) or $2^{O(n)}$ operations, but at the cost of a memory amount of $2^{O(n)}$ bits (sieving algorithms [AKS01] or based on Voronoï cells [MV13]). This state-of-the-art makes SVP_1 intractable with reasonable resources (say, less than 10^{10} years using all the world's computers) as soon as the dimension n becomes larger than a few hundred (the largest record recorded on [Nam] at the time of writing this manuscript is $n = 190$).

Comparing algorithmic problems The central notion for comparing the difficulty of two algorithmic problems is the notion of reduction. In complexity theory, we say that a problem \mathcal{B} is harder or as hard as a problem \mathcal{A} if, given an efficient algorithm solving problem \mathcal{B} , we can write an efficient algorithm¹⁷ (called reduction) to solve problem \mathcal{A} . We'll say that problem \mathcal{A} reduces to problem \mathcal{B} . Two problems \mathcal{A} and \mathcal{B} are said to be equivalent if \mathcal{A} reduces to \mathcal{B} and \mathcal{B} reduces to \mathcal{A} .

The larger γ is, the easier SVP_γ becomes, in the sense that a solution of SVP_γ is also a solution of $\text{SVP}_{\gamma'}$ if $\gamma' \geq \gamma$. It has been shown that SVP_1 is NP-hard¹⁸, which implies that it is plausible that an efficient algorithm to solve every instance of it does not exist, even with a quantum calculator¹⁹. We can go even further in describing its difficulty as a function of γ : for a lattice L of dimension n and an integer β between 1 and n , there exists an algorithm (the BKZ algorithm [SE94]) solving SVP_γ for $\gamma = 2^{\tilde{O}(n/\beta)}$ on input L in time proportional to 2^β (the result is deliberately simplified in this introduction, for conditions on β and a precise value of γ , see Lemma II.1.14). This algorithm enables us to give a difficulty gradient for SVP_γ as a function of γ , shown in Figure I.12.

That being said, it has not been proved that SVP_γ is NP-hard for the approximation factors γ used in cryptographic constructions. In fact, it is unlikely, as it has been proved [AR05] that for $\gamma = \sqrt{n}$, SVP_γ is in the complexity class $\text{NP} \cap \text{coNP}$. The construction of cryptography based on the difficulty of an NP-hard problem is still an open problem.

Another central problem in lattice-based cryptography is SIVP_γ , which asks, roughly speaking, to find a full-rank family of small (their size is controlled by γ as for SVP_γ) vectors in a given lattice. It has been proved that SIVP_1 and SVP_1 are computationally equivalent [GMSS99, Mic08] and that $\text{SIVP}_{\sqrt{n}\cdot\gamma}$ reduces to SVP_γ [Ste15].

Average-case problems

A public-key cryptographic protocol based on lattices typically works as follows: we sample a lattice L at random (according to a certain distribution) with small vectors inside of it ; we then publish a basis of the lattice as our public key, and use the small vectors as our secret key. The security of the system then relies on the fact that an adversary cannot find small non-zero vectors of the lattice given its basis, i.e., on the difficulty of SVP_γ for the lattice that has been drawn. Note that in this case, the security of the protocol is not exactly based on the difficulty of SVP_γ , where you have to be able to find a short vector in *all lattices*, but on the difficulty of SVP_γ on a *randomly chosen lattice*. Solving SVP_γ on a random lattice is a so-called *average-case* problem, which is easier than solving SVP_γ on all lattices, a so-called *worst-case* problem.

One of the interests of lattice-based cryptography is that the difficulty of SVP_γ on some of these distributions can be related to the difficulty of SVP_γ in the worst case. We describe two of these in this introduction.

¹⁷Executing in time polynomial in the size of its input.

¹⁸For randomized reductions [Ajt98], or for the ℓ_∞ norm [Emd].

¹⁹NP-diffculty and its relation to quantum computation is a broad topic that is outside the scope of this manuscript. The interested reader is redirected to [AB09].




Figure I.12: Time/approximation trade-off for SVP with BKZ in dimension n . Approximation factors used in cryptography are shown in green.

Short Integer Solution (SIS). The first example of an average-case problem we give is the Short Integer Solutions problem, introduced in 1996 by Miklós Ajtai [Ajt96]. It consists in solving a linear system with a condition on the size of the solution. In what follows, \mathbb{Z}_q denotes the ring $\mathbb{Z}/q\mathbb{Z}$.

Definition I.2.2 ($\text{SIS}_{q,n,m,\beta}$). Let $q \geq 2$, $n \geq 1$, $m \geq n \log(q)$, and $\beta \geq \sqrt{n \log(q)}$. The problem $\text{SIS}_{q,n,m,\beta}$ is defined as follows. Given \mathbf{A} a uniform matrix in $\mathbb{Z}_q^{n \times m}$, find $\mathbf{x} \in \mathbb{Z}^m \setminus \{\mathbf{0}\}$ such that $\|\mathbf{x}\| \leq \beta$ satisfying

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}.$$

Note that the input to this problem is a matrix *sampled uniformly* from $\mathbb{Z}_q^{n \times m}$: this is an average-case problem. We can relate SIS to a lattice problem by observing that for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, the set of possible solutions of SIS on input \mathbf{A} is the set

$$\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m, \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}\},$$

which is a lattice. Finding a small vector in $\Lambda_q^\perp(\mathbf{A})$ is equivalent to finding a solution of SIS. We can therefore reformulate the problem SIS as “solving SVP in the lattice $\Lambda_q^\perp(\mathbf{A})$ for an uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ ”.

When SIS was introduced [Ajt96], its hardness was related to SVP_{n^c} for a certain $c > 1$. Further results [MR04, GPV08, MP13] have refined this dependency: if m and q satisfy $m = \text{poly}(n)$ and $q \geq \beta \cdot n^\epsilon$ for $\epsilon > 0$, then SVP_γ in dimension n reduces to $\text{SIS}_{q,n,m,\beta}$ for $\gamma = \beta \cdot \tilde{O}(\sqrt{n})$.

Learning With Errors (LWE). Another average-case problem central to modern lattice-based cryptography is the Learning With Errors problem²⁰, introduced by Oded Regev in 2005 [Reg05].

²⁰We present here the LWE problem with uniform secret and Gaussian error; other variants are considered in the literature, notably with restrictions on \mathbf{s} [Mic18] and other distributions on \mathbf{e} .

Definition I.2.3. Let $1 \leq n \leq m$ and $q \geq 2$ three integers and $\alpha \in [0, 1]$ a real parameter. The $\text{LWE}_{n,m,q,\alpha}$ problem ask to distinguish with probability $\geq 2/3$ between the two following distributions:

$$(\mathbf{A}, \mathbf{u}) \text{ and } (\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}),$$

where \mathbf{A} is a uniform matrix in $\mathbb{Z}_q^{m \times n}$, \mathbf{u} a uniform vector in \mathbb{Z}_q^m , \mathbf{s} a uniform vector in \mathbb{Z}_q^n and \mathbf{e} a vector of \mathbb{Z}_q^m sampled from the discrete Gaussian distribution of parameter $\alpha \cdot q$.

The LWE problem can be seen as distinguishing between a uniform vector and a vector close to a lattice, as shown in Figure I.13. It is also present in the literature in a “search” version, where only the pair $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$ is given, and we are asked to find \mathbf{s} .




Figure I.13: The two distribution of LWE.

As with SIS, one can relate the hardness of LWE to the worst-case problem SIVP. If $q > 2\sqrt{n}/\alpha$ and $m = \text{poly}(n)$, then SIVP_γ quantumly reduces²¹ [Reg05] to $\text{LWE}_{n,m,q,\alpha}$ for $\gamma = \tilde{O}(n/\alpha)$.

Figure I.14 summarizes the hardness relationships between the problems presented above. The worst-case part of the diagram is taken from the literature review by Noah Stephens-Davidowitz [Ste15, Page 1] and from [Mic08]. An arrow from problem A to problem B indicates that A reduces to B (so “ A is at most as difficult as B ”). A dashed arrow indicates a quantum reduction.

I.2.4 Structured lattices

Polynomials and lattices

In the problems described previously, the lattices are represented by their bases, in the form of a matrix. For a lattice of dimension n , this makes n^2 integers to be manipulated to perform our cryptographic operations (encryption, decryption, signatures...), which makes protocol running-time important when n becomes large. One way of solving this problem is to use matrices with a structure.

Let us define a version of the SIS problem over the polynomial ring $\mathbb{Z}[X]/(X^n + 1)$.

Definition I.2.4 (Ring-SIS $_{q,n,m,\beta}$ [Mic02, LM06, PR06]). Let $q \geq 2$, $n \geq 1$, $m \geq n \log(q)$, and $\beta \geq \sqrt{n \log(q)}$. The Ring-SIS $_{q,n,m,\beta}$ problem is defined as follows. Given uniform polynomials $P_1, \dots, P_k \in \mathbb{Z}_q[X]/(X^n + 1)$, Find $Q_1, \dots, Q_k \in \mathbb{Z}[X]/(X^n + 1)$ not all equal to zero

²¹The reduction requires a quantum calculator.




Figure I.14: Relative hardness of some lattice problems.

with $\|(Q_i)_i\| \leq \beta$ satisfying²²

$$P_1(X) \cdot Q_1(X) + \dots + P_k(X) \cdot Q_k(X) = 0 \pmod{(X^n + 1, q)}. \quad (\text{I.2})$$

If one chooses to represent polynomials as vectors, Equation (I.2) can be rewritten as follows

$$[\text{nrot}(P_1) | \dots | \text{nrot}(P_k)] \cdot [Q_1, \dots, Q_k]^T = \mathbf{0} \pmod{q},$$

where $\text{nrot}(P)$ is the nega-circulant matrix associated with the polynomial $P(X) = p_0 + p_1 \cdot X + \dots + p_{n-1} \cdot X^{n-1}$:

$$\text{nrot}(P) := \begin{pmatrix} p_0 & -p_{n-1} & -p_{n-2} & \cdots & -p_1 \\ p_1 & p_0 & -p_{n-1} & \cdots & -p_2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ p_{n-2} & p_{n-3} & p_{n-4} & \cdots & -p_{n-1} \\ p_{n-1} & p_{n-2} & p_{n-3} & \cdots & p_0 \end{pmatrix}.$$

With this rephrasing, we can see that the Ring-SIS problem is exactly the SIS problem restricted to matrices of the form $[\text{nrot}(P_1) | \dots | \text{nrot}(P_k)]$ with the P_i uniform in $\mathbb{Z}[X]/(q, X^n + 1)$. As arithmetic in the ring $\mathbb{Z}_q[X]/(X^n + 1)$ is made faster by the use of the discrete Fourier transform (which allows multiplication of two polynomials in quasi-linear time), cryptographic systems built from Ring-SIS will be more efficient than those built from SIS. It is also possible to define the Ring-LWE [SSTX09, LPR10, PRS17] problem in a similar fashion.

That being said, the increased efficiency of protocols using structured lattices comes with stronger security assumptions, which need to be studied specifically.

²²There are multiple ways to define norms over $\mathbb{Z}[X]/(X^n + 1)$, for a presentation of the norm we actually use, see Section II.2.1. If $n = 2^k$, this norm is the same - up to a multiplicative factor - to the Euclidean norm of the coefficient vector of the polynomial.




Figure I.15: Graphical representation of matrices.

Unstructured ($n = 9$) ; Rank-1 module ($d = 9, k = 1$) ; Rank-3 module ($d = 3, k = 3$).

Algebraic number theory

A natural way of defining more generic security assumptions is to extend our structured computational problems to generic polynomial rings of the form $\mathbb{Z}[X]/P(X)$ for any polynomial $P \in \mathbb{Z}[X]$. However, depending on the choice of the polynomial P , this approach can lead to security problems (for example, in the ring $\mathbb{Z}[X]/(X^n - 1)$ the problem Ring-SIS is solvable in reasonable time with good probability under certain conditions on n [PR06]). To construct polynomial rings with good properties, we rely on number theory. If $P \in \mathbb{Z}[X]$ is an irreducible polynomial of degree d , the ring $K = \mathbb{Q}[X]/P(X)$ is a field, called “number field of degree d ”, and we can define its ring of integers \mathcal{O}_K as the set of its elements satisfying a monic integer polynomial equation. This ring of integers is always a sub-ring of a polynomial ring and has a number of properties, including being a Dedekind domain and the only maximal order of K . We do not go into details about what these properties mean; intuitively, they are a guarantee that \mathcal{O}_K does not have algebraic properties “too different from \mathbb{Z} ”. In particular, if K is the cyclotomic field of degree $d = 2^n$, then the ring \mathcal{O}_K is exactly the polynomial ring $\mathbb{Z}[X]/(X^{2^n} + 1)$. Throughout the rest of the manuscript, readers more familiar with polynomial rings than with number theory are encouraged to replace all occurrences of \mathcal{O}_K by $\mathbb{Z}[X]/(X^d + 1)$ for $d = 2^n$.

In the rest of this introduction, K is a number field of degree $d \geq 2$ and ring of integers \mathcal{O}_K . A module over K of rank $k \geq 1$ is a subset of K^k , stable by addition, subtraction and multiplication by elements of \mathcal{O}_K ²³. A module $M \subset K^k$ is associated with a lattice $\Phi(M) \subset \mathbb{R}^{d \cdot k}$ via $\Phi(\cdot)$, the canonical embedding of the field²⁴, we can therefore restrict all the lattice problems presented above to module lattices, with the advantage that the structure of K allows us to perform operations on module lattices more efficiently than on non-module lattices of the same dimension.

The most commonly used field in practice is the cyclotomic field of conductor a power of 2, i.e. the field $K = \mathbb{Q}(\zeta_{2^n})$. With this choice of field, the free²⁵ modules of rank k are (up to rotation and scaling) exactly the lattices generated by block matrices of the form

$$\mathbf{B} = \begin{pmatrix} \text{nrot}(P_{1,1}) & \dots & \text{nrot}(P_{1,k}) \\ \vdots & & \vdots \\ \text{nrot}(P_{k,1}) & \dots & \text{nrot}(P_{k,k}) \end{pmatrix}$$

for $(P_{i,j})_{1 \leq i,j \leq k} \in \mathbb{Q}[X]/(X^{2^{n-1}} + 1)^{k \times k}$ (see Fig. I.15). The structure of this field was no-

²³In this manuscript, we only talk about finitely generated \mathcal{O}_K -modules without torsion

²⁴Or the embedding by coefficients.

²⁵All modules are not free, but we restrict ourselves to free modules for this introduction.

tably used for the signature (CRYSTALS-Dilithium [LDK⁺20]) and key exchange (CRYSTALS-Kyber [ABD⁺19]) algorithms selected by the NIST [NIST] competition.

Some particularities of the cyclotomic field (in particular its large number of sub-fields) have led some authors to propose the use of different fields for cryptographic protocols, in order to “reduce the attack surface” (the field proposed in [BCLV17] is the field defined by the polynomial $X^p - X - 1$, for p a prime number). We have tried to be agnostic about the choice of field in this manuscript. When more precise results are available on cyclotomics, we mention it.

As with unstructured lattices, we can then define the problem SVP_γ restricted to module lattices. If K is a number field, $\gamma \geq 1$ an approximation factor and $k \geq 1$ a rank, the problem $\text{modSVP}_{k,\gamma}^K$ is the problem SVP_γ restricted to rank k modules over the field K . As previously stated, $\text{modSVP}_{k,\gamma}^K$ is a subproblem of SVP_γ in dimension $d \cdot k$.

Note that unlike the case of unstructured lattices, here there are two variables to adjust to increase problem difficulty: the rank k of the lattice and the degree d of the field. Unlike unstructured lattices, here it is the degree of the field that we increase to increase problem difficulty. In particular, we are interested in rank-1 and rank-2 lattices in increasingly large fields. The variables on which the running-time of our algorithms and our approximation factors depend on are the degree d of the field, which corresponds to the size of the lattices, and the root-discriminant $\Delta_K^{1/(2d)}$ of the field, which - roughly speaking - represents the size of the number field.

On id-HSVP

The most elementary case of a module lattice is the rank 1 case, which corresponds to finding a short vector in a fractional ideal of K . This problem is therefore called id-HSVP_γ ²⁶. The difficulty of this problem is not yet precisely understood. It seems that for some large parameters, id-HSVP_γ is easier than SVP_γ : there are algorithms to solve in polynomial time it for approximation factors $\gamma \geq 2^{\tilde{O}(\sqrt{d})}$ for cyclotomic fields [CDPR16, CDW17] or on any field K if pre-calculations have been performed [PHS19]. For $\gamma = \text{poly}(d)$ however, the best algorithms currently known are those used on unstructured lattices (see Figure I.16).

Ideals are the simplest example of lattice modules, so understanding the difficulty of id-HSVP is an important matter. Their algebraic structure has also made it possible to develop worst-case to average-case reductions.

Relationship between id-HSVP and other module lattice problems. The average-case problems Ring-SIS and Ring-LWE have been linked to the search for small vectors in ideal lattices. It has been proved that id-HSVP_γ reduces to the Ring-LWE [SSTX09, LPR10] problem and the Ring-SIS [PR06, LM06] problem for polynomial approximation factors. These results should be seen as lower bounds of difficulty on Ring-SIS and Ring-LWE: as these problems are not defined on ideal lattices directly, attacks on id-HSVP do not extend to Ring-SIS or Ring-LWE.

Worst-case to average-case self reduction for id-HSVP . In [Gen09], Gentry proves that (with a factorization oracle), id-HSVP in the worst-case reduces to solving id-HSVP with good probability when the input is the inverse of a small uniform prime ideal.

²⁶The H - for “Hermite” Short Vector Problem - comes from the fact that for ideal lattices, it is more natural to look for small vectors with respect to the algebraic norm of the ideal and not with respect to λ_1 . It can be shown that those two approaches are equivalent, see Lemma II.4.9.




Figure I.16: Time/approximation trade-off for id-HSVP with $\deg(K) = d$ and exponential pre-computations in d [PHS19, Fig. 2].

A similar result was proven in 2022 by de Boer et. al. [BDPW20], where id-HSVP in the worst-case is reduced to id-HSVP with good probability when the input is the Gaussian rounding of a uniform ideal in the set of norm 1 replete ideals.

The NTRU problem

The NTRU²⁷ problem, introduced in 1998 [HPS98] (in a version slightly different from the one presented here) is another modular polynomial equation problem, where given a polynomial $h \in \mathbb{Z}[X]/(X^n + 1)$, we have to find a writing $h = g/f \pmod q$ with f and g small polynomials (with the promise that such polynomials exist):

Definition I.2.5 ((γ, γ', q) -NTRU). *Let $q \geq 2$ and $\gamma \geq \gamma' > 0$. A (γ, q) -NTRU instance is a polynomial $h \in \mathbb{Z}[X]/(X^n + 1)$ such that $h = g/f \pmod q$ with $\|f\|, \|g\| \leq \sqrt{q}/\gamma$. The (γ, γ', q) -NTRU problem asks, given a (γ, q) -NTRU instance h , to find \tilde{f}, \tilde{g} with norm $\leq \sqrt{q}/\gamma'$ satisfying $h = \tilde{g}/\tilde{f} \pmod q$.*

The NTRU problem can be generalized by taking $f, g, h \in \mathcal{O}_K$ for a number field K . It was noticed early on that the NTRU problem can be interpreted in terms of lattices [HPS98, CS97]. Indeed, the set

$$L_h := \begin{bmatrix} 1 \\ h \end{bmatrix} \mathcal{O}_K + \begin{bmatrix} 0 \\ q \end{bmatrix} \mathcal{O}_K = \left\{ (\tilde{f}, \tilde{g})^T \in \mathcal{O}_K^2, \quad h \cdot \tilde{f} = \tilde{g} \pmod q \right\}$$

is a module lattice of rank 2. This lattice is described by h , from which a basis can be computed, and has a particular property: it contains an unusually short non-zero vector (f, g) . Indeed, for most h 's, we have $\det L_h = \Delta_K \cdot q^d$, where Δ_K refers to the field discriminant; our running example satisfies $\Delta_K = d^d$. As a result, one would expect the shortest non-zero vectors to have ℓ_2 -norm around $q^{1/2}$, up to limited factors depending on Δ_K and d ;²⁸ but $(f, g)^T$ is much shorter,

²⁷The meaning of the name seems to have been lost.

²⁸This is a consequence of the Gaussian heuristic (see, e.g., [GNR10]), which describes how a “typical” looks like.




Figure I.17: On the left a typical lattice, on the right a lattice with an unusually short non-zero vector.

by assumption. We therefore have a rank 2 module on \mathcal{O}_K with the promise that it contains an unusually short non-zero vector, i.e. an unusually dense rank 1 submodule. We call the problem of finding a short non-zero vector in a rank 2 module containing a short vector mod-uSVP_2 (see Figure I.17).

NTRU's hardness. The NTRU and mod-uSVP_2 problems actually exist in two forms. The most natural, described above, involves recovering a short vector of the corresponding rank 2 module. This is the variant we implicitly consider in this introduction when discussing NTRU and mod-uSVP_2 . The other versions considered in this manuscript require finding a basis of the densest submodule (i.e. generated by the unusually short vector) rather than finding the short vector directly: they will be written with a mod exponent: NTRU^{mod} and $\text{mod-uSVP}_2^{\text{mod}}$. Both versions of this problem are equivalent if a id-HSVP oracle is given. As seen above, the NTRU problem can be considered a special case of mod-uSVP_2 , however, it is not clear whether its instances are representative of all mod-uSVP_2 instances. In [Pei16, Section 4.4.4], Peikert sketches a reduction from a decision version of the NTRU problem to the Ring-LWE [SSTX09, LPR10] problem; this reduction can be adapted to the search NTRU problem we are considering here. Note that under certain parameter constraints, the Ring-LWE problem is equivalent to mod-SIVP_2 [LS15, AD17], which is the restriction to rank 2 module lattices of the SIVP problem described above.

Links between NTRU and id-HSVP

In the other direction, Pellet-Mary and Stehlé [PS21] presented a reduction of the shortest vector problem for ideal lattices over \mathcal{O}_K (id-HSVP) to NTRU. Overall, we see that NTRU lies between id-HSVP and mod-SIVP . As noted earlier, id-HSVP admits better algorithms than generic lattice reduction algorithms [LLL82, Sch87] for certain parameter ranges [CDW21, PHS19]. Since such a phenomenon is unknown in the case of mod-SIVP , there could be a big jump in difficulty between id-HSVP and mod-SIVP . It is currently unclear which of these problems captures the true difficulty of NTRU, or whether NTRU lies somewhere in between.

A partial presentation of the difficulty of some module lattices problems is available in Figure I.18. All the problems described are for the same field K ; for the sake of readability, we

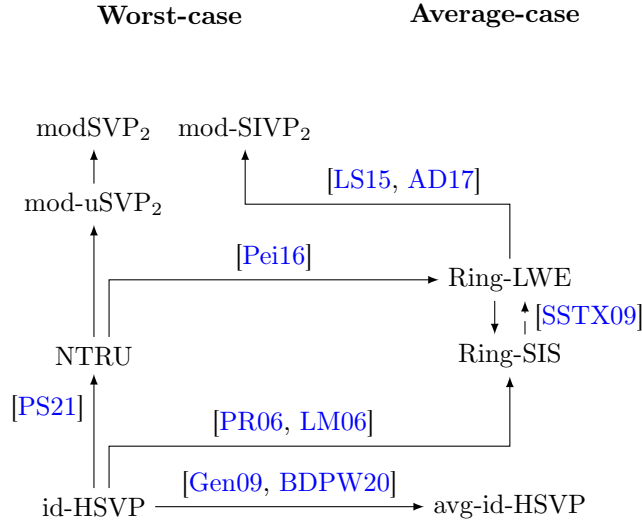


Figure I.18: Relative hardness of some module-lattices problems (approximations factors omitted).

do not show approximation factor losses or application conditions. As before, an arrow from A to B indicates that A reduces to B . Dashed arrows denote quantum reductions, arrows without associated citation are trivial or folklore reductions.

I.2.5 Contribution of this PhD

Publications

This manuscript is based on the two publications produced during my thesis:

- [FPS22] **On Module Unique-SVP and NTRU**. Joël Felderhoff, Alice Pellet-Mary and Damien Stehlé. ASIACRYPT 2022.
- [FPSW23] **Ideal-SVP is Hard for Small-Norm Uniform Prime Ideals**. Joël Felderhoff, Alice Pellet-Mary, Damien Stehlé and Benjamin Wesolowski. TCC 2023.

Hardness of id-HSVP for random integral ideals

The two average-case distributions of ideals for which a worst-case to average-case reduction exists in the literature are not satisfying. The first [Gen09] is a distribution of inverse ideals which, in addition to being unnatural from an algebraic point of view, does not allow us to use the reduction of [PS21] (which is valid only for integral ideals) to describe a distribution NTRU whose difficulty is based on id-HSVP in the worst-case. The second [BDPW20] is indeed a distribution of integral ideals, but the rounding process at the heart of the reduction only allows us to expect ideals of very large algebraic norm (of the order of $2^{O(d^2)}$). Using ideals of this size for cryptography would lead to such high computational costs that the use of structured lattices would lose all interest. In particular, if the reduction of [PS21] is applied with this distribution, it leads to NTRU instances with exponential modulus.

We propose in Chapter IV (from [FPSW23]) a new average-case to average-case reduction (with factorization oracle) for id-HSVP. Let \mathcal{W} be any set of ideals and \mathcal{W}^{-1} the set of inverses of the elements of \mathcal{W} . We prove in Theorem IV.5.1 that solving id-HSVP over a uniform ideal of \mathcal{W}^{-1} reduces to solving id-HSVP over a small uniform integral ideal and over a uniform ideal of \mathcal{W} .

By specializing this result with \mathcal{W} being the set of prime ideals of small norm ($\text{poly}(d, \Delta_K^{1/d})^d$) we show in Corollary IV.5.4 that the difficulty of id-HSVP on the distribution described in [Gen09] is no different from that of id-HSVP on the uniform distribution on small prime ideals, and thus that id-HSVP in the worst-case reduces to id-HSVP on a small uniform prime ideal.

Thanks to the reduction of [PS21], this new distribution allows us to define a distribution on polynomial modulus NTRU instances based on the difficulty of id-HSVP in the worst case (Corollary IV.6.2). However, this distribution requires a factoring oracle in order to be used in a cryptographic context.

Relationship between NTRU and mod-uSVP₂, and average-case distribution for mod-uSVP₂

As mentioned earlier, the NTRU problem is a special case of the mod-uSVP₂ problem, where the module is of the form $(1, h)^T \cdot \mathcal{O}_K + (0, q)^T \cdot \mathcal{O}_K$. In Chapter V, we prove that these NTRU modules are in fact representative of general mod-uSVP₂ lattices, in the sense that any mod-uSVP₂ module can be transformed in polynomial time into a NTRU lattice of similar geometry, this is Theorem V.4.1. We accompany this result with reductions for worst-case and average-case variants of mod-uSVP₂^{mod} (respectively mod-uSVP₂^{vec}) to NTRU^{mod} (respectively NTRU^{vec}).

We then propose in Theorem V.6.1 a worst-case to average-case reduction for mod-uSVP₂ instances, which runs in polynomial time with a call to a factorization oracle (hence in quantum polynomial time). This result allows us to give another average-case distribution on NTRU instances whose hardness is based on a worst-case problem by combining Theorems V.4.1 and V.6.1. However, this distribution requires a factoring oracle to be used and seems rather artificial.

Contributions of this manuscript compared to [FPS22, FPSW23]

As mentioned earlier, this manuscript is largely based on the two papers published during my thesis: [FPS22, FPSW23]. However, we have added an additional result in Chapter III which allows us to refine certain theorems of [FPS22] and [FPSW23].

Let K be a number field of degree ≥ 3 . A classical result of analytic number theory is that, when X tends to infinity, the number of ideals of algebraic norm less than X , denoted by $N_K(X)$, is equivalent to $\rho_K \cdot X$, where ρ_K is the residue in $s = 1$ of the Dedekind zeta function associated with K . In some of our results, we needed to bound the error term in this approximation, but the literature bounds did not explicitly show the dependence on K . In chapter III we propose a, new to our knowledge, expression of the bound on this error where the dependence in the field is specified, this is Theorem III.1.2.

In [FPSW23], the running time of the worst-case reduction from id-HSVP to id-HSVP for a uniform prime ideal [FPSW23, Corollary 5.4] depends on an ad hoc parameter $\tilde{\rho}_A$, which is the inverse proportion of prime ideals in the field of norm less than A . It is known that $\tilde{\rho}_A$ behaves asymptotically as $\rho_K \cdot \ln(A)$, but the error bounds on this approximation in the literature did not allow us to apply this result for the values of A that interest us. Corollary III.1.3 allows us to give a reduction whose running-time depends only on ρ_K , and no longer on $\tilde{\rho}_A$. This leads to larger ideal sizes and approximation factors, which are presented in Corollary IV.5.5. We also use our result about $N_K(\cdot)$ in Chapter V. The reduction of mod-uSVP₂ to NTRU (in this manuscript, Theorem V.4.1 or [FPS22, Theorem 4.1]) had a running-time depending on $\zeta_K(2)$

which, depending on the number field considered, can be exponential in degree. We use the results of Chapter III to propose a new version of this theorem where the running-time no longer depends on $\zeta_K(2)$ (at the cost of a larger approximation factor loss). The theorem in question is Theorem V.4.2, the proof of which is in Section D.5.

A second modification has been made to Theorem V.4.1 in comparison with [FPS22, Theorem 4.1]. In [FPS22, Theorem 4.1], a condition on the number field K is present: we require that $\zeta_K(2) = 2^{o(d)}$. This condition is non-trivial, as it has implications on the arithmetic of K (it can be interpreted as the fact that “ \mathcal{O}_K does not have too many small ideals”). This condition originated from the proof of [FPS22, Lemma 4.3] and was purely technical. We have modified the proof of this lemma (in this manuscript, Lemma V.4.4) to remove it.

Chapter II

Preliminaries

We let \mathbb{Z} denote the set of integers, \mathbb{Q} the rationals, \mathbb{R} the real numbers and \mathbb{C} the complex plane. In this manuscript, the notation \ln will refer to the base- e logarithm and the \log notation will refer to the base-2 logarithm. For any positive function $f : X \rightarrow \mathbb{R}_{\geq 0}$ and $S \subseteq X$ with S countable, we define $f(S) := \sum_{x \in S} f(x)$.

Let f and g two functions depending on several variables x_1, \dots, x_n varying over some set X_1, \dots, X_n . We will use in this work both the Vinogradov notation $f \ll g$ and the most classical Landau notation $f = O(g)$ (we mostly use Vinogradov notation in Chapter III and Landau notation in the rest of the manuscript). Writing $f \ll g$ or $f = O(g)$ means that there exists an absolute constant C (i.e., a constant independent of any variable of the context) such that $|f| \leq C \cdot |g|$ for any x_1, \dots, x_n in X_1, \dots, X_n except maybe a finite number of them independent of any variable of the context.

We consider column vectors (unless they are explicitly transposed). Vectors and matrices are respectively written in bold lowercase and uppercase fonts. For a vector $\mathbf{x} \in \mathbb{C}^k$, we let $\|\mathbf{x}\|$ denote its Hermitian norm. When $\mathbf{M} = [\mathbf{M}_1, \dots, \mathbf{M}_n]$ is a matrix, we denote its norm $\|\mathbf{M}\| = \max_i \|\mathbf{M}_i\|$ the maximal norm of its columns. Note that for any vector $\mathbf{x} \in \mathbb{C}^n$, it holds that

$$\|\mathbf{M} \cdot \mathbf{x}\| \leq \sqrt{n} \cdot \|\mathbf{M}\| \cdot \|\mathbf{x}\|.$$

The Hermitian product will be denoted, for any $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$, by

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n \bar{x}_i \cdot y_i.$$

For any integer $n \geq 0$, we will denote by $[n]$ the set $\{1, \dots, n\}$.

II.1 Lattices

II.1.1 Definitions and bases

Lattices are the mathematical object we will be the more concerned with in this work. We define them and state some of their properties here. For a more detailed presentation of those properties, including proof, the reader is referred to [Coh93] and [MG02].

We define a lattice as follows (other equivalent definitions exist).

Definition II.1.1. *Let $n \geq m \geq 1$ be integers. A lattice in \mathbb{R}^n is a set of the form $L = \sum_{i=1}^m \mathbb{Z} \cdot \mathbf{b}_i$ for $(\mathbf{b}_i)_{1 \leq i \leq m}$ a \mathbb{R} -linearly independent family of vectors in \mathbb{R}^n .*

If, for an \mathbb{R} -linearly independent family $(\mathbf{b})_{1 \leq i \leq m}$, one has $L = \sum_{i=1}^m \mathbb{Z} \cdot \mathbf{b}_i$, the family $(\mathbf{b})_{1 \leq i \leq m}$ is called a basis of L ; we shall write $L = \mathcal{L}(\mathbf{B})$ where \mathbf{B} is the matrix whose columns are the $(\mathbf{b}_i)_i$. The integer m is then called the rank of L , and the lattice is said to be full rank if $m = n$, the lattice L is said to be full rank. We will often give the basis of a lattice $L \subset \mathbb{R}^n$ as a matrix $\mathbf{B} \in \mathbb{R}^{n \times m}$, in that case we have

$$L = \mathcal{L}(\mathbf{B}) := \mathbf{B} \cdot \mathbb{Z}^m$$

Two bases $\mathbf{B}_1, \mathbf{B}_2 \in \mathbb{R}^{n \times m}$ define the same lattice in \mathbb{R}^n if and only if there exists $\mathbf{U} \in GL_m(\mathbb{Z})$ such that $\mathbf{B}_1 = \mathbf{B}_2 \cdot \mathbf{U}$. If L is a lattice in \mathbb{Z}^n , then it has a special basis called its Hermite Normal Form (HNF):

Lemma II.1.2. *For $1 \leq n \leq m$, let $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_m) \in \mathbb{Z}^{n \times m}$ be m integral vectors spanning a full rank lattice $\mathcal{L}(\mathbf{B}) \subseteq \mathbb{Z}^n$. There exists a unique matrix $\mathbf{H} = (h_{i,j}) \in \mathbb{Z}_{\geq 0}^{n \times n}$ satisfying*

- \mathbf{H} is lower triangular;
- $0 \leq h_{i,j} < h_{i,i}$ for any $1 \leq j < i \leq n$;
- $\mathcal{L}(\mathbf{H}) = \mathcal{L}(\mathbf{B})$.

This matrix, that we call the Hermite Normal Form (HNF) of \mathbf{B} , is computable in polynomial time in n and in $\log(\|\mathbf{B}\|)$.

If the matrix associated to a basis of a lattice is in Hermite normal form, we say that the basis is in HNF.

If $L \subseteq \mathbb{Q}^n$, there exists a smallest integer $a \geq 1$ such that $a \cdot L \subset \mathbb{Z}^n$, and hence we call the Hermite normal form of L the matrix \mathbf{H}/a . Note that such HNF does not exist for lattices in \mathbb{R}^n . The Hermite normal form gives the lattice a canonical basis, which allows to decide if two bases span the same lattice in polynomial time.

For any basis $(\mathbf{b}_1, \dots, \mathbf{b}_k) \in \mathbb{R}^{n \times k}$ of a lattice, we let $(\mathbf{b}_1^*, \dots, \mathbf{b}_k^*)$ denote its Gram-Schmidt vectors, that is to say

$$\mathbf{b}_i^* = \mathbf{b}_i - \sum_{1 \leq j < i} \frac{\langle \mathbf{b}_j^*, \mathbf{b}_i \rangle}{\langle \mathbf{b}_j^*, \mathbf{b}_j^* \rangle} \cdot \mathbf{b}_j^*, \quad (\text{II.1})$$

for $1 \leq j \leq k$. We will often let \mathbf{B}^* denotes the Gram-Schmidt vectors of the columns of a matrix \mathbf{B} . Equivalently, Equation (II.1) shows that $\mathbf{B} = \mathbf{B}^* \cdot \mathbf{S}$, where \mathbf{S} is upper-triangular; by dividing each column of \mathbf{B}^* by its norm, this identity becomes $\mathbf{B} = \mathbf{Q} \cdot \mathbf{R}$ where \mathbf{Q} is orthogonal and \mathbf{R} upper-triangular, the QR-factorization of \mathbf{B} .

Definition II.1.3. *Let $L \subset \mathbb{R}^n$ be a lattice, the $(\mathbb{Z}$ -)dual of this lattice is*

$$L^* = \{\mathbf{x} \in \text{span}(L), \langle \mathbf{x}, \mathbf{l} \rangle \in \mathbb{Z} \text{ for any } \mathbf{l} \in L\},$$

where $\text{span}(L) = L \otimes_{\mathbb{Z}} \mathbb{R}$ is the subspace of \mathbb{R}^n spanned by L .

If \mathbf{B} is a basis of L , then $\mathbf{B} \cdot (\mathbf{B}^T \cdot \mathbf{B})^{-1}$ is a basis of L . Note that if L is full rank in \mathbb{R}^n (so that the matrix \mathbf{B} is square), this reduces to $\mathbf{B}^{-T} = (\mathbf{B}^T)^{-1}$.

II.1.2 Lattice invariants

If L is a lattice with basis $\mathbf{B} \in \mathbb{R}^{n \times m}$, we define its (co-)volume as

$$\text{Vol}(L) = \sqrt{\det(\mathbf{B} \cdot \mathbf{B}^T)}.$$

This quantity is independent on the choice of basis. It holds that

$$\text{Vol}(L^*) = \text{Vol}(L)^{-1}.$$

For any $1 \leq i \leq m$, the i th successive minimum of a lattice L is

$$\lambda_i(L) = \min \{ \max(\|\mathbf{x}_1\|, \dots, \|\mathbf{x}_i\|), \mathbf{x}_1, \dots, \mathbf{x}_i \in L \text{ are linearly independent} \}.$$

The relationship between the $\lambda_i(L)$ and the volume is given by Minkowski's theorem.

Theorem II.1.4 (Minkowski's theorems). *Let $L \subset \mathbb{R}^n$ be a lattice of rank k . Then it holds that*

$$\lambda_1(L) \leq \sqrt{k} \cdot \text{Vol}(L)^{\frac{1}{k}}, \quad (\text{Minkowski's First Theorem})$$

$$\prod_{1 \leq i \leq k} \lambda_i(L) \leq k^{\frac{k}{2}} \cdot \text{Vol}(L). \quad (\text{Minkowski's Second Theorem})$$

We denote by $\mu(L)$ the covering radius of a lattice L :

$$\mu(L) = \max_{x \in \mathbb{R}^n} d(x, L).$$

Theorem II.1.5 (Banaszczyk's transference theorem [Ban93]). *For any rank lattice $L \subset \mathbb{R}^n$ of rank k , the following holds for any $1 \leq i \leq k$*

$$1 \leq \lambda_i(L) \cdot \lambda_{k-i+1}(L^*) \leq k.$$

II.1.3 The Gaussian distribution

Let $\varsigma > 0$ be a real number, the Gaussian weight of parameter ς of a vector $\mathbf{x} \in \mathbb{R}^n$ is

$$\rho_\varsigma(\mathbf{x}) = \exp\left(-\pi \cdot \frac{\|\mathbf{x}\|^2}{\varsigma^2}\right)$$

Given an n -dimensional lattice L , a vector $\mathbf{u} \in \mathbb{R}^n$ and a parameter $\varsigma > 0$, we define the discrete Gaussian distribution $D_{L, \varsigma, \mathbf{u}}$ over L with center \mathbf{u} and standard deviation parameter ς by

$$D_{L, \varsigma, \mathbf{u}}(\mathbf{x}) := \rho_\varsigma(\mathbf{x} - \mathbf{u}) / \rho_\varsigma(L - \mathbf{u})$$

for all $\mathbf{x} \in L$.

If L is a lattice and $\varepsilon > 0$, the smoothing parameter is defined as

$$\eta_\varepsilon(L) := \inf \{ \varsigma > 0, \rho_{1/\varsigma}(L^* \setminus \{0\}) \leq \varepsilon \}.$$

This parameter measures how large ς needs to be for the lattice Gaussian distribution to behave like a continuous Gaussian distribution.

Lemma II.1.6 (Proof of [MR07, Lemma 4.4]). *Let L be a rank n lattice, $\mathbf{u} \in \text{span}(L)$ and $\varsigma \geq \eta_\varepsilon(L)$ for some $\varepsilon > 0$. Then it holds that*

$$\rho_\varsigma(L + \mathbf{u}) \in [1 - \varepsilon, 1 + \varepsilon] \cdot \frac{\varsigma^n}{\text{Vol}(L)}.$$

Lemma II.1.7 ([Reg05, Claim 3.8]). *For any $\varepsilon > 0$, lattice L , center $\mathbf{u} \in \text{span}_\mathbb{R}(L)$ and parameter $\varsigma \geq \eta_\varepsilon(L)$, it holds that*

$$\rho_\varsigma(L + \mathbf{u}) \in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, 1 \right] \cdot \rho_\varsigma(L).$$

The smoothing parameter is related to the n th minimum of the lattice and to the Gram-Schmidt of their basis.

Lemma II.1.8 ([GPV08, Lemmas 3.1 and 3.2]). *Let $L \subset \mathbb{R}^n$ be a full rank lattice given by a basis $\mathbf{B} \in \mathbb{R}^{n \times n}$. Let \mathbf{B}^* be the Gram-Schmidt of the columns of \mathbf{B} . For any $\varepsilon \in (0, 1)$, it holds that*

$$\eta_\varepsilon(L) \leq \min(\lambda_n(L), \|\mathbf{B}^*\|) \cdot \sqrt{\frac{\log(2n(1 + \frac{1}{\varepsilon}))}{\pi}}$$

The following lemma shows that most of the weight of the discrete Gaussian distribution is within distance $O(\varsigma \cdot \sqrt{n})$ of its center:

Lemma II.1.9 ([Ban93, Lemma 1.5]). *For any $c > 1/\sqrt{2\pi}$, $\varsigma > 0$, any n -dimensional lattice L and any $\mathbf{u} \in \text{span}(L)$, we have*

$$\frac{\rho_\varsigma((L - \mathbf{u}) \setminus c \cdot \sqrt{n} \cdot \varsigma \cdot \mathcal{B})}{\rho_\varsigma(L)} \leq 2C^n,$$

where \mathcal{B} denotes the Euclidean ball of radius 1 and $C = c\sqrt{2\pi}e \cdot e^{-\pi c^2} < 1$.

Corollary II.1.10. *Let L be a rank n lattice, \mathbf{B} be a basis of L , $\mathbf{u} \in \text{span}(L)$ and $\varsigma \geq \sqrt{n} \cdot \max_i \|\mathbf{b}_i^*\|$. For any $\varepsilon \in (0, 1]$, it holds that*

$$\Pr_{\mathbf{x} \leftarrow D_{L, \varsigma, \mathbf{u}}} (\|\mathbf{x} - \mathbf{u}\| \geq \varsigma \cdot \sqrt{\ln(1/\varepsilon) + 4n}) \leq \varepsilon.$$

Proof. Without loss of generality, we can scale everything so that $\varsigma = 1$. Let us define $c := \sqrt{(1/n) \cdot \ln(1/\varepsilon) + 4}$. Then, we have

$$\Pr_{\mathbf{x} \leftarrow D_{L, 1, \mathbf{u}}} (\|\mathbf{x} - \mathbf{u}\| \geq \sqrt{\ln(1/\varepsilon) + 4n}) = \frac{\rho_1((L - \mathbf{u}) \setminus c\sqrt{n}\mathcal{B})}{\rho_1(L - \mathbf{u})}.$$

Since $c \geq \sqrt{4} > 1/\sqrt{2\pi}$, we can apply Lemma II.1.9 to bound the numerator from above. In order to simplify the computations, we use the fact that $c \cdot e^{-\pi c^2} \leq e^{-c^2}$ for all $c > 1/\sqrt{2\pi}$. Then we see that $6 \cdot C^n \leq 6^n \cdot C^n \leq e^{-\ln(1/\varepsilon) + (\ln(\sqrt{2\pi}e) + \ln(6) - 4)n} \leq e^{-\ln(1/\varepsilon)} = \varepsilon$. Using Lemma II.1.9, we hence obtain the bound

$$\rho_1((L - \mathbf{u}) \setminus c\sqrt{n}\mathcal{B}) \leq \varepsilon/3 \cdot \rho_1(L).$$

Let us now bound the quantity $\rho_1(L - \mathbf{u})$ from below. Using Lemma II.1.7 with $\varepsilon = 1/2$ (observe that $\varsigma \geq \eta_{1/2}(L)$ by observing that $\ln(6n/\pi) \leq n$ for all $n \geq 1$ and Lemma II.1.8), we see that $\rho_1(L - \mathbf{u}) \geq 1/3 \cdot \rho_1(L)$. Combining both inequalities provides the desired result. \square

Sampling along discrete Gaussian distribution for lattices is an active field of research. In this work, we will use two different Gaussian samplers: one which sample from the exact discrete Gaussian distribution, and the other which is tail-cut in order to garanty the size of the output. We now present the exact sampler.

Lemma II.1.11 (Exact Gaussian Sampler, from [BLP⁺13, Lemma 2.3]). *There is a probabilistic polynomial-time algorithm that, given a basis \mathbf{B} of an n -dimensional lattice L , a center $\mathbf{u} \in \text{span}_{\mathbb{R}}(L)$, and a parameter $\varsigma \geq \sqrt{\ln(2n+4)}/\pi \cdot \|\mathbf{B}^*\|$, outputs a sample distributed according to $D_{L,\varsigma,\mathbf{u}}$.*

The following lemma is adapted from [GPV08, Theorem 4.1] and [PS21, Lemma 2.2]. We will notably be interested in values of ε that are $2^{-\omega(n)}$, a case which is not captured in the typical variants of this statement. For the sake of completeness, a proof is available in Appendix A.1.5.

Lemma II.1.12. *There exists a probabilistic polynomial time algorithm that takes as input a basis \mathbf{B} of an n -dimensional lattice L , an error bound $\varepsilon \in (0, 1/2]$, a parameter $\varsigma \geq \sqrt{n} \cdot \|\mathbf{B}^*\|$ and a center $\mathbf{u} \in \text{span}(L)$ and outputs a sample from a distribution $\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}$ such that*

- $\text{SD}(D_{L,\varsigma,\mathbf{u}}, \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}) \leq \varepsilon$;
- for all $\mathbf{v} \leftarrow \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}$, it holds that $\|\mathbf{v} - \mathbf{u}\| < \varsigma \cdot \sqrt{\ln(1/\varepsilon) + 4n}$;

where SD is denotes the statistical distance between two distributions (see Definition II.5.1).

II.1.4 The LLL and BKZ algorithms

Two algorithms will be used in this work to find relatively short vectors in lattices: the LLL (introduced in [LLL82]) and BKZ (introduced in [SE94]) algorithms. We will not present them in full generality, but will rather highlight their relevant properties for our work.

Lemma II.1.13 ([Gal12, Theorem 17.2.12]). *The LLL algorithm takes as input a basis of a full rank lattice $L \subset \mathbb{Q}^n$ and output a basis $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{Q}^{n \times n}$ of L satisfying:*

- $\|\mathbf{b}_i\| \leq 2^{(n-1)/2} \cdot \lambda_i(L)$ for $1 \leq i \leq n$,
- $\|\mathbf{b}_1\| \leq 2^{(n-1)/4} \cdot \text{Vol}(L)^{1/n}$,
- $\prod_{1 \leq i \leq n} \|\mathbf{b}_i\| \leq 2^{n(n-1)/4} \cdot \text{Vol}(L)$.

It runs in polynomial time in n and in $\log(B)$ where B is a bound on the norm of the vectors of the input basis.

In particular, note that the LLL algorithm solves in polynomial time the $2^{(n-1)/2}$ -SVP and the $2^{(n-1)/4}$ -HSVP. Better approximation factors can be obtained at the cost of a higher running time using the BKZ algorithm.

Lemma II.1.14 ([Sch87, Theorem 2.3, Corollary 2.5]). *For any $\beta \geq 1$, there exists $\alpha_\beta \in \mathbb{R}_{>0}$ such that for any $n \geq 1$, such that $\beta - 1 | n - 1$, on input a basis of a full rank lattice $\mathcal{L} \subset \mathbb{Q}^n$, the BKZ algorithm with size block β outputs a basis $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{Q}^{n \times n}$ of L with*

$$\|\mathbf{b}_1\| \leq \alpha_\beta^{\frac{n-1}{2(\beta-1)}} \cdot \lambda_1(L),$$

where $\alpha_\beta \leq \beta^{1+\ln(\beta)}$ for any β . Furthermore, the BKZ algorithm with size block β runs in time polynomial in 2^β , n and $\log(B)$ where B is a bound on the norm of the vectors of the input basis.

In particular, the BKZ algorithm solves SVP_γ for $\gamma = 2^{\tilde{O}(n/\beta)}$ in time polynomial in 2^β . By increasing the dimension n by at most $\beta - 1$, the condition $\beta - 1 | n - 1$ can be removed and the bound on $\|\mathbf{b}_1\|$ becomes

$$\|\mathbf{b}_1\| \leq \alpha_\beta^{\frac{n-1}{2(\beta-1)} + \frac{1}{2}} \cdot \lambda_1(L).$$

II.2 Number Theory

II.2.1 Number fields and their geometry

We briefly introduce here the number theoretic objects we will use throughout this work. For an in-depth introduction to the field and the proofs of the mentioned results, the reader is referred, e.g., to [Coh93, Neu13].

A number field K is a finite field extension of \mathbb{Q} . Its degree is the dimension of the extension, its ring of integer is the integral closure of \mathbb{Z} in K , namely

$$\mathcal{O}_K = \{x \in K, \exists P \in \mathbb{Z}[X] \text{ monic}, P(x) = 0\}.$$

In the rest of this section, we let K be a number field of degree d , \mathcal{O}_K its ring of integers, $\Phi = (\sigma_i)_{1 \leq i \leq d} : K \rightarrow \mathbb{C}^d$ its canonical embedding. We order the σ_i so that $\sigma_i(K) \subset \mathbb{R}$ for $1 \leq i \leq d_{\mathbb{R}}$ (the real embeddings of K), and $\sigma_{d_{\mathbb{R}}+i} = \overline{\sigma_{d_{\mathbb{R}}+d_{\mathbb{C}}+i}}$ for any $1 \leq i \leq d_{\mathbb{C}}$ (the complex embeddings), in particular it holds that $d = d_{\mathbb{R}} + 2d_{\mathbb{C}}$. We let $\mathcal{O}_K^\times = \{x \in \mathcal{O}_K : \mathcal{N}(x) = 1\}$ denote the set of units of \mathcal{O}_K . The structure of the group of units of \mathcal{O}_K is given by the following theorem.

Theorem II.2.1 (Dirichlet's unit theorem). *If K is a number field of degree $d = d_{\mathbb{R}} + 2d_{\mathbb{C}}$, then*

$$\mathcal{O}_K^\times \simeq \mu_K \times \mathbb{Z}^{d_{\mathbb{R}}+d_{\mathbb{C}}-1},$$

where μ_K is the (finite) set of roots of unity of K .

Let $K_{\mathbb{R}} = K \otimes \mathbb{R}$; by abuse of notation, we shall also denote $\Phi = (\sigma_i)_i$ the extension of Φ to $K_{\mathbb{R}}$. The set $K_{\mathbb{R}}$ is a \mathbb{R} -algebra of dimension d containing K , and its image by Φ is the set

$$\Phi(K_{\mathbb{R}}) = \mathbb{R}^{d_{\mathbb{R}}} \times \{(y, \bar{y}), y \in \mathbb{C}^{d_{\mathbb{C}}}\}.$$

The function Φ is a ring homomorphism in the sense that for any $x, y \in K_{\mathbb{R}}$, $\Phi(x \cdot y) = \Phi(x) \odot \Phi(y)$ where \odot is the coordinate-wise multiplication. By abuse of notation, we will often identify $K_{\mathbb{R}}$, K and \mathcal{O}_K with their image by $\Phi(\cdot)$. For $x \in K_{\mathbb{R}}$, we define $\bar{x} \in K_{\mathbb{R}}$ as the element obtained by componentwise complex conjugation of the canonical embedding vector of x . We extend this notation to vectors and matrices over $K_{\mathbb{R}}$, and let \mathbf{x}^\dagger denote $\bar{\mathbf{x}}^T$ for any $\mathbf{x} \in K_{\mathbb{R}}^n$. We define \bar{K} and $\overline{\mathcal{O}_K}$ as the subsets of $K_{\mathbb{R}}$ obtained by applying complex conjugation to elements of K and \mathcal{O}_K , respectively.

The set $\Phi(\mathcal{O}_K)$ is a full-rank lattice in $K_{\mathbb{R}}$. The discriminant of K by $\Delta_K = \text{Vol}(\Phi(\mathcal{O}_K))^2$. The canonical embedding endows $K_{\mathbb{R}}$ with an Euclidean structure, where the norm of $x \in K_{\mathbb{R}}$ is $\|\Phi(x)\|$; by abuse of notation, we will write $\|x\| = \|\Phi(x)\|$ and $\|x\|_\infty = \|\Phi(x)\|_\infty$. We define the algebraic norm of any element of $K_{\mathbb{R}}$ as $\mathcal{N}(x) = \left| \prod_{i=1}^d \sigma_i(x) \right|$; this definition extends the algebraic norm to $K_{\mathbb{R}}$ in the sense that for $x \in K$, $\mathcal{N}(x) = |\mathcal{N}_{K/\mathbb{Q}}(x)|$. We let $K_{\mathbb{R}}^\times = \{x \in K_{\mathbb{R}}, \forall i, \sigma_i(x) \neq 0\}$ denote the set of invertible elements of $K_{\mathbb{R}}$, $K_{\mathbb{R}}^0$ the set of norm-1 elements of $K_{\mathbb{R}}$ and $K_{\mathbb{R}}^+$ the set of elements of $K_{\mathbb{R}}$ whose image by Φ lies in \mathbb{R}_+^d . We define the logarithmic embedding of $K_{\mathbb{R}}$, by taking the natural logarithm of every embedding of an element:

$$\begin{aligned} \text{Ln} : K_{\mathbb{R}}^\times &\longrightarrow \text{Ln}(K_{\mathbb{R}}) \subseteq \mathbb{R}^d \\ x &\longmapsto (\ln |\sigma_i(x)|)_{1 \leq i \leq d} \end{aligned}$$

Note that

$$\text{span}_{\mathbb{R}}(\text{Ln } \mathcal{O}_K^\times) = \text{Ln}(K_{\mathbb{R}}^0) := \{\mathbf{y} \in \mathbb{R}^d : \sum y_i = 0 \wedge \forall i \in [d_{\mathbb{C}}], y_{d_{\mathbb{R}}+d_{\mathbb{C}}+i} = y_{d_{\mathbb{R}}+i}\}.$$

For $\zeta \in \text{Ln}(K_{\mathbb{R}}^\times)$, we define $\text{Exp}(\zeta)$ as the element of $K_{\mathbb{R}}^+$ whose i -th embedding is $\exp(\zeta_i)$, for all i .

II.2.2 Ideals

In this subsection, we introduce the ideal arithmetic needed in this work.

Definition II.2.2. A fractional ideal I of K is a discrete subgroup of $(K, +)$ stable by multiplication by \mathcal{O}_K such that there exists $n \in \mathbb{Z} \setminus \{0\}$ such that $n \cdot I \subseteq \mathcal{O}_K$. A fractional ideal included in \mathcal{O}_K is called integral.

An (oriented) replete ideal is a subset of $K_{\mathbb{R}}$ of the form $I = x \cdot \mathfrak{a}$, where $x \in K_{\mathbb{R}}^\times$ and \mathfrak{a} is an integral ideal.

If a replete ideal can be written $I = x \cdot \mathcal{O}_K$, it is said to be principal.

Equivalently, an oriented replete ideal is a finitely generated \mathcal{O}_K -submodule of $K_{\mathbb{R}}$. Unless specified otherwise, by default, an ideal will refer to an oriented replete ideal.

In this work, we will take the convention that gothic letters (such as $\mathfrak{a}, \mathfrak{b}, \mathfrak{p}$) correspond to integral ideals, while upper-case letters (such as I, J) refer to ideals that are not necessarily integral.

For any replete ideals I, J , we define the product $I \cdot J$ as the ideal generated by all products $a \cdot b$ for $a \in I, b \in J$ and the inverse I^{-1} as the ideal $I^{-1} = \{x \in K_{\mathbb{R}}, x \cdot I \subseteq \mathcal{O}_K\}$. Note that if I is a fractional ideal, then so is I^{-1} . An integral ideal \mathfrak{p} is said to be prime if $\mathcal{O}_K/\mathfrak{p}$ is integral. This is equivalent (because \mathcal{O}_K is a Dedekind domain) to saying that there do not exist \mathfrak{a} and \mathfrak{b} integral and distinct from \mathfrak{p} such that $\mathfrak{p} = \mathfrak{a}\mathfrak{b}$. Similarly to the integers, every fractional ideal can be uniquely written as a product of prime ideals up to reordering.

We define the algebraic norm of an integral ideal \mathfrak{a} by $\mathcal{N}(\mathfrak{a}) = [\mathcal{O}_K : \mathfrak{a}]$. We have $\mathcal{N}(\mathfrak{a}\mathfrak{b}) = \mathcal{N}(\mathfrak{a})\mathcal{N}(\mathfrak{b})$ for all integral ideals \mathfrak{a} and \mathfrak{b} . If I is a replete ideal, there exists $x \in K_{\mathbb{R}}$ such that $x \cdot I$ is integral, and we define $\mathcal{N}(I) = \mathcal{N}(x \cdot I)/\mathcal{N}(x)$ (this is independent of the choice of x). The multiplicativity property of the norm carries over to replete ideals. For some ideal I of K , we define the ideal $\bar{I} = \{\bar{x} : x \in I\}$ of \bar{K} . For any real $2 \leq A \leq B$, we denote by $\mathcal{I}_{A,B}$ the set of integral ideal with norm in $[A, B]$ and by $\mathcal{P}_{A,B}$ the set of prime ideals with norm in $[A, B]$.

II.2.3 Embedding and ideal lattices.

Every non-zero replete ideal I corresponds to a full-rank lattice $\Phi(I)$. Such a lattice is called an ideal lattice (with respect to K). By abuse of notation, we will often identify I and $\Phi(I)$. It holds that $\text{Vol}(\Phi(I)) = \sqrt{\Delta_K} \cdot \mathcal{N}(I)$. We define IdLat_K^0 as the set of replete ideal lattices of norm 1.

Ideals lattices are not typical lattices in the sense that their geometry cannot be too skew. This is summarized in next lemma

Lemma II.2.3. Let I be an ideal lattice, then it holds that

$$\sqrt{d} \cdot \mathcal{N}(I)^{1/d} \leq \lambda_1(I) \leq \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d},$$

and

$$\lambda_d(I) \leq \sqrt{d} \cdot \lambda_d(\mathcal{O}_K) \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d},$$

furthermore, $\lambda_d(\mathcal{O}_K) \leq \sqrt{d} \cdot \Delta_K^{1/d}$.

Proof. The lower bound on λ_1 comes from the arithmetic-geometric inequality, the upper bound on λ_d from [BDPW20, Lemma 2.8]. The upper bound on $\lambda_d(\mathcal{O}_K)$ comes from [Boe22, Theorem A.4] and usual norm inequalities. \square

This particular geometry allows us to bound the covering radius in ℓ_∞ of ideal lattices:

Lemma II.2.4. *Let I be an ideal lattice, then it holds that*

$$\mu_\infty(I) \leq d \cdot \lambda_d^{(\infty)}(I) \leq d \cdot \lambda_1^{(\infty)}(I) \cdot \lambda_d^\infty(\mathcal{O}_K) \leq d \cdot \Delta_K^{3/(2d)} \cdot \mathcal{N}(I)^{1/d},$$

Proof. We bounded $\lambda_1(I)$ by $\Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$ using Minkowski's theorem and $\lambda_d^\infty(\mathcal{O}_K)$ by $\Delta_K^{1/d}$ using [BST⁺20, Theorem 3.1] (adapted to the ℓ_∞ norm in [Boe22, Theorem A.4]). \square

Their geometry also allows us to give bounds on $\eta_\varepsilon(I)$ depending only on the field and on $\mathcal{N}(I)$.

Lemma II.2.5 ([PRS17, Lemma 6.9]). *For any ideal I and $\varepsilon \in (0, 1)$, it holds that*

$$\eta_\varepsilon(I) \leq \Delta_K^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \max\left(1, \sqrt{\frac{\ln(1/\varepsilon)}{d}}\right). \quad (\text{II.2})$$

Corollary II.2.6. *Let I be a fractional ideal, \mathfrak{a} be an integral ideal and $\varepsilon \in (0, 1)$. Let $\mathbf{u} \in \text{span}(I)$ and $\varsigma \geq \Delta_K^{1/d} \cdot \mathcal{N}(\mathfrak{a} \cdot I)^{1/d} \cdot \sqrt{\ln(3/\varepsilon)}$. Then*

$$D_{I, \varsigma, \mathbf{u}}(\mathfrak{a} \cdot I) \in [1 - \varepsilon, 1 + \varepsilon] \cdot \mathcal{N}(\mathfrak{a})^{-1}.$$

Proof. Note that since \mathfrak{a} is integral, then $\mathfrak{a} \cdot I$ is a sub-lattice of I and $D_{I, \varsigma, \mathbf{u}}(\mathfrak{a} \cdot I)$ is well-defined. By Lemma II.2.5 and the lower bound on ς , we have $\varsigma \geq \eta_{\varepsilon'}(\mathfrak{a} \cdot I) \geq \eta_{\varepsilon'}(I)$, for $\varepsilon' = \varepsilon/3$. We can thus apply Lemma II.1.6 to both lattices I and $\mathfrak{a} \cdot I$. We obtain $D_{I, \varsigma, \mathbf{u}}(\mathfrak{a} \cdot I) = \rho_\varsigma(\mathfrak{a} \cdot I - \mathbf{u}) / \rho_\varsigma(I - \mathbf{u}) \in [(1 - \varepsilon')/(1 + \varepsilon'), (1 + \varepsilon')/(1 - \varepsilon')] \cdot \text{Vol}(I) / \text{Vol}(\mathfrak{a} \cdot I)$. We conclude using the fact that $\text{Vol}(\mathfrak{a} \cdot I) / \text{Vol}(I) = \mathcal{N}(\mathfrak{a})$ and the choice of ε' . \square

II.2.4 Riemann hypotheses

Riemann zeta function is defined over $\{s \in \mathbb{C}, \text{Re}(s) > 1\}$ by

$$\zeta(s) = \sum_{n \geq 1} \frac{1}{n^s},$$

and extended meromorphically to $\mathbb{C} \setminus \{1\}$ with a simple pole at $s = 1$. A lot of results have been proven under the following (unproven) hypothesis:

Definition II.2.7 (The Riemann Hypothesis (RH)). *The only zeros of ζ in the (so called) critical strip $\{s \in \mathbb{C}, \text{Re}(s) \in (0, 1)\}$ lie on the line $\{s \in \mathbb{C}, \text{Re}(s) = 1/2\}$.*

Even if not proven, the RH has been verified numerically for $|\text{Im}(s)| \leq 3 \cdot 10^{12}$ ([PT21]). A particular zeta function can be associated to numerous arithmetic objects. For a number field K , the Dedekind zeta function of K is defined over $\{s \in \mathbb{C}, \text{Re}(s) > 1\}$ by

$$\zeta_K(s) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^s},$$

where the sum is taken over the nonzero integral ideals and extended meromorphically to $\mathbb{C} \setminus \{1\}$ with a simple pole at $s = 1$. Note that $\zeta_{\mathbb{Q}}$ is the Riemann zeta function. As for ζ , the analytic properties of the function ζ_K encode many arithmetic properties of the field K . It is often convenient to assume that the following conjecture holds:

Definition II.2.8 (Riemann Hypothesis for ζ_K). *The only zeros of ζ_K in the critical strip lie on the line $\{s \in \mathbb{C}, \operatorname{Re}(s) = 1/2\}$.*

This Dedekind zeta function can be generalized further. Let χ be a Hecke character on K (a precise definition of Hecke character is beyond the scope of this manuscript. See, e.g., [Neu13, Definition 6.1] for an introduction), the Hecke L function associated to χ is defined over $\{s \in \mathbb{C}, \operatorname{Re}(s) > 1\}$ by

$$L(s, \chi) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{\chi(\mathfrak{a})}{\mathcal{N}(\mathfrak{a})^s},$$

where the sum is taken over the nonzero integral ideals and extended meromorphically to $\mathbb{C} \setminus \{1\}$ with a simple pole at $s = 1$. Note that if χ is the trivial character, we recover ζ_K .

Definition II.2.9 (Extended Riemann Hypothesis [Bac90] for the field K). *For any Hecke character χ on K , the only zeros of $L(\cdot, \chi)$ in the critical strip lie on the line $\{s \in \mathbb{C}, \operatorname{Re}(s) = 1/2\}$.*

The Extended Riemann Hypothesis (denoted by ERH in this manuscript) asserts that this holds for all number field K . Several of our results depend on ERH; in this case, only the Extended Riemann Hypothesis for the underlying field K is needed.

Residue at $s = 1$. We let ρ_K denote the residue at $s = 1$ of ζ_K . Then, we have:

$$\rho_K := \lim_{s \rightarrow 1} (s-1)\zeta_K(s) = \frac{2^{d_{\mathbb{R}}} \cdot (2\pi)^{d_{\mathbb{C}}} \cdot R_K \cdot |\operatorname{Cl}_K|}{|\mu_K| \cdot \sqrt{|\Delta_K|}}.$$

Where Cl_K is the class group of K (see Section II.2.5), μ_K is the set of its roots of unity and R_K is its regulator (a field invariant related to the volume of $\operatorname{Ln}(\mathcal{O}_K^{\times})^1$).

This residue ρ_K can be bounded depending on simple field invariants (discriminant, degree).

Theorem II.2.10 ([Lou00, Theorem 1]). *Let K be a number field of degree $d \geq 2$ and discriminant Δ_K . It holds that*

$$\rho_K \leq \left(\frac{e \cdot \log(\Delta_K)}{2(d-1)} \right)^{d-1}.$$

The quantity ρ_K is known to be $\operatorname{poly}(\log \Delta_K)$ for some families of number fields such as cyclotomic family (under ERH, see [Boe22, Theorem A.5]).

II.2.5 Class group and norm-1 ideals

The properties about multiplication and inverse give the set of replete ideals a group structure. The quotient of the group of replete ideals by the subgroup of principal replete ideals is called the class group of K and is denoted by Cl_K . It is a finite group, and for any replete ideal I , we denote its image in Cl_K by $[I] \in \operatorname{Cl}_K$. Two ideals I and J are in the same coset of Cl_K if and only if there exists $x \in K_{\mathbb{R}}^{\times}$ such that $I = x \cdot J$.

¹Not introduced in this work, see [Neu13, § 7.5] for a precise definition.

The set of norm-1 ideals IdLat_K^0 is a compact subgroup of the set of replete ideals, it therefore has a Haar measure $\mathcal{U}(\text{IdLat}_K^0)$. An algorithm to sample from this measure has been proposed in [BDPW20], and consist in performing a random walk in IdLat_K^0 . We present here a specialized version of this theorem where the walk has only one step. We will use algorithms from [BDPW20] to sample among different classes of ideals.

Lemma II.2.11 (Adapted from [BDPW20, Lemma 2.2], Assuming ERH). *There exists an absolute constant c and an algorithm `SamplePrimeIdeal` such that for any $B \geq (\log \Delta_K)^c$, the algorithm `SamplePrimeIdeal` runs in time $\text{poly}(\log B, d)$ on input B and returns a prime ideal uniform among prime ideals of norm $\leq B$.*

We will also rely on Algorithm II.2.1, which is adapted from [BDPW20, Theorem 3.3], to sample (essentially) uniformly in the set of norm-1 ideals lattices, in time polynomial in $\log B$. Note that [BDPW20] considers norm-1 ideals xI with I integral and all $\sigma_i(x)$'s being positive integers. This discrepancy is handled by introducing u at Step 3. The standard deviation in Step 2 and tail-cut may seem arbitrary at first sight: these choices simplify the analysis of the module randomization (in Section V.5.3). A proof of the following lemma is given in Appendix A.1.1.

Algorithm II.2.1 `Ideal-SampleB`

- 1: Let $\mathfrak{p} \leftarrow \text{SamplePrimeIdeal}(B)$ (see Lemma II.2.11) an uniform prime ideals of norm $\leq B$;
 - 2: Sample $\zeta \in \text{Ln}(K_{\mathbb{R}}^{\times})$ from the centered normal law with standard deviation $d^{-3/2}$, conditioned on $\|\zeta\| \leq 1/d$;
 - 3: Sample u uniform in $\{x \in K_{\mathbb{R}}, \forall i \in [d] : |\sigma_i(x)| = 1\}$;
 - 4: Return $u \cdot \text{Exp}(\zeta) \cdot \mathfrak{p} / \mathcal{N}^{1/d}(\mathfrak{p})$.
-

Lemma II.2.12 (Adapted from [BDPW20, Theorem 3.3], Assuming ERH). *There exists an absolute constant c such that for any $B \geq (d^d \Delta_K)^c$, `Ideal-SampleB` runs in time polynomial in $\log B$ and its output distribution is within $2^{-\Omega(d)}$ statistical distance from $\mathcal{U}(\text{IdLat}_K^0)$.*

Lemma II.2.13. *Let J be a replete ideal, then*

$$\Pr_{I \leftarrow \mathcal{U}(\text{IdLat}_K^0)} \left(\exists x \in K_{\mathbb{R}}^{\times} : J = I \cdot (x) \right) = \frac{1}{|\text{Cl}_K|}$$

Proof. For any replete ideal $I = (x) \cdot \mathfrak{a}$ with \mathfrak{a} integral, we recall that $[I] = [\mathfrak{a}] \in \text{Cl}_K$, which does not depend on the choices of \mathfrak{a} and x . The function $I \mapsto [I]$ for $I \in \text{IdLat}_K^0$ is a surjective morphism whose kernel is the set of principal replete ideals of norm 1 in $K_{\mathbb{R}}$. The lemma states that if I is sampled from $\mathcal{U}(\text{IdLat}_K^0)$, then the probability that it belongs to a fixed coset of Cl_K is $|\text{Cl}_K|^{-1}$, which follows directly from the fact that $[\cdot]$ is a surjective morphism. \square

II.2.6 Computations with number theoretic objects

Representing field elements and ideals. In all this work, when working with a number field K , we assume that we know a \mathbb{Z} -basis $\mathbf{B}_{\mathcal{O}_K} = [b_1^{\mathcal{O}_K}, \dots, b_d^{\mathcal{O}_K}]$ of \mathcal{O}_K , and that it is LLL-reduced with respect to the geometry induced by Φ (in some cases, a much better basis could be known). We define $\delta_K := \|\mathbf{B}_{\mathcal{O}_K}\|$. Since $\mathbf{B}_{\mathcal{O}_K}$ is LLL-reduced, we have that $\delta_K \leq 2^d \cdot \lambda_d(\mathcal{O}_K) \leq \sqrt{d} \cdot 2^d \cdot \Delta_K^{1/d}$ by Lemma II.2.3, which implies that $\log \delta_K = O(\log \Delta_K)$.

Elements of K will be represented as vectors of \mathbb{Q}^d , corresponding to their coordinates in the basis $\mathbf{B}_{\mathcal{O}_K}$. Fractional ideals of K will be represented by a \mathbb{Z} -basis, i.e., d elements of K

generating the ideal (each element being represented as a vector of \mathbb{Q}^d as described above). The bases we obtain for a fractional ideal I are in $\mathbb{Q}^{d \times d}$, so they admit a Hermite Normal Form (HNF), which provides a canonical representation for I . When replete ideals are used in algorithms, they will be represented by an arbitrary basis with size polynomial in the log of their norm and in $\log \Delta_K$ (with a polynomial number of bits of precision).

Ideal Arithmetic. We will often manipulate ideals and their bases. We will use the following results on how to derive a short basis from a full-rank set of vectors.

Lemma II.2.14 (Corollary of [MG02, Lemma 7.1]). *There exists a polynomial time algorithm that takes as input a basis \mathbf{B} of an n -dimensional lattice L and a set of n linearly independent vectors $\mathbf{s}_1, \dots, \mathbf{s}_n \in L$ and outputs a new basis \mathbf{C} of L such that $\|\mathbf{C}^*\| \leq \max_i \|\mathbf{s}_i^*\|$ and $\|\mathbf{C}\| \leq \sqrt{n} \cdot \max_i \|\mathbf{s}_i\|$.*

We will use Lemma II.2.14 to perform arithmetic over ideals while bounding the sizes of the outputs.

Lemma II.2.15. *There exist polynomial-time algorithms performing ideal inversion, reduction and multiplications `InvertIdeal`, `ReduceIdeal` and `MultiplyIdeals` with the following specifications.*

- `InvertIdeal` takes as input an integral ideal \mathfrak{a} and outputs a basis \mathbf{B} of \mathfrak{a}^{-1} satisfying $\|\mathbf{B}^*\| \leq \delta_K$ and $\|\mathbf{B}\| \leq \sqrt{d} \cdot \delta_K$.
- `ReduceIdeal` takes as input a basis \mathbf{B} of an ideal $I \subset K_{\mathbb{R}}$ and a vector $v \in I \setminus \{0\}$ and returns a basis \mathbf{B}_I of I such that $\|\mathbf{B}_I^*\| \leq \delta_K \cdot \|v\|$ and $\|\mathbf{B}_I\| \leq \sqrt{d} \cdot \delta_K \cdot \|v\|$.
- `MultiplyIdeals` takes as input bases \mathbf{B}_I and \mathbf{B}_J of two ideals $I, J \subseteq K_{\mathbb{R}}$ and output \mathbf{B}_{IJ} a basis of $I \cdot J$ such that $\|\mathbf{B}_{IJ}^*\| \leq \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$ and $\|\mathbf{B}_{IJ}\| \leq \sqrt{d} \cdot \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$.

Proof. `InvertIdeal` starts by computing a basis \mathbf{B} of \mathfrak{a}^{-1} , which can be done in polynomial time from a representation of \mathfrak{a} by generators. Then, the algorithm runs the algorithm from Lemma II.2.14 with input the basis \mathbf{B} of \mathfrak{a}^{-1} and the vectors of the known basis $\mathbf{B}_{\mathcal{O}_K}$ of \mathcal{O}_K (in the role of the short vectors \mathbf{s}_i). Note that since \mathfrak{a} is integral, we have that $\mathcal{O}_K \subseteq \mathfrak{a}^{-1}$, and hence the vectors of $\mathbf{B}_{\mathcal{O}_K}$ are indeed in \mathfrak{a}^{-1} . Also, the euclidean norm of those vectors is bounded from above by δ_K , by definition. We conclude by using Lemma II.2.14.

For `ReduceIdeal`, note that the set $v \cdot \mathbf{B}_{\mathcal{O}_K}$ is a free subset of I whose vectors have norms $\leq \|v\| \cdot \delta_K$. We can then define `ReduceIdeal` as the application of Lemma II.2.14 with input $\mathbf{B}, v \cdot \mathbf{B}_{\mathcal{O}_K}$.

Let $\mathbf{B}_I = (b_i^{(I)})_i, \mathbf{B}_J = (b_i^{(J)})_i$ be the inputs to `MultiplyIdeals`. Then the set $(b_i^{(I)} \cdot b_j^{(J)})_{i,j}$ generates IJ and has size d^2 , this implies that there exists a \mathbb{Q} -free family $(r_i)_{i=1, \dots, d}$ inside it, which can be found in polynomial time and satisfies $\max_i \|r_i\| \leq \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$. Further, a basis \mathbf{B} of IJ can be found in polynomial time. We then apply Lemma II.2.14 with input $\mathbf{B}, (r_i)_i$. \square

II.3 Modules

II.3.1 General definitions

For $\mathbf{x}, \mathbf{y} \in K_{\mathbb{R}}^n$, we define $\mathbf{x}^\dagger = (\overline{x_1}, \dots, \overline{x_n})^T$, the scalar product in $\langle \mathbf{x}, \mathbf{y} \rangle_{K_{\mathbb{R}}} = \mathbf{x}^\dagger \cdot \mathbf{y} \in K_{\mathbb{R}}$, $\|\mathbf{x}\|_{K_{\mathbb{R}}} = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle_{K_{\mathbb{R}}}} \in K_{\mathbb{R}}^+$ where the square root is taken component-wise, and $\|\mathbf{x}\| = \|\Phi(\|\mathbf{x}\|_{K_{\mathbb{R}}})\|$. The algebraic norm of $\mathbf{x} \in K_{\mathbb{R}}^n$ is defined as $\mathcal{N}(\mathbf{x}) = \mathcal{N}(\|\mathbf{x}\|_{K_{\mathbb{R}}})$. In analogy with

lattices, which are \mathbb{Z} -modules, we define module-lattices, which are \mathcal{O}_K -modules (though we will often just call them modules when the context is clear). The \mathcal{O}_K -modules we will consider throughout this work will always be finitely generated and torsion-free.

Definition II.3.1 (\mathcal{O}_K -module). *Let $1 \leq k \leq m$. In this work, a (finitely generated and torsion-free) module-lattice is a subset of $K_{\mathbb{R}}^m$ of the form $M = \sum_{i \leq k} \mathbf{b}_i I_i$ where the I_i 's are non-zero ideals and the \mathbf{b}_i 's are $K_{\mathbb{R}}$ -linearly independent. This is written compactly as $M = \mathbf{B} \cdot \mathbb{I}$ (where \mathbf{B} is the matrix whose columns are the \mathbf{b}_i and $\mathbb{I} = (I_1, \dots, I_k)$).*

One can see that if $K = \mathbb{Q}$, this is the definition of lattices, and that if $k = m = 1$ a module is simply an ideal. During all this manuscript, when talking about modules, it will refer to modules-lattices (so, finitely generated and torsion-free \mathcal{O}_K -modules).

The tuple $((I_1, \mathbf{b}_1), \dots, (I_k, \mathbf{b}_k))$ is called a pseudo-basis of M and is written compactly as (\mathbf{B}, \mathbb{I}) . Two pseudo-bases (\mathbf{B}, \mathbb{I}) and $(\mathbf{B}', \mathbb{I}')$ define the same module if and only if there exists a matrix $\mathbf{U} = (u_{ij})_{1 \leq i, j \leq k} \in \text{GL}_k(K_{\mathbb{R}})$ with $u_{ij} \in I_i^{-1} \cdot J_j$ such that $\mathbf{B} \cdot \mathbf{U} = \mathbf{B}'$. The integer k is the rank of M .

We define the norm of a module M as $\mathcal{N}(M) = \sqrt{\det(\mathbf{B}^\dagger \cdot \mathbf{B})} \cdot \prod_{i \leq k} \mathcal{N}(I_i)$. Note that for $k = m = 1$, this matches the norm of an ideal. Using the canonical embedding, any rank- k module-lattice is identified to a (kd) -dimensional lattice. In particular, we define $\det(M)$ as the determinant of the module lattice. Note that $\det(M) = \mathcal{N}(M) \cdot \Delta_K^{k/2}$. We will be interested in the module norm-minimum $\lambda_1^{\mathcal{N}}(M) = \inf\{\mathcal{N}(N) : N \text{ rank-1 submodule of } M\}$. A rank-1 submodule of M is said *densest* if it reaches $\lambda_1^{\mathcal{N}}(M)$.

The dual of a module is defined in the same way as the dual of any lattice L , but with base-ring \mathcal{O}_K :

Definition II.3.2. *The dual of a module M is defined as*

$$M^\vee = \{\mathbf{b}^\vee \in \text{span}_{K_{\mathbb{R}}}(M) : \forall \mathbf{b} \in M, \langle \mathbf{b}^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \in \mathcal{O}_K\}.$$

Note that M^\vee is an $\overline{\mathcal{O}_K}$ -module, $\Phi(M^\vee)$ is the dual lattice of $\Phi(M)$ and $(\mathbf{B} \cdot \mathbb{I})^\vee = (\mathbf{B}^{-\dagger} \cdot \mathbb{J})$, where $J_i = (\overline{I_i})^{-1}$ for all $i \leq k$.

The Hermite Normal Form can be generalized to modules over \mathcal{O}_K (because it is a Dedekind domain). For any full-rank torsion-free module $M \subseteq K^m$, there exists a pseudo-basis (\mathbf{B}, \mathbb{I}) such that $\mathbf{B} \in K^{m \times m}$ is lower-triangular with ones on the diagonal. It is called a Hermite Normal Form of M and can be computed in polynomial time from any finite set of pairs $\{(I_i, \mathbf{b}_i)\}_i$ such that $M = \sum_i \mathbf{b}_i I_i$ and the \mathbf{b}_i 's are not necessarily independent (see, e.g., [Coh00, §1.4.2]).

Definition II.3.3. *Let M be a module. A submodule $N \subseteq M$ is said to be primitive if it satisfies any of the three equivalent conditions:*

- *the module N is maximal for the inclusion in the set of submodules of M of rank at most $\text{rank}(N)$;*
- *there is a module N' with $M = N + N'$ and $\text{rank}(M) = \text{rank}(N) + \text{rank}(N')$;*
- *we have $N = M \cap \text{span}_K(N)$.*

In particular, any densest rank-1 submodule of M is primitive.

A proof that the three conditions are equivalent is provided in Appendix A.1.2. The last statement follows from Condition 1.

II.3.2 Rank-2 Modules with a Gap

In this work, we will be interested in rank-2 modules that contain an unexpectedly dense rank-1 submodule, i.e., in modules M with $\lambda_1^{\mathcal{N}}(M)$ significantly smaller than $\sqrt{\mathcal{N}(M)}$. We define the gap of M by

$$\gamma(M) = \left(\frac{\mathcal{N}(M)^{\frac{1}{2}}}{\lambda_1^{\mathcal{N}}(M)} \right)^{\frac{1}{d}}.$$

The following lemma shows that if the gap is sufficiently large, then the densest rank-1 submodule (meaning the rank-1 submodule with the lowest algebraic norm) is unique.

Lemma II.3.4. *Let M be a rank-2 module with gap $\gamma > 0$ and N a densest rank-1 submodule of M . If N' is a rank-1 submodule of M with $\mathcal{N}(N') < \gamma^d \sqrt{\mathcal{N}(M)}$, then $N' \subseteq N$.*

In particular, for $\gamma > 1$, the densest rank-1 submodule is unique and any vector $\mathbf{b} \in M$ with $\|\mathbf{b}\| < \gamma \cdot \mathcal{N}(M)^{1/(2d)}$ belongs to it.

Proof. See Appendix A.1.3. □

In the following, when a rank-2 module M has a gap larger than 1, we will implicitly use Lemma II.3.4 when referring to the densest rank-1 submodule of M . Most rank-2 modules we will consider will have gap larger than 1.

The latter lemma allows us to conclude that the module norm-minimum is reached (see Appendix A.1.4 for a proof).

Lemma II.3.5. *For any module M , there exists a rank-1 submodule N of M such that $\mathcal{N}(N) = \lambda_1^{\mathcal{N}}(M)$.*

Proof. See Appendix A.1.4. □

II.4 Computational Problems

In this section, we introduce the main computational problems over lattices. The complexity of the algorithm on unstructured lattices will be a function of the dimension n of the lattice and the log of the norm of their basis. Let $\gamma \geq 1$ be an approximation factor. All the lattices and modules in this section are supposed full rank and represented by their HNF basis.

II.4.1 The Shortest Vector Problem

Definition II.4.1 (Variants of SVP). *The Shortest Vector Problem SVP_γ asks, given as input a lattice $L \subset \mathbb{Q}^n$, to find a non-zero element $\mathbf{x} \in L$ such that $\|\mathbf{x}\| \leq \gamma \cdot \lambda_1(L)$.*

The Hermite Shortest Vector Problem HSVP_γ asks, given as input a lattice $L \subset \mathbb{Q}^n$, to find a non-zero element $\mathbf{x} \in L$ such that $\|\mathbf{x}\| \leq \gamma \cdot \text{Vol}(L)^{1/n}$.

The Short Independent Vectors Problem SIVP_γ asks, given as input a lattice $L \subset \mathbb{Q}^n$, to find a free set $\mathbf{v}_1, \dots, \mathbf{v}_n \in L$ such that $\|\mathbf{v}_i\| \leq \gamma \cdot \lambda_n(L)$ for $1 \leq i \leq n$.

II.4.2 Structured lattice problems

We now introduce short vector problems for structured lattices, as well as related algorithmic problems.

Definition II.4.2 (Variants of id-HSVP). *The ideal Hermite Shortest Vector Problem id-HSVP $_\gamma$ asks, given as input a fractional ideal I represented by its HNF basis, to find a non-zero element $x \in I$ such that $\|x\| \leq \gamma \cdot \text{Vol}(I)^{1/d}$.*

For a finite set X of fractional ideals, the average-case variant X -avg-id-HSVP $_\gamma$ asks to solve the problem id-HSVP $_\gamma$ when the input ideal I is uniformly sampled in X . The success probability of a probabilistic algorithm \mathcal{A} when solving X -avg-id-HSVP $_\gamma$ is defined as

$$\Pr_{I \leftarrow X} [x \in I \text{ and } \|x\| \leq \gamma \cdot \text{Vol}(I)^{1/d} \mid \mathcal{A}(I) = x],$$

where the probability is taken over the choice of I and the randomness used by \mathcal{A} .

The problem inv-HSVP $_\gamma$ is id-HSVP $_\gamma$ restricted to inverses of integral ideal lattices.

We now define different variants of the unique-SVP problem for rank-2 modules, as well as variants of the NTRU problem.

Definition II.4.3 (γ -mod-uSVP $_2$ instance). *Let $\gamma > 0$. A γ -mod-uSVP $_2$ instance consists in a pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 module $M \subset K^2$ such that M contains a non-zero vector \mathbf{s} with $\|\mathbf{s}\| \leq \gamma^{-1} \cdot \mathcal{N}(M)^{1/(2d)}$.*

Note that any module M associated to a γ -mod-uSVP $_2$ instance contains the rank-1 submodule \mathfrak{sO}_K whose norm is $\leq \sqrt{\mathcal{N}(M)}/\gamma^d$. By Lemma II.3.4, this implies that if $\gamma > 1$, then the module M has a unique densest rank-1 submodule.

Definition II.4.4 ($(\mathcal{D}, \gamma, \gamma')$ -mod-uSVP $_2^{\text{vec}}$ and (γ, γ') -wc-mod-uSVP $_2^{\text{vec}}$). *Let $\gamma \geq \gamma' > 0$ and \mathcal{D} be a distribution over γ -mod-uSVP $_2$ instances. The $(\mathcal{D}, \gamma, \gamma')$ average-case unique SVP problem for rank-2 modules ($(\mathcal{D}, \gamma, \gamma')$ -mod-uSVP $_2^{\text{vec}}$ for short) asks, given as input a pseudo-basis of some rank-2 module M sampled from \mathcal{D} , to compute a vector $\mathbf{s} \in M \setminus \{\mathbf{0}\}$ such that $\|\mathbf{s}\| \leq \mathcal{N}(M)^{1/(2d)}/\gamma'$. The advantage of a probabilistic algorithm \mathcal{A} against the $(\mathcal{D}, \gamma, \gamma')$ -mod-uSVP $_2^{\text{vec}}$ problem is defined as*

$$\text{Adv}(\mathcal{A}) = \Pr_{(\mathbf{B}, \mathbb{I}) \leftarrow \mathcal{D}} \left(\mathcal{A}((\mathbf{B}, \mathbb{I})) = \mathbf{s} \text{ with } \left| \begin{array}{l} \mathbf{s} \in M \setminus \{\mathbf{0}\} \\ \|\mathbf{s}\| \leq \mathcal{N}(M)^{1/(2d)}/\gamma' \end{array} \right. \right),$$

where the probability is also taken over the internal randomness of \mathcal{A} .

The variant (γ, γ') -wc-mod-uSVP $_2^{\text{vec}}$ asks to solve this problem for any γ -mod-uSVP $_2$ instance (\mathbf{B}, \mathbb{I}) .

Definition II.4.5 ((\mathcal{D}, γ) -mod-uSVP $_2^{\text{mod}}$ and γ -wc-mod-uSVP $_2^{\text{mod}}$). *Let $\gamma > 0$ and \mathcal{D} be a distribution over γ -mod-uSVP $_2$ instances. The (\mathcal{D}, γ) unique SVP problem for rank-2 modules ((\mathcal{D}, γ) -mod-uSVP $_2^{\text{mod}}$ for short) asks, given as input a γ -mod-uSVP $_2$ module M sampled from \mathcal{D} , to recover a densest rank-1 submodule $N \subset M$. The advantage of a probabilistic algorithm \mathcal{A} against the (\mathcal{D}, γ) -mod-uSVP $_2^{\text{mod}}$ problem is defined as*

$$\text{Adv}(\mathcal{A}) = \Pr_{(\mathbf{B}, \mathbb{I}) \leftarrow \mathcal{D}} \left(\mathcal{A}((\mathbf{B}, \mathbb{I})) = N \text{ with } \left| \begin{array}{l} N \subset M \text{ with } \text{rk}(N) = 1 \\ \mathcal{N}(N) = \lambda_1^{\mathcal{N}}(M) \end{array} \right. \right),$$

where the probability is also taken over the internal randomness of \mathcal{A} .

The worst-case variant $(\gamma$ -wc-mod-uSVP $_2^{\text{mod}}$) asks to solve this problem for any γ -mod-uSVP $_2$ instance (\mathbf{B}, \mathbb{I}) .

We can now define the NTRU problems, as special cases of the mod-uSVP $_2$ variants above.

Definition II.4.6 (NTRU instance). *Let $q \geq 2$ be an integer, and $\gamma > 0$ a real number. A (γ, q) -NTRU instance is a γ -mod-uSVP $_2$ instance whose pseudo-basis is required to be of the form $((\mathbf{b}_1, \mathcal{O}_K), (\mathbf{b}_2, \mathcal{O}_K))$ with $\mathbf{b}_1 = (1, h)^T$ for some $h \in \mathcal{O}_K$ and $\mathbf{b}_2 = (0, q)^T$.*

Comparison with [PS21]. In order to emphasize the similarities between mod-uSVP_2 and NTRU, we adopted slightly different definitions from [PS21] for NTRU problems. The problems corresponding to those definitions can easily be reduced to one another; we thus elected to keep the same problem names. In [PS21], an NTRU instance consists in the single element $h \in R_q$, whereas we consider it as a basis of a rank-2 module in this work. Both formalisms are equivalent, since one can reconstruct the basis of the rank-2 module from h (and also q , which is a parameter of the problem). A second difference comes from the fact that [PS21] requires the short vector $\mathbf{s} = (s_1, s_2)^T$ to satisfy $\|s_1\|, \|s_2\| \leq \sqrt{q}/\gamma$, whereas we require that $\|\mathbf{s}\| \leq \sqrt{q}/\gamma$. This means that a (γ, q) -NTRU instance for us is a (γ, q) -NTRU instance for [PS21], but the converse does not hold: a (γ, q) -NTRU instance for [PS21] is only guaranteed to be a $(\gamma/\sqrt{2}, q)$ -NTRU instance for us.

Definition II.4.7 (NTRU problems). *Let $q \geq 2$, $\gamma \geq \gamma' > 0$ and \mathcal{D} a distribution over (γ, q) -NTRU instances. The $(\mathcal{D}, \gamma, \gamma', q)$ -NTRU^{vec} problem, (γ, γ', q) -wc-NTRU^{vec} problem, (\mathcal{D}, γ, q) -NTRU^{mod} problem and (γ, q) -wc-NTRU^{mod} problem are the restrictions of the mod-uSVP_2 problems to (γ, q) -NTRU instances.*

From the definitions of the NTRU and mod-uSVP_2 problems, one can see that the average case $\text{mod-uSVP}_2^{\text{vec}}$ and NTRU^{mod} problems reduce to $\text{wc-mod-uSVP}_2^{\text{vec}}$ and $\text{wc-mod-uSVP}_2^{\text{mod}}$. In fact, we will show that the converse also holds, provided we have an oracle solving ideal-SVP.

Finally, we also recall the definition of the Hermite shortest vector problem in ideal lattices (id-HSVP).

Definition II.4.8 (γ -id-HSVP). *Let $\gamma \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$. Given as input a fractional ideal $I \subset K$, the γ -id-HSVP problem asks to find an element $x \in I \setminus \{0\}$ such that $\|x\| \leq \gamma \cdot \mathcal{N}(I)^{1/d}$.*

By Lemma II.2.3, this problem is well-defined for any $\gamma \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$. The problem id-HSVP is equivalent (up to a small change in the approximation factor) to SVP restricted to ideal lattices (which is denoted id-SVP).

Lemma II.4.9. *Let $\gamma \geq 1$, then $\text{id-HSVP}_{\gamma'}$ reduces to id-SVP_{γ} for $\gamma' = \gamma \cdot \sqrt{d} \cdot \Delta_K^{1/(2d)}$.*

Let $\gamma \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, then $\text{id-SVP}_{\gamma'}$ reduces to id-HSVP_{γ} for $\gamma' = \gamma/\sqrt{d}$.

Proof. This is a consequence of Lemma II.2.3. If x is a id-SVP_{γ} solution for an ideal I , then it holds that

$$\|x\| \leq \gamma \cdot \lambda_1(I) \leq \gamma \cdot \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}.$$

If x is now a id-HSVP_{γ} solution for an ideal I , then it holds that

$$\|x\| \leq \gamma \cdot \mathcal{N}(I)^{1/d} \leq \frac{\gamma \cdot \lambda_1(I)}{\sqrt{d}},$$

which concludes the proof. \square

II.4.3 Complexity parameters of algorithms over K

As said, in the unstructured case, the asymptotic complexity of problems are given as functions of the dimension n of the lattice and of the bit-size of the basis. In contrast, in the structured case, we shall study the asymptotic complexity of problems using as parameters the degree d and the root-discriminant of the underlying number field rather than the rank of the modules. The parameters d and Δ_K are related by Minkowski's bound.

$$\left(\frac{\pi}{4}\right)^{\frac{d}{2}} \cdot \frac{d^d}{d!} \leq \sqrt{\Delta_K},$$

which in particular implies that $\log(\Delta_K) = \Omega(d)$. In actual cryptographic constructions, the number field used usually satisfy $\log(\Delta_K) = \tilde{O}(d)$ (e.g. cyclotomics). This implies that we will consider acceptable to have algorithm running-time/approximation factors with a polynomial bound on $\Delta_K^{1/d}$.

However, fields of fixed degree with arbitrarily large Δ_K do exist, the simplest example being the field $\mathbb{Q}(\sqrt{D})$ for large square free $D \in \mathbb{Z}$. On the contrary, there exists a family of fields $(K_n)_{n \geq 1}$ such that $\deg(K_n)$ goes to infinity and $\Delta_{K_n}^{1/(\deg(K_n))}$ is constant. These fields are given by constructing number field with an infinite number of unramified extensions. A construction of such field can be found in [Mai00], but it should be noted that none of these field is actually used in actual cryptographic construction to our knowledge.

The most widely used field for constructions is the power-of-two cyclotomic fields $K_n = \mathbb{Q}(\zeta_{2^n})$, for any $n \geq 1$, where ζ_{2^n} is a 2^n -th root of unity. It holds that for any $n \geq 1$, $\deg(K_n) = 2^{n-1}$, $\mathcal{O}_K = \mathbb{Z}[\zeta_{2^n}]$ and

$$\Delta_{K_n} = \pm 2^{(n-1) \cdot 2^{n-1}} = \pm \deg(K_n)^{\deg(K_n)}.$$

II.5 Probabilities

In this manuscript, we will always work with probability distributions which are discrete or absolutely continuous with respect to the Lebesgue measure over \mathbb{R}^n for some n . When writing

$$\int_{t \in \text{Supp}(D)} f(D(t)) dt,$$

for D a probability distribution and some function f , we will mean it either as a sum over $\text{Supp}(D)$ (if D is discrete), or as an integral over \mathbb{R}^n .

Let X be a set which is finite or has finite Lebesgue measure. Throughout this manuscript we will denote by $\mathcal{U}(X)$ the uniform distribution over X .

In order to measure the differences between two probability distributions, we shall use, depending on the circumstances, both the statistical distance and the Rényi divergence.

Definition II.5.1. *Let D_1 and D_2 be two probability distributions defined over the same σ -algebra. The statistical distance between D_1 and D_2 is defined by:*

$$\text{SD}(D_1, D_2) = \frac{1}{2} \int_{t \in \text{Supp}(D_1) \cup \text{Supp}(D_2)} |D_1(t) - D_2(t)| dt.$$

Assume that $\text{Supp}(D_1) \subseteq \text{Supp}(D_2)$. The Rényi divergence of order 2 is

$$\text{RD}_2(D_1 \parallel D_2) = \int_{t \in \text{Supp}(D_1)} \frac{D_1(t)^2}{D_2(t)} dt,$$

and the Rényi divergence of infinite order is

$$\text{RD}_\infty(D_1 \parallel D_2) = \max_{x \in \text{Supp}(D_1)} D_1(x)/D_2(x).$$

These quantities allow to compare the respective probability of events sampled from different distributions.

Lemma II.5.2 (Data processing inequality for SD and RD). *Let D_1, D_2 be two probability distributions. For any event $E \subseteq \text{Supp}(D_1) \cup \text{Supp}(D_2)$, we have*

$$D_2(E) \geq D_1(E) - \text{SD}(D_1, D_2).$$

If $\text{Supp}(D_1) \subseteq \text{Supp}(D_2)$, then for any event $E \subseteq \text{Supp}(D_1)$, we have

$$D_2(E) \geq \frac{D_1(E)^2}{\text{RD}_2(D_1 \parallel D_2)} \quad \text{and} \quad D_2(E) \geq \frac{D_1(E)}{\text{RD}_\infty(D_1 \parallel D_2)}.$$

Chapter III

Counting Small Ideals

During our work on module lattices, we needed multiple times to have precise estimates on $N_K(B)$, the number of integral ideals whose norm is less or equal than $B \geq 0$. It is known that when B goes to infinity $N_K(B) \sim \rho_K \cdot B$, but the error term in this previous equivalence heavily depends on K . The bounds we found on the literature were unusable for values of B that we used in our work (e.g., $B = \text{poly}(d, \Delta_K^{1/(2d)})^d$) or had their dependence on the field K not made explicit.

In this chapter, we give new (to the extent of our knowledge) result about the size of the error on the approximation $N_K(B) \sim \rho_K \cdot B$ for a number field K of degree $d \geq 3$. In particular, we are explicit about the error's dependence on the field invariants.

III.1 Preliminaries

In this chapter we manipulate functions depending on an underlying number field K . We emphasize that when writing $f \ll g$ or $f = O(g)$, the implicit constant in the \ll and $O(\cdot)$ symbol does not depend on the number field.

For any $B \geq 0$, we let $N_K(B)$ (respectively $\pi_K(B)$) denote the number of integral ideals (respectively prime ideals) with algebraic norm $\leq B$. As in the rational case, it is possible to give estimate on these quantities. The RH for ζ_K allows us to give tighter bounds. The easiest to estimate is $\pi_K(B)$, which is $\sim B/\log(B)$ with explicit bounds on the error.

Lemma III.1.1 ([BS96, Theorem 8.7.4], Assuming ERH). *There exists an absolute constant c_1 such that for any $B \geq (\log \Delta_K)^{c_1}$, we have*

$$\pi_K(B) \in \frac{B}{\log B} \cdot [0.9, 1.1].$$

It seems that $N_K(B)$ is trickier to estimate. It is known that it is equivalent to $\rho_K \cdot B$ when B goes to infinity, but the error term depends on the field K in a non-trivial manner. We prove in this chapter the following two results whose proof can be found in Section III.4 and which are, to the extent our knowledge, new contributions.

Theorem III.1.2 (Assuming ERH). *Let K be a number field of degree $d \geq 3$ and $N_K(\cdot)$ be its ideal-counting function. For any $X \geq 2^d$, it holds that*

$$|N_K(X) - \rho_K X| \leq M(K) \cdot X^{1-\eta},$$

for $\eta = 1/(16 \ln(d))$ and some $M(K)$ satisfying

$$\ln(M(K)) \ll \ln(|\Delta_K|) + d \ln(d).$$

Corollary III.1.3 (Assuming ERH). *Let K be a number field of degree $d \geq 3$ and $N_K(\cdot)$ be its ideal-counting function. There exists a field independent constant $c_2 > 1$ such that for any $C \geq 1$, if*

$$X \geq (C \cdot d^d \cdot |\Delta_K|)^{c_2 \cdot \ln(d)}$$

then it holds that

$$\frac{|N_K(X) - \rho_K \cdot X|}{\rho_K \cdot X} \leq \frac{1}{C}.$$

III.1.1 Analysis preliminaries

Lemma III.1.4. *Let $d \geq 1$ be an integer and $\alpha \in (0, 1)$ a real number. Then, for any $x \geq e$, we have*

$$\frac{(\ln \ln x)^d}{x^\alpha} \leq \left(\ln \left(1 + \frac{d}{\alpha} \right) \right)^d.$$

Proof. See Appendix B.1.1. □

Lemma III.1.5 ([Ten95, Lemma II.2.1.1]). *For any $T, \kappa > 0$ and $x \in \mathbb{R}_{>0} \setminus \{1\}$,*

$$\left| \frac{1}{2i\pi} \int_{\kappa+iT}^{\kappa-iT} \frac{x^s}{s} ds - \mathbf{1}_{>1}(x) \right| \leq \frac{x^\kappa}{\pi \cdot T \cdot |\ln(x)|},$$

where $\mathbf{1}_{>1}$ is the indicator function of the set $\{x \in \mathbb{R}, x > 1\}$.

III.1.2 Analytic number theory preliminaries

Let K be a number field of degree $d = d_{\mathbb{R}} + 2d_{\mathbb{C}} \geq 3$ and of discriminant Δ_K . Recall that $|\Delta_K| \geq 1$ and $d \ll \ln(|\Delta_K|)$. For any $X \geq 1$, we recall the ideal-counting function

$$N_K(X) = |\{\mathfrak{a} \text{ integral ideal of } K, \mathcal{N}(\mathfrak{a}) \leq X\}|.$$

Recall that ζ_K denotes the Dedekind zeta function associated to K . Note that $\zeta_K(\bar{s}) = \overline{\zeta_K(s)}$. We first recall classical results concerning ζ_K .

Lemma III.1.6. *The residue ρ_K satisfies*

$$\rho_K \geq 0.5 \cdot |\Delta_K|^{-1/2}.$$

Proof. See Appendix B.2.1. □

The following theorem follows from [Rad59, Theorem 4] with $\eta = 1/2$.

Theorem III.1.7. *For any $s = \sigma + it \in \mathbb{C}$ with $\sigma \in (-1/2, 3/2)$, we have*

$$|\zeta_K(s)| \leq 3 \cdot \zeta \left(\frac{3}{2} \right)^d \cdot \left| \frac{1+s}{1-s} \right| \cdot |\Delta_K|^{\frac{3/2-\sigma}{2}} \cdot |1+s|^{\frac{3/2-\sigma}{2} \cdot d}.$$

Now, we define some field-dependent objects. Let

$$c_K = 5.545 |\Delta_K|^{1/d} \quad \text{and} \quad T_0 = 10^4 + 10^3 \cdot \ln \ln c_K.$$

Note that Minkowski's lower bound for $|\Delta_K|$ implies that $c_K > e$, and then that $T_0 \geq 10^4$. For any $t \in \mathbb{R} \setminus \{0\}$, we define the set

$$I_t = \left(\frac{1}{2} + \frac{1}{2 \ln \ln(c_K |t|)}, 1 + \frac{1}{2 \ln \ln(c_K |t|)} \right).$$

Theorem III.1.8 (Assuming ERH). *Let $\eta = 1/(16 \ln(d))$. For any t satisfying $|t| \geq T_0$ and $\sigma \in I_t$ satisfying $\sigma \geq 1 - 2\eta$, we have*

$$|\zeta_K(\sigma + it)| \leq B(K) \cdot |t|^{\eta/4}$$

for some $B(K)$ satisfying

$$\ln B(K) \ll \frac{\ln(|\Delta_K|)}{d} + d \ln d.$$

Proof. See Appendix B.2. □

Note that the proof of Theorem III.1.8 gives in fact $\ln B(K) \ll \ln(|\Delta_K|)/(d \ln(d)) + d \ln(d)$, but we simplified the bound to be of the same order of magnitude of the rest of the expressions.

III.2 Bounds on the Dedekind's Zeta Function of K

From now on, we let $X \geq 2^d > \zeta_K(3)^{1/4}/T_0$ ¹ and $T = X^4 \cdot \zeta_K(3) > T_0$. We also define $\eta = 1/(16 \ln(d)) \in (0, 1/16)$ (note that this choice is only valid if $d \geq 3$, as we assume for this work).

Lemma III.2.1 (Assuming ERH). *For any $\sigma \in [1 - 2\eta, 3]$, we have*

$$|\zeta_K(\sigma + iT)| \leq M_1(K) \cdot T^{\eta/4},$$

for some $M_1(K)$ satisfying

$$\ln(M_1(K)) \ll \frac{\ln(|\Delta_K|)}{d} + d \ln d.$$

Proof. Let $\sigma_0 = 1 + \left(6 \ln \ln \left(6|\Delta_K|^{1/d} \cdot T\right)\right)^{-1}$. For $\sigma \in [1 - 2\eta, \sigma_0]$, it can be checked that $\sigma \in I_T$, so that Theorem III.1.8 gives us that $|\zeta_K(\sigma + iT)| \leq B(K) \cdot T^{\eta/4}$.

We will use the inequality $|\zeta_K(s)| \leq \zeta(\operatorname{Re}(s))^d$, which holds for any $s \in \mathbb{C}$ satisfying $\operatorname{Re}(s) > \sigma_0 > 1$. It is obtained using the Eulerian product form of the Dedekind and Riemann zeta functions. Note that for any $x \in (1, 3)$, we have $\zeta(x) \leq 3/(x - 1)$. We have

$$\begin{aligned} |\zeta_K(\sigma + it)| &\leq \zeta(\sigma)^d \leq \zeta(\sigma_0)^d \\ &\leq \left(6 \ln \ln \left(6|\Delta_K|^{1/d} \cdot T\right)\right)^d \\ &\leq 6^d \left(\ln \left(1 + \frac{4d}{\eta}\right)\right)^d \cdot \left(6|\Delta_K|^{1/d} \cdot T\right)^{\frac{\eta}{4}} && \text{(by Lemma III.1.4)} \\ &\leq 6^d (\ln(1 + 64d \ln(d)))^d (6|\Delta_K|)^{\frac{1}{64d \ln(d)}} \cdot T^{\frac{\eta}{4}}. \end{aligned}$$

We have that

$$\begin{aligned} &d \ln(6) + d \ln \ln(1 + 64d \ln(d)) + \ln \left(6|\Delta_K|^{\frac{1}{64d \ln(d)}}\right) \\ &\ll d \ln \ln(d) + \frac{\ln(|\Delta_K|)}{d \ln(d)} \\ &\ll d \ln(d) + \frac{\ln(|\Delta_K|)}{d}, \end{aligned}$$

from which our claim follows. □

¹The bound $\zeta_K(3)^{1/4}/T_0$ is obtained by noticing that $T_0 \geq 1$ and $\zeta_K(3)^{1/4} \leq 2^d$.

Lemma III.2.2 (Assuming ERH). *For any $t \in [-T_0, T_0]$, we have that*

$$|\zeta_K(1 - 2\eta + it)| \leq M_2(K),$$

for some $M_2(K)$ satisfying

$$\ln(M_2(K)) \ll \ln(|\Delta_K|).$$

Proof. Let $s = 1 - 2\eta + it$ with $|t| \leq T_0$. By applying Theorem III.1.7, and noting that:

- $|1 - s| \geq 2\eta$,
- $|1 + s| \leq 2 \cdot T_0$,
- $(3/2 - \operatorname{Re}(s))/2 \leq 1/3$,
- $d/3 + 1 \leq d - 1$,

it holds that

$$|\zeta_K(1 - 2\eta + it)| \leq \frac{3 \cdot \zeta\left(\frac{3}{2}\right)^d}{2\eta} \cdot |\Delta_K|^{\frac{1}{3}} \cdot (2 \cdot T_0)^{d-1} := M_2(K).$$

Then it holds that

$$\ln(M_2(K)) \ll d + \ln \ln(d) + \ln |\Delta_K| + d \ln(T_0).$$

As $d \ll \ln(|\Delta_K|)$ and

$$d \ln(T_0) \ll d \ln\left(|\Delta_K|^{1/d}\right) = \ln(|\Delta_K|),$$

the result follows. □

III.3 Bounds on the integral




Figure III.1: $D_{T,\eta}$

An estimate of $N_K(B)$ is given by integrating the function ζ_K along the contour presented in Fig. III.1. In this subsection, we give bounds on the different parts of the integral. Through all this section, we will keep the notations of Section III.2.

Lemma III.3.1 (Assuming ERH). *It holds that*

$$\left| \int_{\sigma=1-2\eta}^3 \zeta_K(\sigma + iT) \cdot \frac{X^\sigma}{|\sigma + iT|} d\sigma \right| \leq 3 \cdot M_1(K) \cdot X^{-\frac{9}{10}},$$

where $M_1(K)$ is defined in Lemma III.2.1.

Proof. By Lemma III.2.1 we have,

$$\begin{aligned} \left| \int_{\sigma=1-2\eta}^3 \zeta_K(\sigma + iT) \cdot \frac{X^\sigma}{\sqrt{\sigma^2 + T^2}} d\sigma \right| &\leq M_1(K) \cdot T^{\frac{\eta}{4}-1} \int_{\sigma=1-2\eta}^3 \frac{X^\sigma}{\sqrt{(\sigma/T)^2 + 1}} d\sigma \\ &\leq 3 \cdot M_1(K) \cdot T^{\frac{\eta}{4}-1} \cdot X^3. \\ &= 3 \cdot M_1(K) \cdot \zeta_K(3)^{\frac{\eta}{4}-1} \cdot X^{\eta-4+3} \\ &\leq 3 \cdot M_1(K) \cdot X^{-\frac{9}{10}} \end{aligned}$$

where the last inequalities come from the fact that $\eta/4 - 1 < 0$ and $\zeta_K(3) \geq 1$. □

Lemma III.3.2 (Assuming ERH). *We have that*

$$\left| \int_{t=-T_0}^{t=T_0} \zeta_K(1 - 2\eta + it) \frac{X^{1-2\eta+it}}{1 - 2\eta + it} dt \right| \leq M'_2(K) \cdot X^{1-2\eta},$$

for some $M'_2(K)$ satisfying

$$\ln(M'_2(K)) \ll \ln(|\Delta_K|).$$

Proof. By Lemma III.2.2, we have that

$$\begin{aligned} \left| \int_{t=-T_0}^{t=T_0} \zeta_K(1 - 2\eta + it) \frac{X^{1-2\eta+it}}{1 - 2\eta + it} dt \right| &\leq X^{1-2\eta} \cdot M_2(K) \cdot \int_{t=-T_0}^{t=T_0} \frac{1}{\sqrt{(1-2\eta)^2 + t^2}} dt \\ &\leq X^{1-2\eta} \cdot M_2(K) \cdot \int_{t=-T_0}^{t=T_0} \frac{1}{\sqrt{(7/8)^2 + t^2}} dt \\ &\leq X^{1-2\eta} \cdot M_2(K) \cdot \frac{16}{7} T_0. \end{aligned}$$

The definition of T_0 allows us to conclude. □

Lemma III.3.3 (Assuming ERH). *We have that*

$$\left| \int_{t=T_0}^{t=T} \zeta_K(1 - 2\eta + it) \frac{X^{1-2\eta+it}}{1 - 2\eta + it} dt \right| \leq M'_3(K) \cdot X^{1-\eta},$$

for some $M'_3(K)$ satisfying

$$\ln(M'_3(K)) \ll d \ln(d) + \frac{\ln(|\Delta_K|)}{d}.$$

Proof. By Theorem III.1.8, we have that

$$\begin{aligned} \left| \int_{t=T_0}^{t=T} \zeta_K(1-2\eta+it) \frac{X^{1-2\eta+it}}{1-2\eta+it} dt \right| &\leq X^{1-2\eta} \cdot B(K) \cdot \int_{t=T_0}^{t=T} \frac{t^{\frac{\eta}{4}}}{\sqrt{(1-2\eta)^2+t^2}} dt \\ &\leq X^{1-2\eta} \cdot B(K) \cdot \frac{4}{\eta} \cdot T^{\frac{\eta}{4}} \\ &\leq 64 \cdot B(K) \cdot \ln(d) \cdot X^{1-2\eta+\eta} \cdot \zeta_K(3)^{\frac{\eta}{4}} \end{aligned}$$

Now, we have that $\ln(\zeta_K(3)^{3\eta/4}) \ll d/\ln(d)$, which gives the result. \square

Lemma III.3.4. *If $X \in \mathbb{Z}_{>0} + 1/2$, we have*

$$\left| \frac{1}{2i\pi} \int_{3-iT}^{3+iT} \zeta_K(s) \cdot \frac{X^s}{s} ds - N_K(X) \right| \leq 1.$$

Proof. This proof is in large part extracted from the proof of [Lan13, Theorem 6]. For any $s = 3 + it$, the series $\sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \mathcal{N}(\mathfrak{a})^{-s}$ is absolutely convergent. This and Lemma III.1.5 imply that (note that since $X \in \mathbb{Z}_{>0} + 1/2$, it is never equal to the norm of an integral ideal)

$$\begin{aligned} &\left| \frac{1}{2i\pi} \int_{3-iT}^{3+iT} \zeta_K(s) \cdot \frac{X^s}{s} ds - N_K(X) \right| \\ &\leq \left| \sum_{\mathcal{N}(\mathfrak{a}) \leq X} \frac{1}{2i\pi} \int_{3-iT}^{3+iT} \frac{(X/\mathcal{N}(\mathfrak{a}))^s}{s} ds - N_K(X) \right| \\ &\quad + \left| \sum_{\mathcal{N}(\mathfrak{a}) > X} \frac{1}{2i\pi} \int_{3-iT}^{3+iT} \frac{(X/\mathcal{N}(\mathfrak{a}))^s}{s} ds \right| \\ &\leq \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{\left(\frac{X}{\mathcal{N}(\mathfrak{a})}\right)^3}{\pi \cdot T \cdot \left| \ln\left(\frac{X}{\mathcal{N}(\mathfrak{a})}\right) \right|} \\ &= \frac{X^3}{\pi \cdot T} \cdot \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^3 \cdot \left| \ln\left(\frac{X}{\mathcal{N}(\mathfrak{a})}\right) \right|}. \end{aligned}$$

For any \mathfrak{a} ideal of \mathcal{O}_K , we have that if $X < \mathcal{N}(\mathfrak{a})$, then $\lfloor X \rfloor + 1 \leq \mathcal{N}(\mathfrak{a})$. Furthermore, as $X \in \mathbb{Z} + 1/2$, we have $X = \lfloor X \rfloor + 1/2$, and then

$$\left| \ln\left(\frac{X}{\mathcal{N}(\mathfrak{a})}\right) \right| = \ln\left(\frac{\mathcal{N}(\mathfrak{a})}{\lfloor X \rfloor + 1/2}\right) \geq \ln\left(\frac{\lfloor X \rfloor + 1}{\lfloor X \rfloor + 1/2}\right) = \ln\left(1 + \frac{1}{2X}\right).$$

Now, if $X > \mathcal{N}(\mathfrak{a})$, we have $\lfloor X \rfloor \geq \mathcal{N}(\mathfrak{a})$, and hence

$$\left| \ln\left(\frac{X}{\mathcal{N}(\mathfrak{a})}\right) \right| = \ln\left(\frac{\lfloor X \rfloor + 1/2}{\mathcal{N}(\mathfrak{a})}\right) \geq \ln\left(\frac{\lfloor X \rfloor + 1/2}{\lfloor X \rfloor}\right) = \ln\left(1 + \frac{1}{2\lfloor X \rfloor}\right).$$

It can be checked that for any $X > 1$, we have $(\ln(1 + 1/(2\lfloor X \rfloor)))^{-1}$ and $(\ln(1 + 1/(2X)))^{-1}$ are both $\leq 3X$. Finally, we obtain

$$\left| \frac{1}{2i\pi} \int_{3-iT}^{3+iT} \zeta_K(s) \cdot \frac{X^s}{s} ds - N_K(X) \right| \leq \frac{3}{\pi} \cdot \frac{X^4}{T} \cdot \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^3} \leq \frac{X^4 \cdot \zeta_K(3)}{T},$$

which completes the proof. \square

III.4 Bounding the ideal-counting function

We recall the statement of Theorem III.1.2.

Theorem III.1.2 (Assuming ERH). *Let K be a number field of degree $d \geq 3$ and $N_K(\cdot)$ be its ideal-counting function. For any $X \geq 2^d$, it holds that*

$$|N_K(X) - \rho_K X| \leq M(K) \cdot X^{1-\eta},$$

for $\eta = 1/(16 \ln(d))$ and some $M(K)$ satisfying

$$\ln(M(K)) \ll \ln(|\Delta_K|) + d \ln(d).$$

Proof. We first prove the theorem if $X \in \mathbb{Z} + 1/2$. Let $\gamma_1 = [3-iT, 3+iT]$, $\gamma_2 = [3+iT, 1-2\eta+iT]$, $\gamma_3 = [1-2\eta+iT, 1-2\eta-iT]$ and $\gamma_4 = [1-2\eta-iT, 3-iT]$. We let the path obtained by concatenating $\gamma_1, \gamma_2, \gamma_3$ and γ_4 be denoted by $D_{T,2\eta}$ (see Fig. III.1). Let I_j be the integral

$$I_j = \frac{1}{2i\pi} \int_{\gamma_j} \zeta_K(s) \cdot \frac{X^s}{s} ds, \text{ for } j \in \{1, 2, 3, 4\}.$$

By the residue theorem, since $1-2\eta < 1 < 3$ we have that

$$\frac{1}{2i\pi} \int_{D_{T,2\eta}} \zeta_K(s) \cdot \frac{X^s}{s} ds = \rho_K \cdot X.$$

By Lemma III.3.4, we have

$$|I_1 - N_K(X)| \leq 1.$$

By Lemma III.3.1 and the fact that $I_2 = \overline{I_4}$, we have that

$$|I_2|, |I_4| \leq 3 \cdot M_1(K) \cdot X^{-\frac{9}{10}}.$$

By Lemmas III.3.2 and III.3.3, we have that

$$|I_3| \leq 2 \cdot M'_3(K) \cdot X^{1-\eta} + M'_2(K) \cdot X^{1-2\eta}.$$

As $\ln(M_1(K)), \ln(M'_2(K)), \ln(M'_3(K)) \ll \ln|\Delta_K| + d \ln d$, the result follows, we denote by $M'(K)$ the constant for $X \in \mathbb{Z} + 1/2$.

Now, we generalize to $X \in \mathbb{R}_{>0}$. For any $X \in \mathbb{R}_{>0}$, we have that $N_K(X) = N_K(\lfloor X \rfloor + 1/2)$, which implies that

$$\begin{aligned} |N_K(X) - \rho_K X| &\leq |N_K(\lfloor X \rfloor + 1/2) - \rho_K \cdot (\lfloor X \rfloor + 1/2)| + \rho_K |\lfloor X \rfloor + 1/2 - X| \\ &\leq M'(K) \cdot (2X)^{1-\eta} + \rho_K \leq (2M'(K) + \rho_K) \cdot X^{1-\eta}. \end{aligned}$$

By Theorem II.2.10, we have that $\rho_K \leq (e \cdot \log(\Delta_K)/(2(d-1)))^{d-1}$. The result follows by taking $M(K) = 2M'(K) + \rho_K$. \square

We recall the statement of Corollary III.1.3.

Corollary III.1.3 (Assuming ERH). *Let K be a number field of degree $d \geq 3$ and $N_K(\cdot)$ be its ideal-counting function. There exists a field independent constant $c_2 > 1$ such that for any $C \geq 1$, if*

$$X \geq (C \cdot d^d \cdot |\Delta_K|)^{c_2 \cdot \ln(d)}$$

then it holds that

$$\frac{|N_K(X) - \rho_K \cdot X|}{\rho_K \cdot X} \leq \frac{1}{C}.$$

We make use of the following lemma, which transforms the additive error on N_K to a multiplicative one.

Lemma III.4.1 (Assuming ERH). *Let $C \geq 1$ and $X \geq 2^d$. If*

$$X \geq (2C \cdot |\Delta_K|^{\frac{1}{2}} \cdot M(K))^{16 \ln(d)}$$

(where $M(K)$ is defined in Theorem III.1.2), then we have

$$\frac{|N_K(X) - \rho_K \cdot X|}{\rho_K \cdot X} \leq \frac{1}{C}.$$

Proof. Theorem III.1.2 and Lemma III.1.6 give that

$$\frac{|N_K(X) - \rho_K \cdot X|}{\rho_K \cdot X} \leq \frac{2\sqrt{|\Delta_K|} \cdot M(K)}{X^\eta},$$

which is $\leq 1/C$ given our assumption on X . □

Proof of Corollary III.1.3. This follows directly from Lemma III.4.1 and the bound on $\ln(M(K))$ from Theorem III.1.2. □

Chapter IV

Ideal-SVP is Hard for Small-Norm Uniform Prime Ideals

This chapter is extracted from [FPSW23]. This work was done in collaboration with Alice Pellet-Mary, Damien Stehlé and Benjamin Wesolowski. I took care of most of the technical aspects of the reduction, and extracted and rewrote in more modern notations the worst-case to average-case reduction from Gentry's thesis [Gen09].

IV.1 Introduction

Contributions of this chapter.

We describe a new quantum self-reduction for id-HSVP. We prove that if \mathcal{W} is a set of ideals and \mathcal{W}^{-1} is the set of inverses of the ideals of \mathcal{W} , then solving id-HSVP for the uniform distribution over \mathcal{W}^{-1} reduces to solving id-HSVP for the uniform distribution over \mathcal{W} and to solving id-HSVP for a uniform ideal within those having their norms in a prescribed interval. Both the cost of the reduction and the loss in the approximation factor are polynomially bounded in the degree d and the root-discriminant $\Delta_K^{1/d}$ of the number field. The precise statement is provided in Theorem IV.5.1.

When specialized with \mathcal{W} chosen as the set of prime ideals of algebraic norm $\Delta_K^{O(1)} \cdot d^{O(d)}$, our reduction implies that solving id-HSVP for the inverse of uniform primes ideals is no harder than solving it for uniform prime ideals (still for those of algebraic norm $\Delta_K^{O(1)} \cdot d^{O(d)}$). The success probability of this reduction is proportional to the proportion of prime ideals among all integral ideals of norm bounded by some $A = \text{poly}(\Delta_K)$. Combined with Gentry's reduction [Gen09], our work implies the random self-reducibility of id-HSVP for the uniform distribution over prime ideals. As Gentry's original reduction considers the bounded distance decoding problem, we present an adaptation to the shortest vector problem in Appendix C.2. Note that the polynomial dependency in the proportion of prime ideals may have a considerable impact on the cost of this reduction (there exists number fields for which the proportion of prime ideals is exponentially small in the degree).

This new reduction, along with the Karp reduction of [PS21], gives a new distribution over NTRU instances with modulus polynomial in d and $\Delta_K^{1/d}$ whose difficulty relies on the worst-case problem id-HSVP. To our knowledge this is the first time a distribution over NTRU instance with polynomial modulus is based on a worst-case problem, even though this distribution needs a factoring oracle to be sampled from.

Technical overview.

We now give an overview of the average-case to average-case reduction for id-HSVP. Let \mathcal{W} be a set of fractional ideals represented by their Hermite Normal Form. The goal of our reduction is to find (with non-negligible probability) a short non-zero vector in a given uniform element of \mathcal{W}^{-1} , given access to two oracles: $\mathcal{O}_{\mathcal{W}}$ which solves id-HSVP with non-negligible probability for a uniform element of \mathcal{W} , and $\mathcal{O}_{\mathcal{I}}$ which solves it with non-negligible probability for a uniform integral ideal with norm between A and $4A$, for $A = \Delta_K^{O(1/d)} \cdot d^{O(1)}$. In everything that follows we assume that we have a factoring oracle (for integers, or equivalently, for integral ideals). Such an oracle can be instantiated in quantum polynomial time with Shor's algorithm, or in sub-exponential time with the number field sieve algorithm.

Before diving into our contribution, let us explain a key idea developed in [Boe22, Chap. 6]. By *ideal of norm 1*, we mean a (replete) ideal¹ of the form $I/\mathcal{N}(I)^{1/d}$. The space of ideals of norm 1 has a natural notion of uniformity. Let B_r denote the ℓ_∞ ball of radius r . In [Boe22, Theorem 6.21], it is proved that if J is sampled uniformly in the set of ideals of norm 1, and x is uniform in $B_r \cap J$, then the integral ideal $x \cdot J^{-1}$ is almost uniform in the set of integral ideals of norm less than r^d .

Now, our reduction follows the following structure. We are given a uniform $I \in \mathcal{W}$, and tasked with finding a short non-zero vector $v_{I^{-1}} \in I^{-1}$.

1. Find a short non-zero vector $v_I \in I$ with the oracle $\mathcal{O}_{\mathcal{W}}$.
2. Generate a uniform norm-1 ideal I' , together with a non-zero vector $v_{I'} \in I'$ as short as possible. The ideal $J = I' \cdot I/\mathcal{N}(I)^{1/d}$ is also uniform in the space of ideals of norm 1, and we can compute a short basis \mathbf{B}_J of J thanks to the short non-zero vectors v_I and $v_{I'}$.
3. Sample $x \in B_r \cap J$; this uses our knowledge of the good basis \mathbf{B}_J . Hopefully, the integral ideal $\mathfrak{b} = x \cdot J^{-1}$ is almost uniform in the set of integral ideals of bounded norm.
4. Find a short non-zero vector $v_{\mathfrak{b}} \in \mathfrak{b}$ with the oracle $\mathcal{O}_{\mathcal{I}}$.
5. Return the vector $v_{I^{-1}} = x^{-1} \cdot v_{I'} \cdot v_{\mathfrak{b}} \cdot \mathcal{N}(I)^{-1/d} \in I^{-1}$.

One can check that $v_{I^{-1}} \in I^{-1}$, but is it short? Its factors are short by construction, except possibly x^{-1} . Indeed, the element x itself is bounded (it is in the set B_r), but its inverse may not be. To circumvent this issue, we would like to replace the ℓ_∞ ball B_r with another shape X which contains only *balanced* vectors (i.e., close to a vector of the form $\lambda \cdot (1, \dots, 1)$), so that for any short $x \in X$, we have that x^{-1} is small. We prove that the result of [Boe22] holds for general sets X satisfying certain conditions. We consider a new shape $\mathcal{B}_{A,B}^\eta$ (see Figure IV.1 and Definition IV.4.1) that satisfies the conditions, and contains only balanced elements. Now, replacing B_r with $\mathcal{B}_{A,B}^\eta$ in Step 3, we sample an element x such that x^{-1} is small, hence all the factors of $v_{I^{-1}}$ are small, and $v_{I^{-1}}$ is indeed a solution to id-HSVP in I^{-1} .

While Step 3 constitutes the main difficulty of the reduction, and the technical core of our paper, let us briefly comment on Step 2. We need to sample a uniform norm-1 ideal I' , together with a short non-zero vector $v_{I'} \in I'$. In [BDPW20], it is proven that if an ideal \mathfrak{p} is sampled uniformly in the set of prime ideals with norm less than $(d^d \cdot \Delta_K)^c$ for some constant c , then, up to a small Gaussian factor, the ideal $\mathfrak{p}/\mathcal{N}(\mathfrak{p})^{1/d}$ is close to uniform in the set of norm-1 ideals. It is therefore sufficient to generate such a prime ideal \mathfrak{p} together with a short element $v_{\mathfrak{p}} \in \mathfrak{p}$. The technique is extracted from [Gen09, Chap. 17], and requires a factoring oracle. It first samples

¹A replete ideal is a subset of $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$ of the form $\alpha \cdot I$ where $I \subseteq \mathcal{O}_K$ is an integral ideal of \mathcal{O}_K and $\alpha \in K_{\mathbb{R}}^\times$ is invertible. More details can be found in the preliminaries.




Figure IV.1: A plot of $\mathcal{B}_{A,B}^\eta$ intersected with the subspace $K_{\mathbb{R}}^+ := \{x \in K_{\mathbb{R}} \mid \sigma_i(x) \in \mathbb{R}_{>0} \text{ for all } i\}$. Here we have $(d_{\mathbb{R}}, d_{\mathbb{C}}) = (3, 0)$, $A = 20$, $B = 40$ and $\eta = \exp(1)$.

a small element $x \in \mathcal{O}_K$ with the Gaussian distribution. It then factors $(x) = \mathfrak{p}_1^{e_1} \cdots \mathfrak{p}_k^{e_k}$ and uniformly selects one of the factors \mathfrak{p}_i . Finalizing with a rejection sampling step, it can be proved that the chosen \mathfrak{p} is almost uniform in the set of primes of norm $\lesssim \mathcal{N}(x)$.

We now have a reduction from id-HSVP for inverses of ideal of a set \mathcal{W} , to id-HSVP for ideals of \mathcal{W} and id-HSVP for a uniform ideal of norm in some interval $[A, 4A]$ for A as small as $\Delta_K^{O(1)} \cdot d^{O(d)}$. This gives a trivial reduction from id-HSVP for a uniform ideal to id-HSVP for an uniform prime ideal, with a success probability decrease of a factor $O(1/\tilde{\rho}_A)$, where $1/\tilde{\rho}_A$ is the proportion of prime ideals among the set of all integral ideals of norm $\leq A$. We can now combine this last reduction with our main result (taking \mathcal{W} to be the set of prime ideals of norm in $[A, 4A]$) in order to reduce id-HSVP for inverses of prime ideals to id-HSVP for prime ideals. This, combined with the worst-case to average-case reduction of [Gen09] gives a worst-case to average-case reduction for id-HSVP where the average-case is the uniform distribution over prime ideals of norm in $[A, 4A]$.

Finally, note that a reduction from id-HSVP to NTRU was recently given in [PS21]. It transforms an integral ideal I into an NTRU instance of modulus polynomially larger than $\mathcal{N}(I)^{1/d}$. Our self-reduction (in contrast with the one from [Gen09]) applies to integral ideals and can be composed with the one from [PS21]. The distribution of NTRU instances obtained by sampling a uniform prime ideal of norm in $[A, 4A]$ and applying [PS21, Alg. 4.1] is at least as difficult to solve as worst-case id-HSVP. By setting $A = \Delta_K^{O(1)} \cdot d^{O(d)}$, we obtain an NTRU modulus bounded as $\Delta_K^{O(1/d)} \cdot d^{O(1)}$. Note that “overstretched NTRU” attacks [ABD16, CJL16, KF17] do not apply for this distribution as, among others, they require a much larger modulus.

Related works on the hardness of id-HSVP.

On the upper bound front, it has been shown that id-HSVP is susceptible to better algorithms than the generic HSVP. Cramer et al. [CDPR16] described an algorithm for id-HSVP in cyclotomic fields for principal ideals with an approximation factor $\exp \tilde{O}(\sqrt{d})$ in quantum polynomial time. It was later generalized to all ideals [CDW17] of cyclotomic fields and (with pre-processing) to all number fields [PHS19]. Note that in the present work, all our reductions feature polynomial losses on the approximation factor, and hence apply to id-HSVP for polynomial approximation factors, a regime that is not impacted by these algorithms. Still, families of easy instances for id-HSVP have been identified even for polynomial approximation factors [PXWC21, PML21, BEP22], specifically ideals stabilized by many field automorphisms.

While these families are very sparse, their existence further motivates the study of different distributions of id-HSVP instances.

IV.2 Preliminaries

When using oracles with a non-zero probability of failing, we assume without loss of generality that either the oracle returns a valid result or \perp (as in our cases, the validity of the output can always be checked efficiently).

IV.2.1 Balanced elements.

For the reductions presented in this article, it will sometimes be convenient to use balanced elements of $K_{\mathbb{R}}$, i.e., elements whose ℓ_{∞} norm and the one of their inverse are not far from the geometric mean of their coordinates: in other terms they do not have an exceptionally small or large coordinate in comparison to the others. This property is convenient as it implies that multiplying an ideal by one of these elements will not change its geometry significantly, in particular if x is balanced and v is small in the ideal $x \cdot I$, then $x^{-1} \cdot v$ will be small in I . The formal definition is as follows.

Definition IV.2.1. *Let $\eta > 1$. An element x in $K_{\mathbb{R}}$ is said to be η -balanced if*

$$\|x\|_{\infty} \leq \eta \cdot |\mathcal{N}(x)|^{\frac{1}{d}} \quad \text{and} \quad \|x^{-1}\|_{\infty} \leq \eta \cdot |\mathcal{N}(x)|^{-\frac{1}{d}}.$$

IV.2.2 Density of prime ideals.

For any $A \geq 1$, we let $\tilde{\rho}_A$ denote the inverse of the proportion of prime ideals among all integral ideals of K of norm $\leq A$, i.e.,

$$\tilde{\rho}_A := \frac{|\{\mathfrak{a} \subset \mathcal{O}_K \mid \mathcal{N}(\mathfrak{a}) \leq A\}|}{|\{\mathfrak{p} \subset \mathcal{O}_K \text{ prime} \mid \mathcal{N}(\mathfrak{p}) \leq A\}|}.$$

In this article, we will mostly be interested in $\tilde{\rho}_A$ for values of A of the order of $\text{poly}(\Delta_K)$. Unfortunately, we are not aware of estimates for $\tilde{\rho}_A$ when A is this “small”. However, it is known that when the number field K is fixed and A tends to infinity, it holds that

$$\tilde{\rho}_A \underset{A \rightarrow \infty}{\sim} \rho_K \cdot \ln(A),$$

where ρ_K is the residue of the Dedekind zeta function at 1. This comes from the fact that

$$|\{\mathfrak{p} \subset \mathcal{O}_K \text{ prime} \mid \mathcal{N}(\mathfrak{p}) \leq A\}| \sim A / \ln(A)$$

(see [BS96, Theorem 8.7.4]), and that $|\{\mathfrak{a} \subset \mathcal{O}_K \mid \mathcal{N}(\mathfrak{a}) \leq A\}| \sim \rho_K \cdot A$. (see [Web08]). We can though give a bound on $\tilde{\rho}_A$ when $A \geq (2 \cdot \Delta_K \cdot d)^{c_2 \cdot \ln(d)}$ using Corollary III.1.3. In this case, we have

$$\tilde{\rho}_A \in [0.5, 1.5] \cdot \rho_K \ln(A)$$

IV.2.3 Algorithmic problems

Algorithmic problems in ideals.

Lemma IV.2.2 (Folklore). *For any $\gamma \geq 1$, there is a Karp polynomial-time reduction from the problem id-HSVP_{γ} to inv-HSVP_{γ} .*

Proof. Let I be a fractional ideal for which we want to solve the id-HSVP_γ problem. We will show that there exists $x \in \mathbb{Q}$ such that $xI = \mathfrak{a}^{-1}$ is the inverse of an integral ideal $\mathfrak{a} \subseteq \mathcal{O}_K$. If such an element x can be computed efficiently, then the reduction simply computes x , then compute $\mathfrak{a}^{-1} = xI$ and runs the inv-HSVP_γ solver on \mathfrak{a}^{-1} (which is a valid input for inv-HSVP). Since multiplication by $x \in \mathbb{Q}$ consists in scaling the lattice corresponding to I , then a solution to id-HSVP_γ in xI provides a solution to id-HSVP_γ in I (by multiplying it by x^{-1}). Note that the reduction preserves the approximation factor γ .

Let us then show that such an x exists and can be computed in polynomial time. Write $I = \mathfrak{a}\mathfrak{b}^{-1}$, with $\mathfrak{a}, \mathfrak{b} \subseteq \mathcal{O}_K$ integral ideals, and define $x = \mathcal{N}(\mathfrak{a})^{-1}$. Note that such $\mathfrak{a}, \mathfrak{b}$ and x can be computed in polynomial time from I (we do not require that \mathfrak{a} and \mathfrak{b} are coprime, so the choice we make is not unique). Let us show that for such x , it holds that $(xI)^{-1} \subseteq \mathcal{O}_K$ is an integral ideal. By definition, we have $(xI)^{-1} = \mathcal{N}(\mathfrak{a}) \cdot \mathfrak{a}^{-1}\mathfrak{b}$. Since \mathfrak{a} is integral, it holds that $\mathcal{N}(\mathfrak{a}) \cdot \mathcal{O}_K \subseteq \mathfrak{a}$. Indeed, note that the group $\mathcal{O}_K/\mathfrak{a}$ has cardinality $\mathcal{N}(\mathfrak{a})$. Lagrange's theorem then gives that any element of $\mathcal{O}_K/\mathfrak{a}$ has order dividing $\mathcal{N}(\mathfrak{a})$, i.e., for any $x \in \mathcal{O}_K$, we have $\mathcal{N}(\mathfrak{a}) \cdot x \in \mathfrak{a}$. We hence obtain that the ideal $\mathcal{N}(\mathfrak{a}) \cdot \mathfrak{a}^{-1} \subseteq \mathcal{O}_K$ is integral. Since \mathfrak{b} is integral by construction, this proves that $(xI)^{-1}$ is integral. \square

IV.2.4 Algorithms on ideals

For $I = \mathcal{O}_K$, the following lemma states that one can quantumly and efficiently sample a random prime ideal together with a short element in it, hence the name. We give a proof based on [PS21] but note that a similar statement was already given as [Gen09, Theorem 16.3.4, Lemma 17.2.1] (see also [Gen10, Se. 3.3]).

Lemma IV.2.3 (Adapted from [PS21, Lemma C.1]). *There exists an algorithm `SampleWithTrap` that on input integers $2 \leq A < B$, a real $\varepsilon \in (0, 1)$ and a basis \mathbf{B}_I of a fractional ideal I , samples a pair (\mathfrak{p}, w) such that*

1. *the distribution of \mathfrak{p} is within statistical distance ε from the uniform distribution over $\mathcal{P}_{A,B}$;*
2. *the element w belongs to $I \cdot \mathfrak{p} \setminus \{0\}$;*
3. *we have $\|w\| \leq 2\sqrt{4d + \ln(24B/\varepsilon)} \cdot \varsigma$ with $\varsigma = \max(\varsigma_{\text{sample}}, \varsigma_{\text{smooth}})$ and*
 - $\varsigma_{\text{sample}} = \sqrt{d} \cdot \|\mathbf{B}_I^*\|$.
 - $\varsigma_{\text{smooth}} = (\Delta_K \cdot B \cdot \mathcal{N}(I))^{1/d} \cdot \sqrt{\ln(24B/\varepsilon)}$.

Furthermore, if the algorithm is given access to an oracle factoring integral ideals of norm smaller than $(2\sqrt{4d + \ln(24B/\varepsilon)} \cdot \varsigma)^d \cdot \mathcal{N}(I)^{-1}$, then the algorithm runs in expected time polynomial in $B/|\mathcal{P}_{A,B}|, B/A, \log \Delta_K, \log B, \log(1/\varepsilon)$ and in the size of I .

The proof is available in Appendix C.1. Note that we will use this result with $\varepsilon = \exp(-d)$ in order to simplify computations and subsequently omit this input.

Factoring ideals. Factoring an integral ideal \mathfrak{a} in \mathcal{O}_K can be done by factoring the algebraic norm $\mathcal{N}(\mathfrak{a})$ of \mathfrak{a} over the integers; computing, for all the prime factors $p \mid \mathcal{N}(\mathfrak{a})$, the set of prime ideals whose norm is a power of p (there are at most d of those); and testing for each of these prime if they divide \mathfrak{a} . Factoring $\mathcal{N}(\mathfrak{a})$ can be performed quantumly in time polynomial in $\log \mathcal{N}(\mathfrak{a})$ (using Shor's algorithm [Sho94]). Computing the set of prime ideals of norm a given prime integer p can be performed classically in time polynomial in $\log p$ and $\log \Delta_K$ using Buchmann-Lenstra's algorithm [BL94], described in details in [Coh93, Sec. 6.2.5]. Finally, testing whether

a prime ideal \mathfrak{p} divides \mathfrak{a} can be done in time polynomial in the bit-sizes of \mathfrak{p} and \mathfrak{a} . Overall, factoring ideals can be done in quantum-polynomial time (using Shor’s algorithm) or in classical sub-exponential time (using the Number Field Sieve).

IV.2.5 Worst-case to average-case reduction for inverse of primes

In [Gen09, Ch. 16 & 17], Gentry described a self-reduction for a variant of the bounded distance decoding problem, from worst-case ideals to prime ideals taken uniformly at random with their norm in some interval $[A, B]$ (for a suitable choice of A and B). This reduction can be adapted to the shortest vector problem (instead of the bounded distance decoding problem), but it requires to take the inverse of the ideals, implying that the average-case distribution we obtain is over the inverses of prime ideals uniformly chosen in the interval $[A, B]$. Below, we state the result of Gentry’s reduction adapted to SVP, and provide a proof in Appendix C.2 for the sake of completeness.

Theorem IV.2.4 (Adapted from [Gen09, Ch. 16 & 17]). *There exist some field dependent constant $C_{1,K} = \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and $C_{2,K} = \text{poly}(\log \Delta_K, \delta_K)$ such that the following holds. Let $\gamma_{\text{avg}} \in [1, 2^d]$, $A \geq C_{1,K}^d \cdot \gamma_{\text{avg}}^d$ satisfying $A \leq (\Delta_K)^{d^{O(1)}}$ and $\gamma = A^{1/d} \cdot C_{2,K}$. Then*

$$\text{id-HSVP}_\gamma \text{ reduces to } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

The reduction is probabilistic and, assuming it has access to an oracle factoring integral ideals whose norms have bit-size $\text{poly}(\log \Delta_K)$, it runs in expected time polynomial in its input size, $\log \Delta_K$ and $1/\delta$, where δ is the success probability of the $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ oracle.²

IV.3 Self-Reducibility of id-HSVP to Inverses

Let \mathcal{W} be a finite set of fractional ideals. In this section, we provide a framework for reducing id-HSVP for the uniform distribution over \mathcal{W} to id-HSVP for the uniform distribution over $\mathcal{W}^{-1} = \{I^{-1} : I \in \mathcal{W}\}$. The reduction, provided in Theorem IV.3.4 relies on three oracles (beyond the one for id-HSVP for $\mathcal{U}(\mathcal{W})$). The first one factors integral ideals, and can be instantiated with a quantum polynomial-time algorithm. The second one samples from $I \cap X$, where I is an arbitrary norm-1 replete ideal and X is a well-chosen set: this oracle will be instantiated in Section IV.4. The last one finds short non-zero vectors in integral ideals uniformly distributed within those having their norms in a prescribed interval. Overall, this will lead to a quantum polynomial-time reduction from \mathcal{W}^{-1} -avg-id-HSVP to \mathcal{W} -avg-id-HSVP and $\mathcal{I}_{A,4A}$ -avg-id-HSVP for a well-chosen A .

The reduction is built in several steps. First, we show how to map a uniform norm-1 replete ideal to an integral ideal uniform among those with norms in $[A, 4A]$, using a new approach introduced in [Boe22, Sec. 6]. This is parameterized by a set X that will be instantiated in Section IV.4. The second step gives a way to randomize an arbitrary ideal to an integral ideal uniform among those with norms in $[A, 4A]$, along with a hint that allows to map a short vector of the resulting ideal to a short vector in the inverse of the input ideal. Finally, this allows to describe the reduction.

²The choice of $4A$ for the upper bound on the norm of the ideals is not a strict requirement of this theorem. We instantiated the theorem with this value in order to simplify its statement.

IV.3.1 From a uniform norm-1 ideal to a uniform integral ideal

In this subsection, we present a way to sample uniformly among integral ideals whose norms belong to a prescribed interval. Given a compact set X satisfying certain properties, our sampler takes as input a uniform ideal $I \in \text{IdLat}_K^0$, samples a point uniformly in $I \cap X$ and outputs $(x) \cdot I^{-1} \subseteq \mathcal{O}_K$. It holds that if X is well-designed, then the output distribution is close to the uniform distribution over the set of integral ideals in terms of Rényi divergence. Our sampler generalizes [Boe22, Theorem 6.9], where the set X is assumed to be the ℓ_∞ ball. This new degree of freedom will allow us (in Section IV.4) to choose a set X whose points are balanced, which will be essential for the proof of Theorem IV.5.1. Note that we do not use the Arakelov ray divisor formalism to state our results: those of [Boe22, Sec. 6] are stated with respect to a modulus $\mathfrak{m} \subseteq \mathcal{O}_K$ and here we take $\mathfrak{m} = \mathcal{O}_K$.

Definition IV.3.1. *Let $X \subset K_{\mathbb{R}}$. We say that X is compact and invariant by complex rotations if the following hold:*

- $\Phi(X)$ is a compact subset of \mathbb{C}^d ;
- for any $\zeta = (\zeta_1, \dots, \zeta_d) \in \Phi(K_{\mathbb{R}})$ with $|\zeta_1| = \dots = |\zeta_d| = 1$, it holds that $\Phi^{-1}(\zeta) \cdot X = \Phi^{-1}(\zeta) \cdot x \mid x \in X \subseteq X$.

We consider the `IdealRound` algorithm (Algorithm IV.3.1), whose output distribution generalizes the distribution presented in [Boe22, Theorem 6.9]. It is parametrized by an arbitrary compact set $X \subset K_{\mathbb{R}}$, takes as input a norm-1 replete ideal (i.e., an element of IdLat_K^0) and returns an integral ideal. We define $\mathcal{D}_{\text{Ideal}}(X)$ as the distribution `IdealRound` $_X(\mathcal{U}(\text{IdLat}_K^0))$. For the moment, we are not interested in the efficiency of `IdealRound` $_X$, but only in the relationship between $\mathcal{D}_{\text{Ideal}}(X)$ and the uniform distribution over ideals with norms belonging to an interval. This is the purpose of the following result.

Algorithm IV.3.1 IdealRound

Input: $I \in \text{IdLat}_K^0$.

Parameter: $X \subset K_{\mathbb{R}}$ compact.

Output: An integral ideal \mathfrak{a} .

- 1: Sample $x \leftarrow \mathcal{U}(I \cap X)$.
 - 2: Return $\mathfrak{a} = (x) \cdot I^{-1}$.
-

Lemma IV.3.2. *For any $t \in \mathbb{R}$, let $H_t = \{x \in \text{Ln } K_{\mathbb{R}} \mid \sum_i x_i = t\}$. Let X be a compact subset of $K_{\mathbb{R}}$ invariant by complex rotations (as per Definition IV.3.1) and $B > A > 2$. Assume that:*

- There exist some real numbers $C \geq 1$ and $C' > 0$ such that we have $|I \cap X| \in C' \cdot [1, C]$ for any $I \in \text{IdLat}_K^0$;
- there exists $C'' \in \mathbb{R}$ such that for any $t \in [\ln(A), \ln(B)]$ we have

$$\text{Vol}\left(\text{Ln}(X) \cap H_t\right) = C'';$$

- for any $t \notin [\ln(A), \ln(B)]$, we have $\text{Vol}(\text{Ln}(X) \cap H_t) = 0$.

Then the support of $\mathcal{D}_{\text{Ideal}}(X)$ is contained in $\mathcal{I}_{A,B}$ and

$$\text{RD}_\infty(\mathcal{U}(\mathcal{I}_{A,B}) \parallel \mathcal{D}_{\text{Ideal}}(X)) \leq C.$$

We now comment the conditions of Lemma IV.3.2. The second and third conditions state that, when embedded in $\text{Ln}(K_{\mathbb{R}})$ the set $\text{Ln}(X)$ should be contained between the two hyperplanes $H_{\log(A)}$ and $H_{\log(B)}$, and that between those hyperplanes, the slices $\text{Ln}(X) \cap H_t$ should have constant volume. Those conditions will yield the bounds on the norm of the output ideal. The first condition states that for any norm-1 replete ideal I , the number of points in $X \cap I$ should be non-zero and almost independent of I . Conditions 1 and 2 will imply the near-uniformity of the output distribution. The proof below is adapted from [Boe22, Theorem 6.9].

Proof. Fix an integral ideal \mathfrak{b} and a norm-1 replete ideal I . We are going to compute bounds on

$$p_{I,\mathfrak{b}} = \Pr_x((x) \cdot I^{-1} = \mathfrak{b}) = \Pr_x((x) = I \cdot \mathfrak{b}) = \Pr_x(x \text{ generates } I \cdot \mathfrak{b}),$$

where the randomness is over $x \leftarrow \mathcal{U}(I \cap X)$. For an ideal J , we define $G_J = \{x \in K_{\mathbb{R}} : (x) = J\}$ as the set of generators of J (if J is not principal, it is the empty set). Note that $G_{I \cdot \mathfrak{b}} = \{x \in K_{\mathbb{R}} : (x) = I \cdot \mathfrak{b}\} \subseteq I$. We have

$$p_{I,\mathfrak{b}} = \frac{|G_{I \cdot \mathfrak{b}} \cap X|}{|I \cap X|} \in \left| G_{I \cdot \mathfrak{b}} \cap X \right| \cdot C'^{-1} \cdot [C^{-1}, 1],$$

where the inclusion follows from the first assumption of the lemma. For any I that is not in the class of \mathfrak{b}^{-1} modulo principal ideals, we have that $G_{I \cdot \mathfrak{b}}$ is empty, since $I \cdot \mathfrak{b}$ is not principal. Let $[\mathfrak{b}^{-1}]^0$ be the set of all norm-1 replete ideals of the form $(\alpha) \cdot \mathfrak{b}^{-1}$ for some $\alpha \in K_{\mathbb{R}}$ (i.e., the coset of \mathfrak{b}^{-1} in IdLat_K^0 modulo principal ideals). Let $I_0 = \mathcal{N}(\mathfrak{b})^{1/d} \cdot \mathfrak{b}^{-1}$, which belongs to $[\mathfrak{b}^{-1}]^0$. There is a bijection between $K_{\mathbb{R}}^0/\mathcal{O}_K^\times$ and $[\mathfrak{b}^{-1}]^0$ given by $u \mapsto (u) \cdot I_0$. This implies that

$$\begin{aligned} \mathbb{E}_{I \leftarrow \mathcal{U}(\text{IdLat}_K^0)} \left(\left| G_{I \cdot \mathfrak{b}} \cap X \right| \right) &= \Pr_{I \leftarrow \mathcal{U}(\text{IdLat}_K^0)} (I \in [\mathfrak{b}^{-1}]^0) \cdot \mathbb{E}_u \left(\left| G_{\mathcal{N}(\mathfrak{b})^{1/d} \cdot (u)} \cap X \right| \right) \\ &= \frac{1}{|\text{Cl}_K|} \cdot \mathbb{E}_u \left(\left| G_{\mathcal{N}(\mathfrak{b})^{1/d} \cdot (u)} \cap X \right| \right), \end{aligned}$$

where $u \leftarrow \mathcal{U}(K_{\mathbb{R}}^0/\mathcal{O}_K^\times)$ and the second equality comes from Lemma II.2.13. Let μ_K be the set of roots of unity in K . Using the fact X is invariant by complex rotations and that the Ln function is $|\mu_K|$ -to-1 when its input is restricted to generators of a principal replete ideal I , it holds that the Ln function is $|\mu_K|$ -to-1 on $G_I \cap X$, and then we have:

$$\forall I \in \text{IdLat}_K^0 : \left| G_I \cap X \right| = |\mu_K| \cdot \left| \text{Ln}(G_I) \cap \text{Ln}(X) \right|.$$

In our context, this implies that for any $u \in K_{\mathbb{R}}^0$,

$$\begin{aligned} \left| G_{\mathcal{N}(\mathfrak{b})^{1/d} \cdot (u)} \cap X \right| &= |\mu_K| \cdot \left| \text{Ln}(X) \cap \left\{ \text{Ln}(x) : x = v \cdot u \cdot \mathcal{N}(\mathfrak{b})^{1/d}, v \in \mathcal{O}_K^\times \right\} \right| \\ &= |\mu_K| \cdot \left| \text{Ln}(X) \cap (\Lambda_K + \text{Ln}(u) + \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \right| \\ &= |\mu_K| \cdot \left| (\text{Ln}(X) - \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \cap (\Lambda_K + \text{Ln}(u)) \right|, \end{aligned}$$

where $\Lambda_K = \text{Ln} \mathcal{O}_K^\times$. Note that Λ_K is full rank in H_0 , and that $\text{Ln}(u) \in H_0$ for any $u \in K_{\mathbb{R}}^0$. Moreover, the vector $\text{Ln}(u)$ is uniform in H_0/Λ_K when u is uniform in $K_{\mathbb{R}}^0/\mathcal{O}_K^\times$. We are hence considering a uniform lattice shift and, for any measurable set $\mathcal{S} \subseteq H_0$, we have:

$$\mathbb{E}_u \left(\left| (\Lambda_K + \text{Ln}(u)) \cap \mathcal{S} \right| \right) = \frac{\text{Vol}(\mathcal{S})}{\text{Vol}(\Lambda_K)}.$$

Applying this to the set $\mathcal{S} = (\text{Ln}(X) - \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \cap H_0$, we obtain

$$\mathbb{E}_u \left(\left| G_{\mathcal{N}(\mathfrak{b})^{1/d}, (u)} \cap X \right| \right) = |\mu_K| \cdot \frac{\text{Vol}((\text{Ln}(X) - \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \cap H_0)}{\text{Vol}(\Lambda_K)}.$$

Observe that by definition of H_t for $t \in \mathbb{R}$, it holds that

$$(\text{Ln}(X) - \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \cap H_0 = \left(\text{Ln}(X) \cap H_{\ln \mathcal{N}(\mathfrak{b})} \right) - \text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d}).$$

Since shifting by $\text{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})$ does not change the volume, we obtain

$$\mathbb{E}_u \left(\left| G_{\mathcal{N}(\mathfrak{b})^{1/d}, (u)} \cap X \right| \right) = |\mu_K| \cdot \frac{\text{Vol}(\text{Ln}(X) \cap H_{\ln \mathcal{N}(\mathfrak{b})})}{\text{Vol}(\Lambda_K)}.$$

Recall from the second and third assumptions that

$$\text{Vol} \left(\text{Ln}(X) \cap H_{\ln \mathcal{N}(\mathfrak{b})} \right) = \begin{cases} C'' & \text{if } \ln \mathcal{N}(\mathfrak{b}) \in [\ln A, \ln B], \\ 0 & \text{otherwise.} \end{cases}$$

Let $p = C'' \cdot |\mu_K| / (C' \cdot |\text{Cl}_K| \cdot \text{Vol}(\Lambda_K))$. Combining everything, this proves that

$$p_{\mathfrak{b}} := \mathbb{E}_{I \leftarrow \mathcal{U}(\text{IdLat}_K^0)} (\mathfrak{p}_{\mathfrak{b}, I}) \in \begin{cases} p \cdot [C^{-1}, 1] & \text{if } \mathcal{N}(\mathfrak{b}) \in [A, B], \\ \{0\} & \text{otherwise.} \end{cases}$$

Observe that $p_{\mathfrak{b}}$ is equal to $\mathcal{D}_{\text{Ideal}}(X)(\mathfrak{b})$, the probability of the ideal \mathfrak{b} for the distribution $\mathcal{D}_{\text{Ideal}}(X)$. The equation above then means that $\mathcal{D}_{\text{Ideal}}(X)$ outputs ideals with norm in $[A, B]$ with probability essentially equal to p (up to a factor C), and other ideals with probability 0. We quantify this using the Rényi divergence. As $1 = \sum_{\mathfrak{b} \in \mathcal{I}_{A,B}} p_{\mathfrak{b}} \in p \cdot |\mathcal{I}_{A,B}| \cdot [C^{-1}, 1]$, we have that $p \in |\mathcal{I}_{A,B}|^{-1} \cdot [1, C]$, and hence:

$$\forall \mathfrak{b} \in \mathcal{I}_{A,B} : \frac{p_{\mathfrak{b}}}{\mathcal{U}(\mathcal{I}_{A,B})(\mathfrak{b})} \in [C^{-1}, C],$$

hence $\text{RD}_{\infty}(\mathcal{U}(\mathcal{I}_{A,B}) \parallel \mathcal{D}_{\text{Ideal}}(X)) \leq C$, which complete the proof. \square

IV.3.2 From an arbitrary ideal to a uniform integral ideal

Below, we give an algorithm, `RandomizeIdeal` _{A, X} (see Algorithm IV.3.2), which on input an arbitrary ideal I , returns a uniform integral ideal \mathfrak{b} and a short non-zero vector $y \in \mathfrak{b}^{-1} \cdot I^{-1}$. The algorithm is parameterized by an integer A and a set X satisfying the conditions of Lemma IV.3.2. `RandomizeIdeal` _{A, X} starts by sampling a uniform norm-1 ideal J , i.e., with distribution equal to $\mathcal{U}(\text{IdLat}_K^0)$, along with a small element v_J in it, using the `SampleWithTrap` algorithm. Since $\mathcal{U}(\text{IdLat}_K^0)$ is the Haar distribution on a compact group, the ideal $I' = J \cdot (I/\mathcal{N}(I)^{1/d})$ is also uniform. We then use `IdealRound` to map $\mathcal{U}(\text{IdLat}_K^0)$ to the uniform distribution over integral ideals with norms in $[A, 4A]$. In more details, a uniform point x in $I' \cap X$ is sampled and Lemma IV.3.2 implies that $\mathfrak{b} := x \cdot I'^{-1}$ is almost uniform, and $v_J \cdot x^{-1}$ is a small element in $\mathfrak{b}^{-1} \cdot \mathcal{N}(I)^{1/d} \cdot I^{-1}$ if x is balanced. We note that Steps 7 and 8 below are exactly the `IdealRound` algorithm applied to the ideal I' . However, we cannot call this algorithm in a blackbox way, as we need to know the intermediate value x for Step 9 of the algorithm.

Algorithm IV.3.2 RandomizeIdeal**Input:** A basis B_I of an ideal I .**Parameters:** A integer and $X \subset K_{\mathbb{R}} \setminus \{0\}$ compact.**Oracles:** \mathcal{F} for factoring integral ideals, \mathcal{S} for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{IdLat}_K^0$.**Output:** \mathfrak{b} an integral ideal, $y \in \mathfrak{b}^{-1} \cdot I^{-1} \setminus \{0\}$.

- 1: Sample $(\mathfrak{q}, v_{\mathfrak{q}}) \leftarrow \text{SampleWithTrap}_{A,4A}(B_{\mathcal{O}_K})$, using \mathcal{F} .
- 2: Sample $\zeta \leftarrow \mathcal{G}(0, d^{-3/2})$ in $\text{span}(\text{Ln}(\mathcal{O}_K^{\times}))$ conditioned on $\|\zeta\| \leq 1/d$.
- 3: Sample u uniform in $\{x \in K_{\mathbb{R}}^{\times} : \forall i \leq d, |x_i| = 1\}$.
- 4: Let $J = u \cdot \text{Exp}(\zeta) \cdot \mathcal{N}(\mathfrak{q})^{-1/d} \cdot \mathfrak{q}$ and $v_J = u \cdot \text{Exp}(\zeta) \cdot \mathcal{N}(\mathfrak{q})^{-1/d} \cdot v_{\mathfrak{q}}$.
- 5: Compute $B_J = \text{ReduceIdeal}(J, v_J)$.
- 6: Let $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$ and $B_{I'} = \text{MultiplyIdeals}(B_J, \mathcal{N}(I)^{-1/d} \cdot B_I)$.
- 7: Sample $x \leftarrow \mathcal{U}(I' \cap X)$, using \mathcal{S} .
- 8: Let $\mathfrak{b} = x \cdot I'^{-1}$.
- 9: Let $y = x^{-1} \cdot \mathcal{N}(I)^{-1/d} \cdot v_J$.
- 10: Return (\mathfrak{b}, y) .

$$\begin{array}{ccccccc}
D_1 & \xleftarrow{\text{SD}=2^{-\Omega(d)}} & D_2 & \xrightarrow{\text{RD}_{\infty}=O(1)} & D_3 & \xleftarrow{\text{SD}=2^{-\Omega(d)}} & D_4 \\
\text{IdealRound}(\cdot) \downarrow & & \downarrow & & \downarrow & & \downarrow \\
D = \widetilde{D}_1 & \xleftarrow{\text{SD}=2^{-\Omega(d)}} & \widetilde{D}_2 & \xrightarrow{\text{RD}_{\infty}=O(1)} & \widetilde{D}_3 & \xleftarrow{\text{SD}=2^{-\Omega(d)}} & \widetilde{D}_4 \xrightarrow{\text{RD}_{\infty}=O(1)} \mathcal{U}(\mathcal{I}_{A,4A})
\end{array}$$

Figure IV.2: Relations between the distributions of the proof of Lemma IV.3.3.

Lemma IV.3.3. *Let $A \geq \max(\delta_K^d, d^d \Delta_K^c)$ for c as in Lemma II.2.12. Let X be a compact subset of $K_{\mathbb{R}} \setminus \{0\}$ whose elements are η -balanced for some $\eta > 1$ and satisfy the assumptions of Lemma IV.3.2 for A and $B = 4A$. Assume that $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}| \leq c'$ for some $c' < 1$. On input a basis B_I of an ideal I , $\text{RandomizeIdeal}_{A,X}$ runs in time polynomial in $\log A$, $\log \Delta_K$, $A/|\mathcal{P}_{A,4A}|$ and the size of its input, and returns (\mathfrak{b}, y) satisfying*

$$\begin{aligned}
\mathfrak{b} &\in \mathcal{I}_{A,4A}, \\
y &\in \mathfrak{b}^{-1} I^{-1} \setminus \{0\}, \\
\|y\| &\leq 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathfrak{b})^{-1/d}.
\end{aligned}$$

Finally, if D and \mathcal{U} respectively denote the distribution of \mathfrak{b} and the uniform distribution over $\mathcal{I}_{A,4A}$, then the following holds for any event $E \subseteq \mathcal{I}_{A,4A}$:

$$D(E) \geq \frac{\mathcal{U}(E)}{\Theta(1)} - 2^{-\Omega(d)}.$$

Proof. We first bound the Euclidean norms of the variables occurring during the execution of the algorithm. By Lemma IV.2.3 and the assumption that $A \geq \delta_K^d$, we have that $0 < \|v_{\mathfrak{q}}\| \leq 51 \cdot d \cdot (A\Delta_K)^{1/d}$. Now, note that $\|u\|_{\infty} = 1$, $\|\text{Exp}(\zeta)\|_{\infty} \leq \exp(1/2)$ and $\mathcal{N}(\mathfrak{q})^{-1/d} \leq A^{-1/d}$. We then have $\|v_J\| \leq 85 \cdot d \cdot \Delta_K^{1/d}$ (and $v_J \neq 0$). Then, by Lemma II.2.15, we have $0 < \|B_J\| \leq 85 \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/d}$ and

$$\|B_{I'}\| \leq 85 \cdot d^2 \cdot \delta_K \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I)^{-1/d} \cdot \|B_I\|.$$

As elements of X are non-zero and η -balanced, we have that $\|x^{-1}\|_\infty \leq \eta \cdot \mathcal{N}(x)^{-1/d}$. Also, note that since $\mathcal{N}(I') = 1$, we have $\mathcal{N}(\mathbf{b}) = \mathcal{N}(x)$. As a result, we obtain that $y \neq 0$ and:

$$\begin{aligned} \|y\| &\leq \mathcal{N}(I)^{-1/d} \cdot \|x^{-1}\|_\infty \cdot \|v_J\| \\ &\leq \mathcal{N}(I)^{-1/d} \cdot \eta \cdot \mathcal{N}(x)^{-1/d} \cdot 85 \cdot d \cdot \Delta_K^{1/d} \\ &= 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathbf{b})^{-1/d}. \end{aligned}$$

The latter and the fact that $\mathcal{N}(\mathbf{b}) = \mathcal{N}(x)$ belongs to $[A, 4A]$ (by assumption on X) provide the first statement on the output.

The previous computations show that every quantity manipulated by the algorithm has size polynomial in $\log A$, $\log \Delta_K$ and the bit-size of the input. Note that `SampleWithTrap` _{$A, 4A$} runs in polynomial time in $A/|\mathcal{P}_{A, 4A}|$. The overall running time is then polynomial in $\log A$, $\log \Delta_K$, $A/|\mathcal{P}_{A, 4A}|$ and the size of the input.

We now analyze the distribution of \mathbf{b} . For this purpose, we define the following distributions (see also Figure IV.2):

- D_1 is the distribution of J at Step 4;
- D_2 is the distribution $u \cdot \text{Exp}(\zeta) \cdot \mathbf{q} \cdot \mathcal{N}(\mathbf{q})^{-1/d}$ where \mathbf{q} is uniform in $\mathcal{P}_{A, 4A}$, and u, ζ are sampled as in Steps 2 and 3;
- D_3 is the same as D_2 but with \mathbf{q} uniform in $\mathcal{P}_{0, 4A}$;
- D_4 is $\mathcal{U}(\text{IdLat}_K^0)$.

Note that we have the following relationships between the D_i 's:

- $\text{SD}(D_1, D_2) = 2^{-\Omega(d)}$, thanks to Lemma IV.2.3 and the data processing inequality;
- $\text{RD}_\infty(D_3 \parallel D_2) = \Theta(1)$, thanks to the assumption on $|\mathcal{P}_{0, A}|/|\mathcal{P}_{0, 4A}|$;
- $\text{SD}(D_3, D_4) = 2^{-\Omega(d)}$ thanks to Lemma II.2.12.

We also define \widetilde{D}_i (for $i \leq 4$) as the distribution of \mathbf{b} obtained by sampling J from D_i , setting $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$, sampling x from $\mathcal{U}(I' \cap X)$ and returning $x \cdot I'^{-1}$. Note that \widetilde{D}_1 is D and that \widetilde{D}_4 is $\mathcal{D}_{\text{Ideal}}(X)$. Indeed, as $\mathcal{U}(\text{IdLat}_K^0)$ is invariant by multiplication by a fixed norm-1 replete ideal, the ideal $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$ is then distributed from $\mathcal{U}(\text{IdLat}_K^0)$. The data-processing inequalities of the statistical distance and Rényi divergence imply that the above relations also hold for \widetilde{D}_i in place of D_i , for all i . Furthermore, by choice of X , the Rényi divergence from $\mathcal{U}(\mathcal{I}_{A, 4A})$ to \widetilde{D}_4 is equal to $\Theta(1)$.

Using the probability preservation properties of the statistical distance and Rényi divergence, we obtain that for any event $E \subseteq \mathcal{I}_{A, 4A}$, we have:

$$\widetilde{D}_1(E) \geq \frac{\mathcal{U}(E) - 2^{-\Omega(d)}}{\Theta(1)} - 2^{-\Omega(d)} = \frac{\mathcal{U}(E)}{\Theta(1)} - 2^{-\Omega(d)}$$

which completes the proof. □

IV.3.3 From ideal to their inverses

Let \mathcal{W} be a set of fractional ideals. Below, we reduce \mathcal{W}^{-1} -avg-id-HSVP to \mathcal{W} -avg-id-HSVP and $\mathcal{I}_{A,4A}$ -avg-id-HSVP for some appropriate integer A and approximation factors. Recall that \mathcal{W}^{-1} refers to the set $\{I^{-1}, I \in \mathcal{W}\}$.

The reduction is described as an algorithm, $\text{InverseToIntegral}_{A,X}^{\mathcal{W}}$ (Algorithm IV.3.3), which takes as input the inverse I^{-1} of an integral ideal $I \in \mathcal{W}$ and returns a short non-zero element of I^{-1} . It is parameterized by an integer A and a compact set X satisfying the conditions of Lemma IV.3.2. It relies on four oracles: oracle $\mathcal{O}_{\mathcal{W}}$ for solving \mathcal{W} -avg-id-HSVP, oracle $\mathcal{O}_{\mathcal{I}}$ for $\mathcal{I}_{A,4A}$ -avg-id-HSVP, oracle \mathcal{F} for factoring integral ideals; and oracle \mathcal{S} for sampling from $I \cap X$ for $I \in \text{IdLat}_K^0$. Recall that \mathcal{F} can be instantiated as a quantum polynomial time algorithm. An instantiation of oracle \mathcal{S} will be provided in Section IV.4, based on the design of a nice set X for Lemma IV.3.2. The reduction first uses $\mathcal{O}_{\mathcal{W}}$ on the inverse I of its input, which gives a short non-zero vector $v_I \in I$. Then $\text{RandomizeIdeal}_{A,X}$ (introduced in the previous subsection) is invoked to randomize I into a uniform integral ideal \mathfrak{b} with norm in $[A, 4A]$. $\text{RandomizeIdeal}_{A,X}$ also returns a short non-zero $y_{(I\mathfrak{b})^{-1}}$ in $(I\mathfrak{b})^{-1}$. Then $\mathcal{O}_{\mathcal{I}}$ is invoked on \mathfrak{b} and returns a short non-zero $v_{\mathfrak{b}}$ in \mathfrak{b} . The reduction finally outputs $v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}} \in I^{-1}$ that is short and non-zero.

The astute reader will notice that, in the above description, the vector v_I and hence the oracle $\mathcal{O}_{\mathcal{W}}$ do not seem to be used in the subsequent steps. In fact, we will be able to instantiate \mathcal{S} only if given a short basis of I (see Lemma IV.4.9). The approximation factor reached by $\mathcal{O}_{\mathcal{W}}$ will lead to a lower bound condition on A : the smaller the approximation factor, the smaller the lower bound on A .

Algorithm IV.3.3 $\text{InverseToIntegral}_{A,X}^{\mathcal{W}}$

Input: I^{-1} with $I \in \mathcal{W}$.

Parameters: A integer and $X \subset K_{\mathbb{R}} \setminus \{0\}$ compact.

Oracles: $\mathcal{O}_{\mathcal{W}}$ for \mathcal{W} -avg-id-HSVP $_{\gamma_{\mathcal{W}}}$, $\mathcal{O}_{\mathcal{I}}$ for $\mathcal{I}_{A,4A}$ -avg-id-HSVP $_{\gamma_{\mathcal{I}}}$,

\mathcal{F} for factoring integral ideals and \mathcal{S} for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{IdLat}_K^0$.

Output: $x \in I^{-1} \setminus \{0\}$.

- 1: Compute $v_I \leftarrow \mathcal{O}_{\mathcal{W}}(I)$.
 - 2: If $v_I = \perp$, then return \perp .
 - 3: Compute $\mathbf{B}_I = \text{ReduceIdeal}(I, v_I)$.
 - 4: Sample $(\mathfrak{b}, y_{(I\mathfrak{b})^{-1}}) \leftarrow \text{RandomizeIdeal}_{A,X}(\mathbf{B}_I)$, using \mathcal{F} and \mathcal{S} .
 - 5: Compute $v_{\mathfrak{b}} \leftarrow \mathcal{O}_{\mathcal{I}}(\mathfrak{b})$.
 - 6: If $v_{\mathfrak{b}} = \perp$, then return \perp .
 - 7: Return $v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}}$.
-

Theorem IV.3.4. *Let \mathcal{W} be a finite set of fractional ideals. Let $\gamma_{\mathcal{W}}, \gamma_{\mathcal{I}} \geq 1$ and A satisfying $A \geq \max(\delta_K^d, d^d \Delta_K^c)$ for c as in Lemma II.2.12. Let X be a compact subset of $K_{\mathbb{R}} \setminus \{0\}$ whose elements are η -balanced for some $\eta > 1$ and satisfy the assumptions of Lemma IV.3.2 for A and $B = 4A$. Assume that $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}| \leq c'$ for some constant $c' < 1$. Let $\mathcal{O}_{\mathcal{W}}$ an oracle for \mathcal{W} -avg-id-HSVP $_{\gamma_{\mathcal{W}}}$ with success probability $\varepsilon_{\mathcal{W}}$ and $\mathcal{O}_{\mathcal{I}}$ an oracle for $\mathcal{I}_{A,4A}$ -avg-id-HSVP $_{\gamma_{\mathcal{I}}}$ with success probability $\varepsilon_{\mathcal{I}}$.*

When given access to $\mathcal{O}_{\mathcal{W}}$, $\mathcal{O}_{\mathcal{I}}$, an integral ideal-factoring oracle \mathcal{F} and an oracle \mathcal{S} for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{IdLat}_K^0$, $\text{InverseToIntegral}_{A,X}^{\mathcal{W}}$ runs in expected time polynomial in $\log A, \log \Delta_K, A/|\mathcal{P}_{A,4A}|$ and the size of its input. Further, if its input I is such that I is distributed from $\mathcal{U}(\mathcal{W})$, it outputs $x \neq \perp$ with probability $\geq \varepsilon_{\mathcal{I}} \cdot (\varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)})$. If $x \neq \perp$,

then we have

$$x \in I^{-1} \setminus \{0\} \quad \text{and} \quad \|x\| \leq \gamma' \cdot \text{Vol}(I^{-1})^{1/d},$$

for $\gamma' = 85 \cdot \gamma_{\mathcal{I}} \cdot \Delta_K^{1/d} \cdot d \cdot \eta$.

Proof. Assume first that neither v_I nor $v_{\mathfrak{b}}$ is equal to \perp . As the assumptions of Lemma IV.3.3 are satisfied, we have $y_{(I\mathfrak{b})^{-1}} \in (I\mathfrak{b})^{-1} \setminus \{0\}$ and

$$\|y_{(I\mathfrak{b})^{-1}}\| \leq 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathfrak{b})^{-1/d}.$$

Now, by assumption, we have that $v_{\mathfrak{b}} \in \mathfrak{b} \setminus \{0\}$ satisfies $\|v_{\mathfrak{b}}\| \leq \gamma_{\mathcal{I}} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{b})^{1/d}$. We then obtain that $x = v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}} \in I^{-1}$ is non-zero and satisfies:

$$\|x\| \leq \|v_{\mathfrak{b}}\| \cdot \|y_{(I\mathfrak{b})^{-1}}\| \leq \gamma_{\mathcal{I}} \cdot \Delta_K^{3/(2d)} \cdot 85 \cdot d \cdot \eta \cdot \mathcal{N}(I)^{-1/d}.$$

Towards completing the proof, note that the algorithm succeeds if and only if neither v_I nor $v_{\mathfrak{b}}$ is equal to \perp . The probability that v_I is not \perp is exactly $\varepsilon_{\mathcal{I}}$. Using Lemma IV.3.3 with the event E set to $\mathcal{O}_{\mathcal{I}}(\mathfrak{b})$ succeeding, we obtain that $v_{\mathfrak{b}}$ is not \perp with probability $\geq \varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)}$. Note that the second probability is over the internal randomness of $\text{RandomizeIdeal}_{A,X}(\mathbf{B}_I)$. \square

IV.4 The Sampling Set

Lemma IV.3.2 states that if a compact X satisfies a certain number of conditions, then the output distribution of IdealRound_X resembles the uniform distribution over integral ideals whose norms belong to a prescribed interval. In this subsection, we show that the set $\mathcal{B}_{A,B}^{\eta}$ defined below satisfies those constraints. We will later also use the fact that its elements are η -balanced. An instantiation of the set $\mathcal{B}_{A,B}^{\eta}$ can be visualized in Figure IV.1.

Definition IV.4.1. Let $B > A > 0$ and $\eta > 1$. We define the set:

$$\mathcal{B}_{A,B}^{\eta} = \left\{ x \in K_{\mathbb{R}}^{\times} \mid \mathcal{N}(x) \in [A, B], \left\| \text{Ln} \left(\frac{x}{\mathcal{N}(x)^{1/d}} \right) \right\|_2 \leq \ln(\eta) \right\}.$$

The purpose of this section is to prove the following theorem.

Theorem IV.4.2. Let $A, B, \eta, \delta > 0$ satisfying $A^{1/d} \geq d^3 \cdot \eta \cdot \max(\Delta_K^{3/(2d)}, \delta)$, $B/A \geq 4$ and $\eta \geq e$. The set $\mathcal{B}_{A,B}^{\eta}$ is compact and invariant by complex rotations, satisfies the conditions of Lemma IV.3.2 and its elements are η -balanced. Furthermore, there exists an algorithm $\text{SampleUniform}_{A,B}^{\eta}$ that, given as input a basis \mathbf{B}_I of a norm-1 replete ideal satisfying $\|\mathbf{B}_I^*\| \leq \delta$, samples uniformly in $I \cap \mathcal{B}_{A,B}^{\eta}$ and whose expected running time is polynomial in $\log B$, d and B/A .

IV.4.1 Volume of the set $\mathcal{B}_{A,B}^{\eta}$

Before proving that the assumptions of Lemma IV.3.2 are satisfied by $\mathcal{B}_{A,B}^{\eta}$, we first study its volume and its approximate invariance under translation.

Lemma IV.4.3. For any $B > A > 0$ and $\eta > 1$, we have

$$\text{Vol}(\mathcal{B}_{A,B}^{\eta}) = \frac{2^{d_{\mathbb{R}}} \cdot (2\sqrt{2}\pi)^{d_{\mathbb{C}}} \cdot V_{d_{\mathbb{R}}+d_{\mathbb{C}}-1}}{\sqrt{d}} \cdot (B - A) \cdot (\ln \eta)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1},$$

where V_n is the volume of the n -dimensional unit ℓ_2 hyperball for any $n \geq 1$.

The computation of the volume proceeds by a change of variable, between \mathbb{R}^d and $\Phi(K_{\mathbb{R}})$. The relevant aspect of the volume formula for the present work is the linear dependency in $(B - A) \cdot (\ln \eta)^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1}$.

Proof. Let \mathbf{e}_j denote the j -th elementary unit vector, we fix $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_d) \in \mathbb{C}^{d \times d}$ an orthonormal \mathbb{R} -basis of $\Phi(K_{\mathbb{R}}) \subset \mathbb{C}^d$ defined by

$$\begin{aligned} \mathbf{c}_j &= \mathbf{e}_j && \text{for } 1 \leq j \leq d_{\mathbb{R}}, \\ \mathbf{c}_{d_{\mathbb{R}}+j} &= 1/\sqrt{2} \cdot (\mathbf{e}_{d_{\mathbb{R}}+j} + \mathbf{e}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}) && \text{for } 1 \leq j \leq d_{\mathbb{C}}, \\ \mathbf{c}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j} &= i/\sqrt{2} \cdot (\mathbf{e}_{d_{\mathbb{R}}+j} - \mathbf{e}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}) && \text{for } 1 \leq j \leq d_{\mathbb{C}}. \end{aligned}$$

We let ϕ be the isomorphism sending an element of $\Phi(K_{\mathbb{R}})$ to its coordinates in the basis \mathbf{C} . Since \mathbf{C} is orthonormal, the map ϕ preserves the geometry. In particular, the volume of $\mathcal{B}_{A,B}^{\eta}$ is the same as the volume of $\phi(\mathcal{B}_{A,B}^{\eta})$ (which is a d -dimensional object in \mathbb{R}^d).

In order to compute this volume, we first introduce the set

$$\mathcal{B}_{A,B}^{\eta+} = \left\{ x \in \mathcal{B}_{A,B}^{\eta} \mid \sigma_i(x) > 0 \text{ for all } 1 \leq i \leq d_{\mathbb{R}} \right\},$$

i.e., the elements of $\mathcal{B}_{A,B}^{\eta}$ whose real embeddings are all positive. Since $\mathcal{B}_{A,B}^{\eta}$ is invariant by complex rotations, the set $\mathcal{B}_{A,B}^{\eta+}$ is the union of $2^{d_{\mathbb{R}}}$ distinct copies of $\mathcal{B}_{A,B}^{\eta+}$, hence we can focus on computing the volume of the latter. In order to compute this volume, we will exhibit a function F transforming a ‘‘nice box’’ of \mathbb{R}^d into the set $\phi(\mathcal{B}_{A,B}^{\eta+})$. We will then use this function to perform a change of variable and compute the volume of $\phi(\mathcal{B}_{A,B}^{\eta+})$ from the volume of the nice box.

Defining the function F . Let H be the $(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)$ -dimensional subspace of \mathbb{R}^d spanned by $\text{Ln}(\mathcal{O}_{K_{\mathbb{R}}}^{\times})$, i.e.,

$$H = \left\{ x \in \mathbb{R}^d \mid \sum_{j \leq d} x_j = 0 \wedge \forall d_{\mathbb{R}} < j \leq d_{\mathbb{R}} + d_{\mathbb{C}} : x_{j+d_{\mathbb{C}}} = x_j \right\},$$

and let $\mathbf{B} = (b_{i,j}) \in \mathbb{R}^{d \times (d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}$ be any orthonormal basis of H . We define the following function f from $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times \mathbb{R}^{d_{\mathbb{C}}}$ to $\Phi(K_{\mathbb{R}})$ as:

$$f(N, \mathbf{z}, \boldsymbol{\theta}) = \exp(N/d) \cdot \exp(\mathbf{B} \cdot \mathbf{z} + i\hat{\boldsymbol{\theta}}),$$

where the second function \exp is applied coordinate-wise to the vector $\mathbf{B} \cdot \mathbf{z} + i\hat{\boldsymbol{\theta}}$, and where $\hat{\boldsymbol{\theta}} = (\mathbf{0}^{d_{\mathbb{R}}} \mid \boldsymbol{\theta}^T \mid -\boldsymbol{\theta}^T)^T \in \mathbb{R}^d$. Note that f is injective on the set $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$, and that its image indeed lies in $\Phi(K_{\mathbb{R}})$ (it even lies in the subset of $\Phi(K_{\mathbb{R}})$ whose first $d_{\mathbb{R}}$ coordinates are positive).

In order to obtain a transformation from \mathbb{R}^d to itself, we compose the above function f with the function ϕ , and we obtain $F = \phi \circ f : \mathbb{R}^d \rightarrow \mathbb{R}^d$, which is injective on $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$. Moreover, by letting $\text{Ball}^{(n)}(R)$ denote the Euclidean ball of radius R in \mathbb{R}^n , we have that

$$\phi(\mathcal{B}_{A,B}^{\eta+}) = F\left([\ln A, \ln B] \times \text{Ball}^{(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}\right).$$

Indeed, let $(N, \mathbf{z}, \boldsymbol{\theta}) \in \mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$ and let $x = \Phi^{-1}(f(N, \mathbf{z}, \boldsymbol{\theta})) \in K_{\mathbb{R}}$. Then $\mathcal{N}(x) = \exp(N)$ (because $\mathbf{B} \cdot \mathbf{z}$ belongs to H , so the sum of its coordinates is zero) and $\text{Ln}(x/\mathcal{N}(x)^{1/d}) = \mathbf{B} \cdot \mathbf{z}$, which implies that $\|\text{Ln}(x/\mathcal{N}(x)^{1/d})\|_2 = \|\mathbf{z}\|_2$ since \mathbf{B} is orthonormal. The inclusion from right to left follows from these two observations and the definition of $\mathcal{B}_{A,B}^{\eta}$. For the inclusion

from left to right, it suffices to observe that a pre-image of $x \in \mathcal{B}_{A,B}^{\eta+}$ is obtained by taking $N = \ln(\mathcal{N}(x))$, \mathbf{z} equal to the coordinates of $\text{Ln}(x/\mathcal{N}(x)^{1/d})$ in basis \mathbf{B} , and $\boldsymbol{\theta}$ equal to the arguments of $\sigma_i(x)$.

The set $[\ln A, \ln B] \times \text{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$ is the “nice set” we mentioned above. To compute the volume of $\mathcal{B}_{A,B}^{\eta+}$, we will change variables, using the function F , in order to transform $\phi(\mathcal{B}_{A,B}^{\eta+})$ into this nice set. In order to perform this change of variables, we first compute the Jacobian matrix of F and its determinant.

Computing the Jacobian matrix of F . For $1 \leq i \leq d$, let F_i be the function corresponding to the i -th coordinate of F (note that F_i goes from \mathbb{R}^d to \mathbb{R}). By definition of F and choice of \mathbf{C} (which defines ϕ), one can check that the following holds, for $(N, \mathbf{z}, \boldsymbol{\theta}) \in \mathbb{R}^d$:

$$\begin{aligned} F_i(N, \mathbf{z}, \boldsymbol{\theta}) &= \exp\left(N/d + \sum_j b_{i,j} z_j\right) && \text{for } 1 \leq i \leq d_{\mathbb{R}}, \\ F_{d_{\mathbb{R}}+i}(N, \mathbf{z}, \boldsymbol{\theta}) &= \sqrt{2} \cdot \exp\left(N/d + \sum_j b_{d_{\mathbb{R}}+i,j} z_j\right) \cdot \cos(\theta_i) && \text{for } 1 \leq i \leq d_{\mathbb{C}}, \\ F_{d_{\mathbb{R}}+d_{\mathbb{C}}+i}(N, \mathbf{z}, \boldsymbol{\theta}) &= \sqrt{2} \cdot \exp\left(N/d + \sum_j b_{d_{\mathbb{R}}+d_{\mathbb{C}}+i,j} z_j\right) \cdot \sin(\theta_i) && \text{for } 1 \leq i \leq d_{\mathbb{C}}. \end{aligned}$$

Note that here, we used the fact that for all $1 \leq j \leq d_{\mathbb{R}} + d_{\mathbb{C}} - 1$ and all $1 \leq i \leq d_{\mathbb{C}}$, it holds that $b_{d_{\mathbb{R}}+i,j} = b_{d_{\mathbb{R}}+d_{\mathbb{C}}+i,j}$ since the columns of \mathbf{B} are in H . We then obtain that:

$$\begin{aligned} \partial_N(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= F_i(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \frac{1}{d} && \text{for } i \leq d, \\ \partial_{z_j}(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= F_i(N, \mathbf{z}, \boldsymbol{\theta}) \cdot b_{i,j} && \text{for } i \leq d \text{ and } j \leq d_{\mathbb{R}} + d_{\mathbb{C}} - 1, \\ \partial_{\theta_j}(F_{d_{\mathbb{R}}+j}(N, \mathbf{z}, \boldsymbol{\theta})) &= F_{d_{\mathbb{R}}+j}(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \frac{-\sin \theta_j}{\cos \theta_j} && \text{for } j \leq d_{\mathbb{C}}, \\ \partial_{\theta_j}(F_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}(N, \mathbf{z}, \boldsymbol{\theta})) &= F_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \frac{\cos \theta_j}{\sin \theta_j} && \text{for } j \leq d_{\mathbb{C}}, \\ \partial_{\theta_j}(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= 0 && \text{else.} \end{aligned}$$

In short, the Jacobian matrix of F is $D_F(N, \mathbf{z}, \boldsymbol{\theta}) = \text{diag}_i(F_i(N, \mathbf{z}, \boldsymbol{\theta})) \cdot \mathbf{M}$, where

$$\mathbf{M} = \left(\begin{array}{c|c|c} \frac{1}{d} & & \mathbf{0} \\ \vdots & \mathbf{B} & \text{diag}_i\left(\frac{-\sin \theta_i}{\cos \theta_i}\right) \\ \frac{1}{d} & & \text{diag}_i\left(\frac{\cos \theta_i}{\sin \theta_i}\right) \end{array} \right).$$

Computing the Jacobian determinant. We now compute the determinant of the matrix $D_F(N, \mathbf{z}, \boldsymbol{\theta})$. First, we have

$$\det\left(\text{diag}_i\left(F_i(N, \mathbf{z}, \boldsymbol{\theta})\right)\right) = \exp(N) \cdot (\sqrt{2})^{2d_{\mathbb{C}}} \cdot \prod_{i \leq d_{\mathbb{C}}} (\sin \theta_i \cdot \cos \theta_i),$$

where we used again the fact that the sum of the coordinates of $\mathbf{B} \cdot \mathbf{z}$ is zero. We now focus on \mathbf{M} . Let $\widehat{\mathbf{M}}_{i,j}$ be the matrix \mathbf{M} where we have removed the i -th line and j -th column. Observe that for $i \leq d_{\mathbb{C}}$, the $(d_{\mathbb{R}} + i)$ -th row of \mathbf{M} and the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$ -th row of \mathbf{M} coincide except in the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$ -th column. So developing the determinant of \mathbf{M} on the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$ -th column leads to

$$\begin{aligned} |\det \mathbf{M}| &= \left| \frac{\sin \theta_i}{\cos \theta_i} + \frac{\cos \theta_i}{\sin \theta_i} \right| \cdot \left| \det \widehat{\mathbf{M}}_{d_{\mathbb{R}}+d_{\mathbb{C}}+i, d_{\mathbb{R}}+d_{\mathbb{C}}+i} \right| \\ &= \frac{1}{|\sin \theta_i \cdot \cos \theta_i|} \cdot \left| \det \widehat{\mathbf{M}}_{d_{\mathbb{R}}+d_{\mathbb{C}}+i, d_{\mathbb{R}}+d_{\mathbb{C}}+i} \right|. \end{aligned}$$

Here, to check that the signs are correct, we observe that we can permute the rows of \mathbf{M} without changing the absolute value of the determinant, and move the row with index $d_{\mathbb{R}} + i$ to position $d_{\mathbb{R}} + d_{\mathbb{C}} + i + 1$ (so that it follows directly the row $d_{\mathbb{R}} + d_{\mathbb{C}} + i$). This ensures that both minor matrices are the same, and that the signs are opposite when we develop according to the $(d_{\mathbb{R}} + d_{\mathbb{R}} + i)$ -th column. Repeating the process on the $d_{\mathbb{C}}$ last columns of \mathbf{M} , we obtain that

$$|\det \mathbf{M}| = \left(\prod_i \frac{1}{|\sin \theta_i \cdot \cos \theta_i|} \right) \cdot |\det \widehat{\mathbf{M}}|,$$

where $\widehat{\mathbf{M}}$ is the top-left square sub-matrix of \mathbf{M} of dimension $d_{\mathbb{R}} + d_{\mathbb{C}}$. Let $\mathbf{B}_0 \in \mathbb{R}^{d_{\mathbb{R}} \times (d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}$ and $\mathbf{B}_1 \in \mathbb{R}^{d_{\mathbb{C}} \times (d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}$ be sub-blocks of the matrix \mathbf{B} such that $\mathbf{B} = (\mathbf{B}_0^T | \mathbf{B}_1^T | \mathbf{B}_1^T)^T$ (recall that \mathbf{B} is an arbitrary orthonormal basis of H). Then all the entries of the first column of $\widehat{\mathbf{M}}$ are equal to $1/d$ and the remaining $d_{\mathbb{R}} + d_{\mathbb{C}} - 1$ are $(\mathbf{B}_0^T | \mathbf{B}_1^T)^T$. Let us consider the following distortion of $\widehat{\mathbf{M}}$:

$$\mathbf{N} = \left(\begin{array}{c|c} \frac{1}{d} \cdot \mathbf{1}_{d_{\mathbb{R}}} & \mathbf{B}_0 \\ \hline \frac{\sqrt{2}}{d} \cdot \mathbf{1}_{d_{\mathbb{C}}} & \sqrt{2} \cdot \mathbf{B}_1 \end{array} \right),$$

where $\mathbf{1}_k$ refers to the k -dimensional all-1 vector. Then $\det \widehat{\mathbf{M}} = \sqrt{2}^{-d_{\mathbb{C}}} \cdot \det \mathbf{N}$. Furthermore, note that $\mathbf{N}^T \cdot \mathbf{N} = \text{diag}(1/d, 1, \dots, 1)$, because the columns of \mathbf{B} are orthonormal and in H (so the sums of their coordinates are zero). This gives us that $|\det \mathbf{N}| = 1/\sqrt{d}$. Unrolling the above, we obtain

$$|\det \mathbf{M}| = \frac{1}{\sqrt{d} \cdot \sqrt{2}^{d_{\mathbb{C}}} \cdot \prod_i |\sin \theta_i \cdot \cos \theta_i|},$$

and

$$|\det(D_F(N, \mathbf{z}, \boldsymbol{\theta}))| = \frac{\exp(N) \cdot \sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}}.$$

Change of variables. We finally perform the change of variables using the function F to compute the volume of $\mathcal{B}_{A,B}^{\eta+}$ (recall that $\text{Vol}(\mathcal{B}_{A,B}^{\eta}) = 2^{d_{\mathbb{R}}} \cdot \text{Vol}(\mathcal{B}_{A,B}^{\eta+})$). Letting $\mathbf{1}_S(\cdot)$ denote the indicator function of a set S , we have

$$\begin{aligned} \text{Vol}(\mathcal{B}_{A,B}^{\eta+}) &= \int_{\mathbf{x} \in \mathbb{R}^d} \mathbf{1}_{\phi(\mathcal{B}_{A,B}^{\eta+})}(\mathbf{x}) \, d\mathbf{x} \\ &= \int_{\substack{N \in [\ln A, \ln B] \\ \mathbf{z} \in \text{Ball}^{(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}(\ln \eta) \\ \boldsymbol{\theta} \in [0, 2\pi]^{d_{\mathbb{C}}}}} |\det(D_F(N, \mathbf{z}, \boldsymbol{\theta}))| \, d\boldsymbol{\theta} \, d\mathbf{z} \, dN \\ &= \frac{\sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}} \cdot \int_{N \in [\ln A, \ln B]} \exp(N) \, dN \cdot \int_{\mathbf{z} \in \text{Ball}^{(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}(\ln \eta)} d\mathbf{z} \cdot \int_{\boldsymbol{\theta} \in [0, 2\pi]^{d_{\mathbb{C}}}} d\boldsymbol{\theta} \\ &= \frac{\sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}} \cdot (B - A) \cdot V_{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \cdot (\ln \eta)^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \cdot (2\pi)^{d_{\mathbb{C}}}, \end{aligned}$$

as desired. □

The proof of Lemma IV.4.3 gives us the volume of the set $\mathcal{B}_{A,B}^{\eta}$, but it also a way to sample uniformly in it.

Lemma IV.4.4. *There exists a probabilistic algorithm that samples from $\mathcal{U}(\mathcal{B}_{A,B}^{\eta})$ for any $B > A > 0$ and $\eta > 1$. The expected running time of this algorithm is polynomial in $\log B$, d (the degree of K) and B/A .*

Proof. Let ϕ and F be the same functions as in the proof of Lemma IV.4.3. Recall that

$$F\left([\ln A, \ln B] \times \text{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}\right) = \phi(\mathcal{B}_{A,B}^{\eta,+}), \quad (\text{IV.1})$$

and that F is injective on this set. It can be observed from their definitions that F , ϕ and ϕ^{-1} can be computed in time polynomial in d .

Note that if we can sample from $\mathcal{U}(\phi(\mathcal{B}_{A,B}^{\eta,+}))$ in time T , then we can sample from $\mathcal{U}(\mathcal{B}_{A,B}^{\eta,+})$ in time $T + \text{poly}(d)$. Indeed, it suffices to sample x from $\mathcal{U}(\phi(\mathcal{B}_{A,B}^{\eta,+}))$; compute $\phi^{-1}(x)$ (which can be done in time $\text{poly}(d)$); sample uniform signs $(\varepsilon_i)_i \in \{-1, 1\}^{d_{\mathbb{R}}}$; and finally output $\phi^{-1}(x) \cdot \Phi^{-1}((\varepsilon_1, \dots, \varepsilon_{d_{\mathbb{R}}}, 1, \dots, 1))$. In the rest of this proof, we then focus on sampling the random variable $\mathcal{U}(\phi(\mathcal{B}_{A,B}^{\eta,+}))$.

Let Y be a random variable distributed over $[\ln A, \ln B] \times \text{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$ with density probability $f_Y(N, \mathbf{z}, \theta)$ proportional to $|\det(D_F(N, \mathbf{z}, \theta))|$, i.e., proportional to $\exp(N)$. From Equation (IV.1) above, we know that $F(Y)$ is distributed as $\mathcal{U}(\phi(\mathcal{B}_{A,B}^{\eta,+}))$. We are then reduced to sampling such a random variable Y .

Note that the domain of Y is a ‘‘nice box’’: $[\ln A, \ln B] \times \text{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$. In this domain, we can sample a uniformly random variable Z in time $\text{poly}(d)$ (to sample from the ball $\text{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta)$, one can sample a Gaussian element and then renormalize it inside the ball). To obtain a sample from Y , we then keep Z with probability $\exp(Z_N)/B$, where Z_N is the first coordinate of Z . Note that the rejection probability is indeed between 0 and 1 since $Z_N \leq \ln(B)$.

It only remains to estimate the cost of the rejection step. Since $Z_N \geq \ln(A)$, the probability of keeping Z is at least A/B , and so the expected number of rejections before acceptance is bounded from above by B/A . \square

IV.4.2 Properties of the set $\mathcal{B}_{A,B}^{\eta}$

The goal of this subsection is to prove that the set $\mathcal{B}_{A,B}^{\eta}$ satisfies the properties needed to apply Lemma IV.3.2.

Lemma IV.4.5. *For any $B > A > 0$ and $\eta > 1$, the set $\mathcal{B}_{A,B}^{\eta}$ is compact, invariant by complex rotations and its elements are η -balanced.*

Proof. Compactness follows from the fact that $\mathcal{B}_{A,B}^{\eta}$ is closed and contained in the ball in infinity norm with radius $\eta \cdot B^{1/d}$. Invariance by complex rotations follows from the fact that both $\mathcal{N}(\cdot)$ and $\text{Ln}(\cdot)$ are invariant by complex rotations (i.e., we have $\mathcal{N}(\zeta x) = \mathcal{N}(x)$ and $\text{Ln}(\zeta x) = \text{Ln}(x)$ if $x \in K_{\mathbb{R}}$ and $\zeta \in K_{\mathbb{R}}$ is such that $|\sigma_i(\zeta)| = 1$ for all i 's). Let $x \in \mathcal{B}_{A,B}^{\eta}$, we have that

$$\left\| \frac{x}{\mathcal{N}(x)^{1/d}} \right\|_{\infty} = \exp\left(\left\| \text{Ln}\left(\frac{x}{\mathcal{N}(x)^{1/d}}\right) \right\|_{\infty}\right) \leq \exp\left(\left\| \text{Ln}\left(\frac{x}{\mathcal{N}(x)^{1/d}}\right) \right\|_2\right) \leq \eta.$$

The same holds for $\mathcal{N}(x)^{1/d}/x$ since $\left\| \text{Ln}\left(x/\mathcal{N}(x)^{1/d}\right) \right\|_{\infty} = \left\| \text{Ln}\left(\mathcal{N}(x)^{1/d}/x\right) \right\|_{\infty}$, which proves that x is η -balanced. \square

We now prove that the slices $\text{Ln}(\mathcal{B}_{A,B}^{\eta}) \cap H_t$ are empty when $t \notin [\ln A, \ln B]$ and have constant volume otherwise.

Lemma IV.4.6. *Let $B > A > 0$ and $\eta > 1$. For $t \in \mathbb{R}$, we define $H_t = \{x \in \text{Ln } K_{\mathbb{R}} \mid \sum_i x_i = t\}$. Then $\text{Ln}(\mathcal{B}_{A,B}^{\eta}) \cap H_t = \emptyset$ for $t \notin [\ln A, \ln B]$, and the volume of $\text{Ln}(\mathcal{B}_{A,B}^{\eta}) \cap H_t$ is constant for $t \in [\ln A, \ln B]$.*

Proof. By definition of $\mathcal{B}_{A,B}^\eta$, we have that

$$\text{Ln}(\mathcal{B}_{A,B}^\eta) = \left\{ \mathbf{x} \in \text{Ln}(K_{\mathbb{R}}) : \sum_{i \leq d} x_i \in [\ln A, \ln B], \left\| x - \left(\sum_{i \leq d} x_i \right) \cdot \mathbf{1}_d \right\|_2 \leq \ln \eta \right\},$$

where $\mathbf{1}_d$ refers to the d -dimensional all-1 vector. The intersection with H_t is the empty set if $t \notin [\ln A, \ln B]$. Otherwise, it is the ball centered in $t \cdot \mathbf{1}$ with radius $\ln(\eta)$, whose volume do not depend on t . \square

At this stage, only the first condition of Lemma IV.3.2 remains to be proved. We start by an auxiliary lemma, where we prove that if we shift the set $\mathcal{B}_{A,B}^\eta$ by some small vector, then the resulting set is included in another slightly larger set $\mathcal{B}_{A',B'}^{\eta'}$. The parameter f in the lemma below quantifies how small the shift vector needs to be, as a function of the parameters A and η . For the rest of the article, one can think of f as being of the order of $\text{poly}(d)$.

Lemma IV.4.7. *Let $B > A > 0$, $\eta > 1$ and $v \in K_{\mathbb{R}}$. Assume that $A^{1/d} \geq \eta \cdot f \cdot \|v\|_\infty$ for some $f > 1$. Then*

$$\mathcal{B}_{A,B}^\eta + v \subset \mathcal{B}_{A',B'}^{\eta'}$$

with $A' = A \cdot (1 - 1/f)^d$, $B' = B \cdot (1 + 1/f)^d$ and $\eta' = \eta \cdot \exp(2\sqrt{d}/(f - 1))$.

Proof. Let $x \in \mathcal{B}_{A,B}^\eta$, we are going to show that $x + v \in \mathcal{B}_{A',B'}^{\eta'}$. The definition of $\mathcal{B}_{A,B}^\eta$ and the fact that $A^{1/d} \geq \eta \cdot f \cdot \|v\|_\infty$ imply that we have, for every i ,

$$\frac{|v_i|}{|x_i|} \leq \frac{\|v\|_\infty}{|x_i|} \leq \frac{\|v\|_\infty \cdot \eta}{\mathcal{N}(x)^{1/d}} \leq \frac{\|v\|_\infty \cdot \eta}{A^{1/d}} \leq \frac{1}{f}.$$

The triangle inequality then gives that $|x_i + v_i| > 0$ for all i , and hence that $x + v \in K_{\mathbb{R}}^\times$. Further, note that

$$\frac{\mathcal{N}(x + v)}{\mathcal{N}(x)} = \prod_i \left| \frac{x_i + v_i}{x_i} \right| = \prod_i \left| 1 + \frac{v_i}{x_i} \right|.$$

Since $|v_i/x_i| \leq 1/f$ holds for all i , this implies that $\mathcal{N}(x + v)/\mathcal{N}(x) \in [(1 - 1/f)^d, (1 + 1/f)^d]$, which in turn shows that $\mathcal{N}(x + v) \in [A', B']$.

Towards completing the proof, note that

$$\begin{aligned} \left\| \text{Ln} \left(\frac{x + v}{\mathcal{N}(x + v)^{1/d}} \right) - \text{Ln} \left(\frac{x}{\mathcal{N}(x)^{1/d}} \right) \right\| &= \left\| \text{Ln} \left(1 + \frac{v}{x} \right) - \frac{1}{d} \ln \left(\frac{\mathcal{N}(x + v)}{\mathcal{N}(x)} \right) \cdot \mathbf{1} \right\| \\ &\leq \left\| \text{Ln} \left(1 + \frac{v}{x} \right) \right\| + \frac{1}{\sqrt{d}} \cdot \left| \ln \left(\frac{\mathcal{N}(x + v)}{\mathcal{N}(x)} \right) \right| \\ &\leq \sqrt{d} \cdot \frac{\|v/x\|_\infty}{1 - \|v/x\|_\infty} + \sqrt{d} \cdot \frac{1/f}{1 - 1/f} \\ &\leq \frac{2\sqrt{d}}{f - 1}, \end{aligned}$$

where we used the fact that

$$|\ln(1 + y)| = \max \left(\ln(1 + y), \ln \left(1 + \frac{-y}{1 + y} \right) \right) \leq \frac{|y|}{1 - |y|},$$

for any $y \in (-1, 1)$. This implies that

$$\left\| \text{Ln} \left(\frac{x+v}{\mathcal{N}(x+v)^{1/d}} \right) \right\| \leq \ln(\eta) + \frac{2\sqrt{d}}{f-1} = \ln(\eta').$$

We conclude that $x+v$ belongs to $\mathcal{B}_{A',B'}^{\eta'}$. \square

We are now ready to prove that the first condition of Lemma IV.3.2 is satisfied. To count the number of points of the ideal lattice I that belong to $\mathcal{B}_{A,B}^\eta$, we tile the space with shifts of a fundamental domain of the lattice (concretely, the Voronoi cell for the ℓ_∞ norm). Using Lemma IV.4.7, we show that the union of Voronoi cells corresponding to elements of $I \cap \mathcal{B}_{A,B}^\eta$ contains a smaller version $\mathcal{B}_{A_0,B_0}^{\eta_0}$ of the set, and is contained in a larger version $\mathcal{B}_{A_1,B_1}^{\eta_1}$. By carefully choosing parameters, we can ensure that the ratio of volumes of these two sets is bounded from above by a constant. In the lemma statement, note that C' is independent of the ideal I , but may depend on the other parameters, such as A, B, η and K . This proof is an adaptation of [Boe22, Lemma 6.13] with $\mathcal{B}_{A,B}^\eta$ instead of the ℓ_∞ ball.

Lemma IV.4.8. *Let A, B, η satisfying $A^{1/d} \geq \eta \cdot d^3 \cdot \Delta_K^{3/(2d)}$, $B/A \geq 4$ and $\eta \geq e$. There exists $C' > 0$ such that for any replete ideal $I \in \text{IdLat}_K^0$, we have*

$$\left| I \cap \mathcal{B}_{A,B}^\eta \right| \in C' \cdot [1, 340].$$

Proof. Let I be a norm-1 ideal, and let $V_\infty(I)$ be its ℓ_∞ -norm Voronoi cell, i.e., $V_\infty(I) = \{y \in K_\mathbb{R} : \forall x \in I \setminus \{0\}, \|y+x\|_\infty \geq \|y\|_\infty\}$. We let $\mu_\infty(I)$ denote the (ℓ_∞ -norm) radius of $V_\infty(I)$. By (II.2.4), we have that $\mu_\infty(I) \leq d \cdot \Delta_K^{3/(2d)}$. As a consequence, Lemma IV.4.7 instantiated with $f = d^2$ gives that since $A^{1/d} \geq \eta \cdot d^3 \cdot \Delta_K^{3/(2d)}$

$$\mathcal{B}_{A,B}^\eta + V_\infty(I) \subset \mathcal{B}_{A_1,B_1}^{\eta_1},$$

with $A_1 = A \cdot (1 - 1/d^2)^d$, $B_1 = B \cdot (1 + 1/d^2)^d$ and $\eta_1 = \eta \cdot \exp(2\sqrt{d}/(d^2 - 1))$. Recall that we assumed that $d \geq 2$ in all the article, so that we have $f > 1$ as needed for Lemma IV.4.7.

Let $A_0 = A \cdot (1 - 1/d^2)^{-d}$, $B_0 = B \cdot (1 + 1/d^2)^{-d}$ and $\eta_0 = \eta \cdot \exp(-2\sqrt{d}/(d^2 - 1))$. From the lower bound on η (and $d \geq 2$), one can check that $\eta_0 > 1$. Moreover, we have that $B_0/A_0 \geq 1/3 \cdot B/A \geq 4/3$ and hence that $B_0 > A_0$. Finally, from $A_0 \geq A$ and $\eta_0 \leq \eta$, we obtain that $A_0^{1/d} \geq \eta_0 \cdot f \cdot \mu_\infty(I)$ with $f = d^2$. This implies that we can apply Lemma IV.4.7 again on A_0, B_0, η_0 and $f = d^2$ and we obtain:

$$\mathcal{B}_{A_0,B_0}^{\eta_0} + V_\infty(I) \subset \mathcal{B}_{A,B}^\eta.$$

Note that for any $x \in \mathcal{B}_{A_0,B_0}^{\eta_0}$, there exists some (not necessarily unique) $\ell_x \in I$ such that $x - \ell_x \in V_\infty(I)$. This implies that $\ell_x \in (\mathcal{B}_{A_0,B_0}^{\eta_0} + V_\infty(I)) \cap I$. Therefore, we have

$$\mathcal{B}_{A_0,B_0}^{\eta_0} \subseteq \bigcup_{\ell \in (\mathcal{B}_{A_0,B_0}^{\eta_0} + V_\infty(I)) \cap I} \ell + V_\infty(I) \subseteq \bigcup_{\ell \in \mathcal{B}_{A,B}^\eta \cap I} \ell + V_\infty(I).$$

The above union is made of sets that are disjoint except for volume-0 intersections, so we have

$$\begin{aligned} \text{Vol}(\mathcal{B}_{A_0,B_0}^{\eta_0}) &\leq \text{Vol} \left(\bigcup_{\ell \in \mathcal{B}_{A,B}^\eta \cap I} \ell + V_\infty(I) \right) = \left| \mathcal{B}_{A,B}^\eta \cap I \right| \cdot \text{Vol}(V_\infty(I)) \\ &= \left| \mathcal{B}_{A,B}^\eta \cap I \right| \cdot \sqrt{\Delta_K}. \end{aligned}$$

Similarly, we have:

$$\left| \mathcal{B}_{A,B}^\eta \cap I \right| \cdot \sqrt{\Delta_K} \leq \text{Vol}(\mathcal{B}_{A_1,B_1}^{\eta_1}).$$

This gives us

$$\left| \mathcal{B}_{A,B}^\eta \cap I \right| \in C' \cdot \left[1, \frac{\text{Vol}(\mathcal{B}_{A_1,B_1}^{\eta_1})}{\text{Vol}(\mathcal{B}_{A_0,B_0}^{\eta_0})} \right],$$

where $C' = \text{Vol}(\mathcal{B}_{A_0,B_0}^{\eta_0})/\sqrt{\Delta_K} > 0$. It remains to bound the right boundary of the interval. By using Lemma IV.4.3, we obtain that

$$\frac{\text{Vol}(\mathcal{B}_{A_1,B_1}^{\eta_1})}{\text{Vol}(\mathcal{B}_{A_0,B_0}^{\eta_0})} = \frac{(B_1 - A_1) \cdot (\ln \eta_1)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1}}{(B_0 - A_0) \cdot (\ln \eta_0)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1}} \leq \frac{B_1 - A_1}{B_0 - A_0} \cdot \left(\frac{\ln \eta_1}{\ln \eta_0} \right)^{d-1}.$$

Recall that we have already seen that $B_0/A_0 \geq 4/3$. This implies that

$$\frac{B_1 - A_1}{B_0 - A_0} \leq \frac{B_1}{B_0 - A_0} = \left(1 + \frac{1}{d^2} \right)^{2d} \cdot \frac{1}{1 - (A_0/B_0)} \leq \frac{5}{2} \cdot \frac{1}{1 - 3/4} = 10.$$

Using the fact that $\ln \eta \geq 1$, we also have:

$$\left(\frac{\ln \eta_1}{\ln \eta_0} \right)^{d-1} = \left(\frac{\ln \eta + 2\sqrt{d}/(d^2 - 1)}{\ln \eta - 2\sqrt{d}/(d^2 - 1)} \right)^{d-1} \leq \left(\frac{1 + 2\sqrt{d}/(d^2 - 1)}{1 - 2\sqrt{d}/(d^2 - 1)} \right)^{d-1} \leq 34.$$

This completes the proof. \square

IV.4.3 Sampling uniform ideal elements in $\mathcal{B}_{A,B}^\eta$

We now show how to uniformly sample in $I \cap \mathcal{B}_{A,B}^\eta$, where I is a norm-1 ideal. For this purpose, `SampleUniform` _{A,B} ^{η} (Algorithm IV.4.1) uniformly samples in a larger $\mathcal{B}_{A_1,B_1}^{\eta_1}$ (using Lemma IV.4.4) and deterministically round to I using Babai's nearest plane algorithm [Bab86, Theorem 3.1]. The sample is kept if it belongs to $\mathcal{B}_{A,B}^\eta$.

Algorithm IV.4.1 `SampleUniform` _{A,B} ^{η}

Input: \mathbf{B}_I a basis of an ideal $I \in \text{IdLat}_K^0$.

Output: $x \in I \cap \mathcal{B}_{A,B}^\eta$.

- 1: Let $A_1 = A \cdot (1 - 1/d^2)^d$, $B_1 = B \cdot (1 + 1/d^2)^d$ and $\eta_1 = \eta \cdot \exp(2\sqrt{d}/(d^2 - 1))$.
 - 2: **repeat**
 - 3: Sample $y \leftarrow \mathcal{U}(\mathcal{B}_{A_1,B_1}^{\eta_1})$.
 - 4: Run Babai's nearest plane algorithm on (\mathbf{B}_I, y) ; let $x \in I$ be the output.
 - 5: **until** $x \in \mathcal{B}_{A,B}^\eta$.
 - 6: Return x .
-

Lemma IV.4.9. *Let A, B, η with $B/A \geq 4$ and $\eta \geq e$. Let $I \in \text{IdLat}_K^0$ given by a basis \mathbf{B}_I and $\delta = \|\mathbf{B}_I^*\|$. Assume that $A^{1/d} \geq d^{2.5} \cdot \eta \cdot \delta$. Then `SampleUniform` _{A,B} ^{η} samples uniformly in $I \cap \mathcal{B}_{A,B}^\eta$ and its expected running time is polynomial in $\log B$, d and B/A .*

Proof. Let $\mathcal{P}(\mathbf{B}_I) = \mathbf{B}_I^* \cdot (-1/2, 1/2]^d$ be the rounding cell of Babai's nearest plane algorithm. In order to prove that the output distribution is uniform, it suffices to prove that for any point $x \in$

$I \cap \mathcal{B}_{A,B}^\eta$, we have $\mathcal{P}(\mathbf{B}_I) + x \subset \mathcal{B}_{A_1,B_1}^{\eta_1}$. The definition of the nearest-plane algorithm's rounding cell implies that the ℓ_∞ norm of vectors in $\mathcal{P}(\mathbf{B}_I)$ is at most $\sqrt{d}\delta$. The definitions of A_1, B_1, η_1 and Lemma IV.4.7 (with $f = d^2$) allow us to conclude.

The running time follows from Lemma IV.4.4 and from bounding the probability that after Step 4, we have $x \notin \mathcal{B}_{A,B}^\eta$. This occurs if $y \notin \cup_{x \in \mathcal{B}_{A,B}^\eta \cap I} (x + \mathcal{P}(\mathbf{B}_I))$. As in the proof of Lemma IV.4.8, we have that:

$$\mathcal{B}_{A_0,B_0}^{\eta_0} \subset \sum_{x \in \mathcal{B}_{A,B}^\eta \cap I} x + \mathcal{P}(\mathbf{B}_I),$$

where $A_0 = A \cdot (1 + 1/d^2)^d$, $B_0 = B \cdot (1 - 1/d^2)^d$ and $\eta_0 = \eta \cdot \exp(-2\sqrt{d}/(d^2 - 1))$. The probability of exiting the loop is then bounded from below by

$$\frac{\text{Vol}(\mathcal{B}_{A_0,B_0}^{\eta_0})}{\text{Vol}(\mathcal{B}_{A_1,B_1}^{\eta_1})} \geq \frac{B_0 - A_0}{B_1 - A_1} \cdot \frac{(\ln \eta_0)^{d-1}}{(\ln \eta_1)^{d-1}} \geq \Omega(1),$$

where the inequalities are as in the proof of Lemma IV.4.8. \square

IV.5 Wrapping Up

We combine Theorems IV.3.4 and IV.4.2 to obtain the main result from this work. To simplify the statement, we instantiate the integral ideal-factoring oracle with a quantum polynomial-time algorithm, and use the Extended Riemann Hypothesis. The latter allows us to bound $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}|$ by a constant that is < 1 when $A \geq (\log \Delta_K)^{\Omega(1)}$ and $A/|\mathcal{P}_{A,4A}|$ by $O(\ln A)$ (see [BS96, Theorem 8.7.4]).

Theorem IV.5.1 (Assuming ERH). *There exists $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ such that the following holds. Let \mathcal{W} be a finite set of fractional ideals. Let $\gamma_{\mathcal{W}}, \gamma_{\mathcal{I}} \geq 1$ and A with $A^{1/d} \geq C_K \cdot \gamma_{\mathcal{W}}$. Let $\mathcal{O}_{\mathcal{W}}$ an oracle for \mathcal{W} -avg-id-HSVP $_{\gamma_{\mathcal{W}}}$ with success probability $\varepsilon_{\mathcal{W}}$ and $\mathcal{O}_{\mathcal{I}}$ an oracle for the problem $\mathcal{I}_{A,4A}$ -avg-id-HSVP $_{\gamma_{\mathcal{I}}}$ with success probability $\varepsilon_{\mathcal{I}}$.*

There exists a quantum algorithm making one call to $\mathcal{O}_{\mathcal{W}}$ and one call to $\mathcal{O}_{\mathcal{I}}$ whose running time is polynomial in $\log A$, $\log \Delta_K$ and the size of its input, such that the following holds. Given as input $I \sim \mathcal{U}(\mathcal{W})$, it outputs $x \in I^{-1} \setminus \{0\}$ with probability $\geq \varepsilon_{\mathcal{I}} \cdot (\varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)})$ such that

$$\|x\| \leq \gamma' \cdot \text{Vol}(I^{-1})^{1/d} \quad \text{with} \quad \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma_{\mathcal{I}}.$$

Proof. The algorithm is `InverseToIntegral` $^{\mathcal{W}}$ (Algorithm IV.3.3) instantiated with the set $\mathcal{B}_{A,B}^\eta$ with $B = 4A$ and $\eta = e$.

Note that at Step 3 of `InverseToIntegral` $^{\mathcal{W}}$, we have $\|\mathbf{B}_I\| \leq \delta_K \cdot \|v_I\|$ (by Lemma II.2.15). By definition of $\mathcal{O}_{\mathcal{W}}$, this implies that $\|\mathbf{B}_I\| \leq \delta_K \cdot \gamma_{\mathcal{W}} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$. `InverseToIntegral` $^{\mathcal{W}}$ then calls `RandomizeIdeal` (Algorithm IV.3.2), which at its Step 6 computes a basis \mathbf{B}_J of an ideal J that was showed in the proof of Lemma IV.3.3 to satisfy:

$$\begin{aligned} \|\mathbf{B}_J\| &\leq 85 \cdot d^2 \cdot \delta_K \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I)^{-1/d} \cdot \|\mathbf{B}_I\| \\ &\leq 85 \cdot d^2 \cdot \delta_K^2 \cdot \Delta_K^{3/(2d)} \cdot \gamma_{\mathcal{W}} =: \delta. \end{aligned}$$

The result follows from Theorems IV.3.4 and IV.4.2, using δ as above. \square

As a corollary, we obtain two quantum self-reductions, one from the problem $\mathcal{I}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma'}$ to the problem $\mathcal{I}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$ and the other from $\mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma'}$ to $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$ if $A^{1/d} \geq (d\delta_K \Delta_K^{1/d})^{\Omega(1)} \cdot \gamma$ and $\gamma' = O(d\Delta_K^{1/d}) \cdot \gamma$. Note that in the case of prime ideals, the success probability decreases with $\tilde{\rho}_A$ (the inverse of the proportion of prime ideals among all ideals of norm $\leq A$), which may or may not be small depending on the choice of the field K . This dependency arises from hoping that a uniform integral ideal is prime.

Corollary IV.5.2 (Assuming ERH). *There exists $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ such that the following holds. Let $\gamma \geq 1$ and A with $A^{1/d} \geq C_K \cdot \gamma$. Let \mathcal{O} an oracle for $\mathcal{I}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$ with success probability $\varepsilon \geq 2^{-\Omega(d)}$.*

There exists a quantum algorithm making two calls to \mathcal{O} whose running time is polynomial in $\log A$, $\log \Delta_K$ and the size of its input, such that, given as input $\mathfrak{a} \sim \mathcal{U}(\mathcal{I}_{A,4A})$, it outputs $x \in \mathfrak{a}^{-1} \setminus \{0\}$ with probability $\Omega(\varepsilon^2)$ with

$$\|x\| \leq \gamma' \cdot \text{Vol}(\mathfrak{a}^{-1})^{1/d} \quad \text{with } \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma.$$

Corollary IV.5.3 (Assuming ERH). *There exists $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ such that the following holds. Let $\gamma \geq 1$ and A with $A^{1/d} \geq C_K \cdot \gamma$. Let \mathcal{O} an oracle for $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$ with success probability $\varepsilon \geq 2^{-\Omega(d)}$.*

There exists a quantum algorithm making two calls to \mathcal{O} whose running time is polynomial in $\log A$, $\log \Delta_K$ and the size of its input, such that, given as input $\mathfrak{p} \sim \mathcal{U}(\mathcal{P}_{A,4A})$, it outputs $x \in \mathfrak{p}^{-1} \setminus \{0\}$ with probability $\Omega(\varepsilon^2/\tilde{\rho}_A)$ with

$$\|x\| \leq \gamma' \cdot \text{Vol}(\mathfrak{p}^{-1})^{1/d} \quad \text{with } \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma.$$

Combining Corollary IV.5.3 with Theorem IV.2.4, we obtain a quantum worst-case to average-case reduction for ideal-HSVP, where the average-case distribution is the uniform distribution over prime ideals with norm in some interval $[A, 4A]$.

Corollary IV.5.4 (Assuming ERH). *Let $\gamma \geq 1$. There exists some $\gamma' = \gamma \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and $A = \gamma^d \cdot \text{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ such that*

$$\text{id-HSVP}_{\gamma'} \quad \text{reduces to } \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}.$$

The reduction is quantum and runs in expected time polynomial in its input size, $\log \Delta_K$, $\tilde{\rho}_A$ and $1/\varepsilon$, where ε is the success probability of the oracle solving $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$.

Proof. We assume without loss of generality that $\gamma \leq 2^d$, since otherwise we can solve id-HSVP $_{\gamma'}$ in polynomial time using the LLL algorithm, for $\gamma' = \gamma \cdot \sqrt{d}$. We also assume that the success probability ε of the oracle solving $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$ is $\geq 2^{-\Omega(d)}$, since otherwise one can run an exact SVP solver in time $1/\varepsilon$.

Let $C'_{1,K}$ be the max of the $C_{1,K}$ from Theorem IV.2.4 and the C_K from Corollary IV.5.3. Then $C'_{1,K} = \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ since both quantities are. Let $A = (C'_{1,K} \cdot (232d \cdot \Delta_K^{1/d}) \cdot \gamma)^d$. One can check that $A = \gamma^d \cdot \text{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ as desired. Let also $\gamma' = A^{1/d} \cdot C_{2,K} = \gamma \cdot (232d \cdot \Delta_K^{1/d}) \cdot C'_{1,K} \cdot C_{2,K}$, where $C_{2,K}$ is as in Theorem IV.2.4. Similarly, one can check that $\gamma' = \gamma \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ as desired. Finally, let $\gamma_{\text{avg}} = 232d \cdot \Delta_K^{1/d} \cdot \gamma$.

Note that A , γ and ε satisfy the conditions from Corollary IV.5.3. So there is a quantum reduction

$$\text{from } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}} \quad \text{to } \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma},$$

which succeeds with probability $\delta = \Omega(\varepsilon^2/\tilde{\rho}_A)$ and runs in time polynomial in $\log \Delta_K$. Now, observe that γ_{avg} , A and γ' satisfy the conditions from Theorem IV.2.4, so there is a quantum reduction

$$\text{from id-HSVP}_{\gamma'} \text{ to } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}},$$

which runs in expected time polynomial in its input size, $\log \Delta_K$ and $1/\delta$. Combining both reductions and instantiating with the lower bound on δ completes the proof. \square

We now give a version of Corollary IV.5.4 with larger A , whose running time does not depend on the ad hoc parameter $\tilde{\rho}_A$.

Corollary IV.5.5 (Assuming ERH). *Let $\gamma \geq 1$. There exists $\gamma' = \gamma \cdot \text{poly}(\Delta_K^{\log(d)/d}, (\log \Delta_K)^{\log(d)}, \delta_K)$ and $A = \gamma^d \cdot \text{poly}(\Delta_K^{\log(d)}, (\log \Delta_K)^{d \ln(d)}, \delta_K^d)$ such that*

$$\text{id-HSVP}_{\gamma'} \text{ reduces to } \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}.$$

The reduction is quantum and runs in expected time polynomial in its input size, $\log \Delta_K$, ρ_K and $1/\varepsilon$, where ε is the success probability of the oracle solving $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma}$.

Proof. The proof is the same as the one of Corollary IV.5.4. We highlight the changes in blue. We define

$$A = \max \left((C'_{1,K} \cdot 232d \cdot \Delta_K^{1/d} \cdot \gamma)^d, (2 \cdot d^d \cdot \Delta_K)^{c_2 \ln(d)} \right),$$

where $C'_{1,K}$ is the max of the $C_{1,K}$ from Theorem IV.2.4 and C_K from Corollary IV.5.3, and c_2 is defined in Corollary III.1.3. It holds that $A = \gamma^d \cdot \text{poly}(\Delta_K^{\log(d)}, (\log \Delta_K)^{d \ln(d)}, \delta_K^d)$ as desired. Let also $\gamma' = A^{1/d} \cdot C_{2,K}$, where $C_{2,K}$ is as in Theorem IV.2.4. Similarly, one can check that $\gamma' = \gamma \cdot \text{poly}(\Delta_K^{\log(d)/d}, (\log \Delta_K)^{\log(d)}, \delta_K) = \gamma \cdot \text{poly}(\Delta_K^{\log(d)/d}, (\log \Delta_K)^{\log(d)}, \delta_K)$. Finally, let $\gamma_{\text{avg}} = 232d \cdot \Delta_K^{1/d} \cdot \gamma$ as desired. Note that A , γ and ε satisfy the conditions from Corollary IV.5.3. So there is a quantum reduction

$$\text{from } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}} \text{ to } \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\gamma},$$

which succeeds with probability $\delta = \Omega(\varepsilon^2/\tilde{\rho}_A)$ and runs in time polynomial in $\log \Delta_K$. We have by Corollary III.1.3 (with $C = 2$) and the lower bound on A that $\tilde{\rho}_A \geq 0.5 \cdot \rho_K \cdot \ln(A)$. Now, observe that γ_{avg} , A and γ' satisfy the conditions from Theorem IV.2.4, so there is a quantum reduction

$$\text{from id-HSVP}_{\gamma'} \text{ to } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}},$$

which runs in expected time polynomial in its input size, $\log \Delta_K$ and $1/\delta$. Combining both reductions and instantiating with the lower bound on δ completes the proof. \square

IV.6 NTRU with Polynomial Modulus

The main result of this section is Corollary IV.6.2. It gives a distribution over NTRU instances with small modulus q that is hard on average, under the worst-case id-HSVP hardness assumption.

Note that in this work we are only interested in the vector version of NTRU. We denote [PS21, Alg. 4.1] by `IdealToNTRU`. It takes as input a basis of an integral ideal \mathfrak{a} and a modulus q and outputs an instance of (γ, q) -NTRU whose solution is related to a short non-zero vector of \mathfrak{a} . The following result is a consequence of [PS21, Lemma 4.3], whose proof is very similar to [PS21, Theorem 4.1]. We provide a proof for the sake of completeness.

Theorem IV.6.1 (Adapted from [PS21, Theorem 4.1]). *Let $\gamma \geq \gamma' \geq 1$ be real numbers, $q \geq 2$ be an integer, and*

$$N = \frac{1}{2^{d+2}} \cdot \left(\frac{\sqrt{q}}{\sqrt{2} \cdot \gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}} \right)^d.$$

Let \mathfrak{a} be an integral ideal of norm in $[N, 2^{d+2} \cdot N]$ and $h = \text{IdealToNTRU}(\mathfrak{a}, q)$. Then h is a (γ, q) -NTRU instance. If (f, g) is a solution to (γ, γ', q) -NTRU on instance h , then g is a solution to $\gamma_{\text{HSVP-id-HSVP}}$ for instance \mathfrak{a} , where $\gamma_{\text{HSVP}} = \gamma/\gamma' \cdot 4\sqrt{2} \cdot d^{1.5} \cdot \delta_K$.

Further, IdealToNTRU runs in time polynomial in its input size and in $\log \Delta_K$.

Note that the statement is void if $2^{d+2} \cdot N < 1$ (no integral ideal has norm in $(0, 1)$): an extra parameter constraint is implicitly required for it to be meaningful.

Proof. Recall that for our definition of NTRU, a (γ, q) -NTRU instance for [PS21] is a $(\gamma/\sqrt{2}, q)$ -NTRU instance for us. We adapted the values on the theorem to take this fact into account. The running time of IdealToNTRU is stated in [PS21, Lemma 4.3].

By [PS21, Lemma 4.3], there exists $(f, g) \in \mathcal{O}_K^2 \setminus \{(0, 0)\}$ such that $g \cdot h = f \pmod{q}$ and $\|f\|, \|g\| \leq d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d}$. (Note that δ_K in the present work is an upper bound on the quantity δ_K from [PS21].) Using $\mathcal{N}(\mathfrak{a}) \leq 2^{d+2} \cdot N$ and the definition of N , this gives that h is a (γ, q) -NTRU instance.

Assume now that $(f, g) \in \mathcal{O}_K \setminus \{(0, 0)\}$ is a solution to (γ', γ, q) -NTRU for instance h . Then we have

$$\begin{aligned} \|f\|, \|g\| &\leq \frac{\sqrt{q}}{\gamma'} \leq \frac{q}{\sqrt{2} \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot (2^{d+2} \cdot N)^{1/d}} \\ &\leq \frac{q}{\sqrt{2} \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d}}, \end{aligned}$$

where the second inequality comes from the definition of N , and the third one comes from the assumption $\mathcal{N}(\mathfrak{a}) \leq 2^{d+2} \cdot N$. By [PS21, Lemma 4.3], we obtain that $g \in \mathfrak{a} \setminus \{0\}$. Finally, the fact that g is a solution to γ_{HSVP} follows from

$$\begin{aligned} \|g\| &\leq \frac{\sqrt{q}}{\gamma'} = \frac{2^{(d+2)/d} \cdot N^{1/d} \cdot \sqrt{2} \cdot \gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}}{\gamma'} \\ &\leq \frac{4\sqrt{2} \cdot \gamma \cdot d^{1.5} \cdot \delta_K}{\gamma'} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d}, \end{aligned}$$

where the last inequality follows from the inequalities $\mathcal{N}(\mathfrak{a}) \geq N$ and $d \geq 2$. \square

For $A, q \geq 2$, we define $D_{\text{NTRU}}^{A,q} = \text{IdealToNTRU}(\mathcal{U}(\mathcal{P}_{A,4A}), q)$. Theorem IV.6.1 implies a polynomial-time reduction from $\mathcal{I}_{A,4A}\text{-avg-id-HSVP}$ to $(D_{\text{NTRU}}^{A,q}, \gamma, \gamma', q)$ -NTRU for well chosen γ, γ', A and q . Combining Corollary IV.5.4 and Theorem IV.6.1 give the following result.

Corollary IV.6.2 (Assuming ERH). *Let $\gamma \geq \gamma' \geq 1$. There exists an integer $q = (\gamma^A/\gamma'^2) \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$, and real numbers $\gamma_{\text{HSVP}} = (\gamma/\gamma') \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and $A = (\gamma/\gamma')^d \cdot \text{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ such that*

$$\text{id-HSVP}_{\gamma_{\text{HSVP}}} \text{ reduces to } (D_{\text{NTRU}}^{A,q}, \gamma, \gamma', q)\text{-NTRU}.$$

The reduction is quantum and runs in expected time polynomial in its input size, $\log q, \log \Delta_K, 1/\tilde{\rho}_A$ and $1/\varepsilon$, where ε is the success probability of the oracle solving $(D_{\text{NTRU}}^{A,q}, \gamma, \gamma', q)$ -NTRU.

Proof. Without loss of generality, we can assume that $\gamma/\gamma' \leq 2^d$, since otherwise we have a polynomial time algorithm solving $\text{id-HSVP}_{\gamma_{\text{HSVP}}}$ for $\gamma_{\text{HSVP}} = \gamma/\gamma'$. Let $\Gamma = (\gamma/\gamma') \cdot 4d^{1.5} \cdot \delta_K$. Let $A = \Gamma^d \cdot \text{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ be as in Corollary IV.5.4, with “ $\gamma = \Gamma$ ”. Similarly, let $\gamma_{\text{HSVP}} = \Gamma \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ be the quantity γ' from Corollary IV.5.4, with “ $\gamma = \Gamma$ ”. Finally, let $X = \gamma \cdot 2 \cdot (4A)^{1/d} \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}$ and $q = \lfloor X^2 \rfloor$. Note that $q \geq X^2/4$. Note that $\gamma_{\text{HSVP}} = (\gamma/\gamma') \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and that $q = (\gamma^4/\gamma'^2) \cdot \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$.

Let $N = \frac{1}{2^{d+2}} \cdot \left(\frac{\sqrt{q}}{\gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}} \right)^d$ be as in Theorem IV.6.1. Using the fact that $X/2 \leq \sqrt{q} \leq X$ and the definition of X , we have that $[A, 4A] \subseteq [N, 2^{d+2} \cdot N]$. Hence, the support of the distribution $\mathcal{U}(\mathcal{P}_{A,4A})$ is contained in the set of integral ideals with norm in $[N, 2^{d+2} \cdot N]$.

Recall that $D_{\text{NTRU}}^{A,q}$ is the distribution $\text{IdealToNTRU}(\mathcal{U}(\mathcal{P}_{A,4A}), q)$. By Theorem IV.6.1, there is a reduction from $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\Gamma}$ to $(D_{\text{NTRU}}^{A,q}, \gamma, \gamma', q)\text{-NTRU}$, which runs in time polynomial in $\log q$, $\log \Delta_K$ and $\log A = \text{poly}(\log \Delta_K)$ (since $\gamma/\gamma' \leq 2^d$) and preserves the success probability of the algorithm. Moreover, from Corollary IV.5.4, $\text{id-HSVP}_{\gamma_{\text{HSVP}}}$ reduces to $\mathcal{P}_{A,4A}\text{-avg-id-HSVP}_{\Gamma}$, which is quantum and runs in expected time polynomial in its input size, $\log \Delta_K$, $1/\tilde{\rho}_A$ and $1/\varepsilon$. Combining both reductions gives the desired result. \square

Note that the distribution $D_{\text{NTRU}}^{A,q}$ can be sampled from along with a trapdoor by running `SampleWithTrap` with appropriate parameters (in order to generate an ideal from $\mathcal{U}(\mathcal{P}_{A,4A})$ together with a short non-zero vector in it), and then running the `IdealToNTRU` algorithm. This, however, requires an access to a factoring oracle (for the `SampleWithTrap` algorithm). Finding a classical algorithm to efficiently sample from $D_{\text{NTRU}}^{A,q}$ with a trapdoor is an interesting open problem.

Chapter V

On Module Unique-SVP and NTRU

This chapter is extracted from [FPS22], which is a joint work with Alice Pellet-Mary and Damien Stehlé. My main contribution to this work is the worst-case to average-case reduction for mod-uSVP₂. I also adapted the mod-uSVP₂-to-NTRU reduction for using on ideal-counting of Chapter III.

V.1 Introduction

Contributions of this chapter

We give evidence that the NTRU problem is not just a particular case of mod-uSVP₂, but actually representative of it. More precisely, we show that worst-case NTRU is computationally equivalent to worst-case mod-uSVP₂, and that worst-case and an appropriately defined average-case mod-uSVP₂ are also computationally equivalent, provided we have an oracle for id-HSVP in both cases (and up to reduction losses). Together, these results imply that worst-case mod-uSVP₂ reduces to average-case NTRU, provided we have an oracle for id-HSVP. Combining this result with the reduction from worst-case id-HSVP to worst-case NTRU from [PS21], this also implies that worst-case NTRU is computationally equivalent to worst-case mod-uSVP₂, without an id-HSVP oracle.

Our first result is a collection of four reductions from the four variants of mod-uSVP₂ (average case vs worst-case and vector vs module) to the corresponding four variants of NTRU, relying on an approximate id-HSVP oracle. We give below a simplified version of one of these reductions, in the special case of power-of-two cyclotomic fields. More details and the other reductions can be found in Theorem V.4.1.

Theorem V.1.1 (Simplified version of Theorem V.4.1). *Let K be a power-of-two cyclotomic field of degree d . Let $\gamma_{\text{SVP}}, \gamma^+, \gamma_{\text{NTRU}} > 1$. For all $q \geq 2^d \cdot \text{poly}(\gamma^+)$ and $\gamma^- \geq \text{poly}(d) \cdot \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}}$, (worst-case) mod-uSVP₂^{mod} with gap in $[\gamma^-, \gamma^+]$ reduces in polynomial time to (worst-case) NTRU^{mod} with modulus q and gap $\geq \gamma_{\text{NTRU}}$ and (worst-case) id-SVP with approximation factor γ_{SVP} .*

More concretely, when starting from a mod-uSVP₂ instance for which the shortest non-zero vectors are $\approx \gamma$ times smaller than the root determinant, the reduction produces an NTRU instance satisfying $\sqrt{q}/(\|f\| + \|g\|) \approx \gamma^{O(1)}$, up to factors depending on field invariants. This transformation can be used to derive a reduction from average-case mod-uSVP₂ to average-case NTRU (where the NTRU distribution is induced by the mod-uSVP₂ distribution) and a reduction from worst-case mod-uSVP₂ to worst-case NTRU (and similarly for the variants searching a

dense rank-1 submodule). To achieve this transformation, an id-HSVP oracle is required to find non-zero vectors in ideals within a factor $\gamma^{O(1)}$ from optimal. Note that for cyclotomic fields, the algorithm from [CDW21] allows to implement the oracle in quantum polynomial time when $\gamma \approx 2^{\sqrt{d}}$. Note also that [PS21] showed a reduction from worst-case id-HSVP to worst-case NTRU, which is compatible with the reduction from worst-case mod-uSVP₂ to worst-case NTRU (relying on an id-HSVP oracle). Combining both, we then obtain a reduction from worst-case mod-uSVP₂ to worst-case search NTRU which *does not* rely on an id-HSVP oracle. A drawback of the reduction is that it results in an NTRU modulus q of the order of $\approx 2^d$, even for small gap parameters γ . The modulus can be decreased by allowing the reduction to be more costly. Using lattice reduction algorithms [Sch87], one can reach $q \approx \gamma^{O(1)} \cdot \beta^{O(d/\beta)}$ if allowing for a reduction that runs in time polynomial in d , 2^β , $\log \Delta_K$ and $\zeta_K(2)$ (where ζ_K refers to the Dedekind zeta function). The quantities $\log \Delta_K$ and $\zeta_K(2)$ depend on the number field, and may not be polynomially bounded in the field degree d . In our running example, we have $\log \Delta_K = O(d)$ and $\zeta_K(2) = O(1)$ (see [SS13]).

Second, we exhibit a random self-reducibility property for mod-uSVP₂^{mod}. More explicitly, we give a reduction from worst-case mod-uSVP₂^{mod} for rank-2 modules to an average-case version of itself, which can be efficiently sampled. The reduction preserves the gap parameter γ , up to factors depending on field invariants, and runs in time polynomial in $\log \Delta_K$.

Theorem V.1.2 (Simplified version of Theorem V.6.1, under ERH). *Let K be a power-of-two cyclotomic field of degree d . For any gap $\text{poly}(d) < \gamma \leq 2^{O(d)}$, there exists a distribution $D_\gamma^{\text{mod-uSVP}_2}$ over mod-uSVP₂ instances with gap $\geq \gamma$ which can be sampled efficiently such that worst-case mod-uSVP₂^{mod} with gap $\geq \gamma' = \gamma \cdot \text{poly}(d)$ reduces in polynomial time to average-case mod-uSVP₂^{mod} for instance distribution $D_\gamma^{\text{mod-uSVP}_2}$.*

Combined with the first reduction, the above allows to map a worst-case instance of NTRU^{mod} to an average-case instance of NTRU^{mod}, where the NTRU^{mod} instance distribution is inherited from the average-case mod-uSVP₂ distribution. This reduction relies on an id-HSVP oracle. Since mod-uSVP₂^{mod} and mod-uSVP₂ are computationally equivalent (up to polynomial losses) when we have an id-HSVP oracle, this also provides a reduction from worst-case mod-uSVP₂ to average-case NTRU. Contrary to the reduction from worst-case uSVP to worst-case NTRU, we cannot use the result of [PS21] to get rid of the id-HSVP oracle. This is because the average-case distribution of NTRU instances that is produced by our reduction may not be compatible with the one used in [PS21].

We summarize the known reductions between variants of mod-uSVP₂ and NTRU in Figure V.1. Note that the reductions may not be composable due to incompatible parameter restrictions or instance distributions.

Technical overview

The NTRU problem is a restriction of mod-uSVP₂ modules with a basis of a specific shape. In general, a rank-2 module M is represented by a pseudo-basis, i.e., two vectors $(\mathbf{b}_1, \mathbf{b}_2)$ in K^2 and two ideals I_1, I_2 of \mathcal{O}_K such that $M = \mathbf{b}_1 I_1 + \mathbf{b}_2 I_2$. When the two ideals I_1 and I_2 are both equal to \mathcal{O}_K , the pseudo-basis is a basis, and the module is said to be free (note that a free module is a module that has at least one basis, but not all of its pseudo-bases will satisfy $I_1 = I_2 = \mathcal{O}_K$). In the NTRU problem, the instance is a basis $(\mathbf{b}_1, \mathbf{b}_2)$ of a free module contained in \mathcal{O}_K^2 , with $\mathbf{b}_1 = (1, h)^T$ for some $h \in \mathcal{O}_K$ and $\mathbf{b}_2 = (0, q)^T$ for some integer q which is a parameter of the NTRU problem. Hence, the only degree of freedom in this basis comes from the choice of h . The NTRU problem then asks to solve mod-uSVP₂ in this very specific module.




Figure V.1: Known reductions between NTRU and mod-uSVP variants. Dashed arrows require an id-HSVP oracle. Blue arrows are proven in [PS21] and red arrows are proven in this article. The black arrows are folklore.

In the reduction from mod-uSVP₂ to NTRU, we start with an arbitrary pseudo-basis of an arbitrary module M , and transform it into an NTRU basis. We then call the NTRU solver on this NTRU instance and lift the solution back to the original mod-uSVP₂ module. In order to meaningfully lift a short vector (or a dense rank-1 submodule) back, we require our transformation to preserve the geometry of the rank-2 module M as much as possible. Our transformation proceeds in four main steps.

First of all, we transform the input module $M \subset K^2$ into an integral module whose volume is bounded from below and above by quantities depending only on the parameters of the reduction (NTRU modules are in \mathcal{O}_K^2 and have volume q^d). This is done by scaling M to the desired volume, and then rounding it to an integral module with a very close geometry. This rounding is performed by sampling two quasi-orthogonal vectors in the dual of M , and multiplying M on the left by the matrix whose rows are these two vectors. Multiplication on the left corresponds to a distortion of the ambient space, but since the two vectors are quasi orthogonal, this does not change the geometry too much. Also, as the row vectors of the sampled matrix belong to the dual of M , the resulting module is integral.

Our second step aims at obtaining the triangular shape of the NTRU basis. To do so, we compute the Hermite Normal Form of the pseudo-basis. With some probability, the two coefficients on the first row of the pseudo-basis will be coprime, leading to an HNF basis with a 1 as a top-left coefficient, exactly what we need for an NTRU instance. This is where $\zeta_K(2)$ comes into play, as it closely relates to the probability that two random elements of \mathcal{O}_K are coprime.

At this point, our pseudo-basis still has coefficient ideals. We remove them with an id-HSVP solver: we compute short x_1 and x_2 in the ideals I_1 and I_2 , respectively, and then replace the pseudo-basis $((\mathbf{b}_1, \mathbf{b}_2), (I_1, I_2))$ by the basis $(x_1\mathbf{b}_1, x_2\mathbf{b}_2)$. This step has the effect of slightly sparsifying the module, i.e., it leads to a rank-2 submodule whose determinant is not much larger. If our gap is sufficiently large compared to the approximation factor of the id-HSVP solver, our sparsified module will still contain an unexpectedly short non-zero vector.

We now have a basis of a free module with vectors of the form $(1, h')^T$ and $(0, b)^T$, with h' and b in \mathcal{O}_K . Our last step consists in replacing b by the NTRU parameter q . This is done by multiplying the second coordinates of both our basis vectors by q/b . If $q/b \approx 1$ (which we can ensure thanks to the id-HSVP solver), then this does not change the geometry of the module too much.

To conclude, the transformation we have described allows us to transform any module of rank-2 with an unexpectedly short vector into an NTRU module with roughly the same geometry. The transformation is reversible, hence, we can lift any short vector or dense module found in the NTRU module back to the original rank-2 module. Since this transformation is a Karp reduction, it can be used to reduce average-case variants of mod-uSVP₂ to average-case variants of NTRU where the NTRU distribution is inherited from the one on the mod-uSVP₂ instances.

For the random self-reducibility of mod-uSVP₂^{mod}, we start with an arbitrary rank-2 module M and want to randomize it so that the distribution of the output module M' does not depend on M . Once again, we design the transformation so that it preserves the geometry of the module, to be able to meaningfully lift any dense rank-1 submodule of M' back to a dense rank-1 submodule of M . For this reduction, we assume that all our worst-case modules live in $K_{\mathbb{R}}^2 = (K \otimes_{\mathbb{Q}} \mathbb{R})^2$ and have fixed volume (which we can always achieve by scaling the module). We also assume that the ℓ_2 -norm of their shortest non-zero vectors is exactly $1/\gamma < 1$. This restriction to modules with a known gap can be waived, by guessing the gap and sparsifying the module (see Section V.6).

Let us explain the main ideas behind the randomization in the simpler case of $K = \mathbb{Q}$. We have a lattice $M \subset \mathbb{R}^2$ with volume 1 and shortest non-zero vector \mathbf{s} with $\|\mathbf{s}\| = 1/\gamma$. Up to rotation of the ambient space, we can assume that $\mathbf{s} = (1/\gamma, 0)^T$. Let us take $\mathbf{t} \in \mathbb{R}^2$ such that (\mathbf{s}, \mathbf{t}) forms a basis of M . Since the volume of M is 1, we know that $\mathbf{t} = (t_0, \gamma)^T$ for some $t_0 \in \mathbb{R}$. Up to the rotation of the ambient space, the quantity t_0 is the only degree of freedom. Note also that the lattice only depends on $t_0 \bmod 1/\gamma$. Let $\pi_{\mathbf{s}}(\mathbf{t})$ denote the quantity t_0 , i.e., the norm of the orthogonal projection of \mathbf{t} onto $\text{span}(\mathbf{s})$. This discussion shows that the lattice M is uniquely determined by the span of its shortest non-zero vector and the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) \bmod 1$. Hence, to “hide” the lattice M , it suffices to “hide” these two quantities. Note that we use the vectors \mathbf{s} and \mathbf{t} for our reasoning, but we usually do not have access to them: we randomize our module by performing only operations that can be done on any of the bases of M (for $K_{\mathbb{R}}^2$ instead of \mathbb{R}^2 , we expect that finding the analogue of (\mathbf{s}, \mathbf{t}) is difficult).

In order to hide the span of \mathbf{s} , one can apply a uniform orthonormal transformation to the ambient space. To hide the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) \bmod 1$, we “blur” the ambient space, by applying to it a transformation that is close to orthogonal, but not fully so. By appropriately choosing the transformation, one can obviously transform the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t})$ into $x \cdot \gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) + y$, where x and y are some random variables. Recall that this quantity only matters modulo 1. Hence, if the standard deviation of y is sufficiently large compared to 1, then $y \bmod 1$ will be uniformly distributed and will hide the original value of $\pi_{\mathbf{s}}(\mathbf{t})$. The existence of a gap ensures that a close-to-orthogonal transformation suffices for this purpose.

This intuition over \mathbb{R}^2 explains one component of our randomization procedure, which we call the geometric randomization (see Section V.5.2). Another important part of our randomization, which we call the coefficient randomization (Section V.5.1), focuses on the coefficient ideals of the pseudo-basis (which are just \mathbb{Z} for lattices). The transformation described above will have the effect of randomizing the vectors \mathbf{b}_1 and \mathbf{b}_2 of a pseudo-basis of our module M , but will have no impact on the coefficients ideals I_1 and I_2 .

In order to hide those ideals, the first step is to multiply the module M by some uniformly distributed ideal I , using [BDPW20]. Our new coefficient ideals $I \cdot I_1$ and $I \cdot I_2$ will then be uniformly distributed too. This is however not sufficient to fully hide the ideals, since the quotient $(I \cdot I_1)/(I \cdot I_2)$ is constant. In order to hide this last quantity, or decouple the ideals, we sparsify the module with respect to some prime ideal \mathfrak{p} : concretely, we take a uniformly random rank-2 submodule of M among those of index \mathfrak{p} .¹ This process generalizes lattice sparsification

¹For two rank-2 modules $M' \subseteq M$ with pseudo-bases $((\mathbf{b}'_1, I'_1), (\mathbf{b}'_2, I'_2))$ and $((\mathbf{b}_1, I_1), (\mathbf{b}_2, I_2))$ respectively, we

as introduced in [Kho06]. Lattice sparsification is a classic tool to remove one (or several) annoying vectors in a lattice. Here, the purpose is different: it has the effect of obliviously multiplying I_1 by \mathfrak{p} while leaving I_2 unchanged (with probability close to 1). By [BDPW20], the uniform distribution over bounded-norm prime ideals is close to the uniform distribution over norm-1 ideals (after renormalization of their norm), in the sense that little remains to be done to obtain the latter distribution. As a result, this sparsification enables us to (almost) randomize both I_1 and I_2 , independently of one another. The gap to perfect randomization is handled by carefully studying the distribution resulting from the geometric and coefficient randomization (Section V.5.3).

Summing up, our randomization consists in two main steps: a distortion of the ambient space, which randomizes the vectors $(\mathbf{b}_1, \mathbf{b}_2)$ and a sparsification, which hides the coefficient ideals I_1 and I_2 (together with the multiplication of the module by a random ideal I). Interestingly, we note that these two operations are similar (though adapted to rank-2 modules) to the ones that were used in [BDPW20] to randomize ideal lattices.

The transformation described above allows us to transform an arbitrary module M of $K_{\mathbb{R}}^2$ into a random module M' of $K_{\mathbb{R}}^2$ whose distribution is independent of the input module. One last subtlety to handle in order to have a full worst-case to average-case reduction is to compute a canonical representation of the module M' . Indeed, the pseudo-basis of the properly distributed module M' that we have at the end of the randomization procedure might leak information about the input module M . Unfortunately, one cannot compute HNF bases in $K_{\mathbb{R}}^2$ (the HNF gives a canonical representation of rational lattices). In order to obtain a canonical representation of M' , we then round it to a close module in \mathcal{O}_K^2 for which we will be able to compute an HNF pseudo-basis. The rounding procedure is the same as the one described in the reduction from mod-uSVP₂ to NTRU, and the distribution of the output pseudo-basis only depends on the input module and not on the specific pseudo-basis that is provided to represent it.

Discussion

A question arising from our reduction concerns the possibility to sample an NTRU instance from the distribution obtained at the end of the reduction, together with a short secret vector of the corresponding NTRU module. The difficulty stems from the fact that the output NTRU distribution we obtain after the reduction is not easy to describe, except as “the distribution obtained by running the reduction”. The same difficulty also appeared in [PS21], where it was tackled by running the reduction to sample from the average-case NTRU distribution (and keeping in mind some quantities generated during the reduction in order to create a short vector of the output NTRU module). In our case, we face two additional difficulties when trying to apply the same strategy. First, we note that even sampling from the NTRU distribution, without asking for a short vector of the corresponding module, does not seem straightforward. Since our mod-uSVP₂ to NTRU reduction requires an id-HSVP solver and takes subexponential time if one wants to reach small NTRU modulus q , it does not provide an efficient sampling algorithm for our final NTRU distribution. Secondly, our reduction allows us to lift a short vector from the NTRU module back to the mod-uSVP₂ module, but it is not so clear whether the converse is also possible (i.e., starting with a known vector of the mod-uSVP₂ module and obtaining a short vector of the final NTRU module). This is because of the sparsification step: when we sparsify a lattice, we can lift a vector from the sparser lattice back to the denser lattice (this is actually the same vector), but the converse seems more difficult.

Another question we leave open is about the compatibility of our reduction with those from [PS21]. Our worst-case mod-uSVP₂^{mod} to average-case NTRU^{mod} reduction produces a

say that M' has index \mathfrak{p} in M if $\det_K(\mathbf{b}'_1, \mathbf{b}'_2) \cdot I'_1 I'_2 = \mathfrak{p} \cdot \det_K(\mathbf{b}_1, \mathbf{b}_2) \cdot I_1 I_2$.

new distribution over NTRU instances. It is unclear whether this distribution is compatible with the search to decision reduction for NTRU from [PS21, Definition 7.1]. It is also unclear how it compares to the one produced by the worst-case id-HSVP to average-case NTRU reduction from [PS21].

It should be noted that the regime where NTRU is provably secure (see [SS13]) is completely distinct from the regime required by our reductions. Indeed, the regime of [SS13] requires that f and g are slightly larger than \sqrt{q} , whereas our reduction requires f and g to be significantly smaller than \sqrt{q} . In other words, we are in a regime where NTRU is a mod-uSVP₂ instance (and we are trying to show that in this regime, it is representative of all mod-uSVP₂ instances), whereas [SS13] works in a regime where an NTRU instance is statistically close to uniform; in particular, in that regime, the underlying lattice is not a mod-uSVP₂ instance. The regime of the overstretch-NTRU attacks (including [KF17]) is also distinct from ours, but in the opposite direction. In these attacks, it is assumed that $\|f\|$ and $\|g\|$ are poly(d) and q grows; whereas in our case, we have $\|f\|$ and $\|g\|$ of the form $\sqrt{q}/\text{poly}(d)$. Said differently, in those attacks, the short vector is short in absolute terms, whereas in our case it is short relative to what it would be for a random lattice of the same volume. We leave as an open problem to check whether these two regimes can be made to intersect.

V.2 Preliminaries

V.2.1 Number Fields

For $x \in K_{\mathbb{R}}^+$, we define $x^{1/2}$ as the element of $K_{\mathbb{R}}^+$ obtained by taking the square-roots of the embeddings.

Recall that elements of $K_{\mathbb{R}}$ are represented as vector of \mathbb{R}^d corresponding to their coordinates in the \mathbb{Z} -basis $\mathbf{B}_{\mathcal{O}_K} = [b_1^{\mathcal{O}_K}, \dots, b_d^{\mathcal{O}_K}]$ of \mathcal{O}_K . For $x = \sum_i x_i b_i \in K_{\mathbb{R}}$, we define $\lfloor x \rfloor = \sum_i \lfloor x_i \rfloor b_i$. We will use the notation $\{x\} = x - \lfloor x \rfloor$. We have $\|\{x\}\|_{\infty} \leq d \cdot \delta_K$, and hence $\|\{x\}\| \leq d^{3/2} \cdot \delta_K$.

We will consider the following distributions over $K_{\mathbb{R}}$.

Definition V.2.1. Let $\varsigma \in \mathbb{R}_{>0}^{d_{\mathbb{R}}+d_{\mathbb{C}}}$. We define the normal distribution $\mathcal{D}_{K_{\mathbb{R}}}(c, \varsigma)$ of center $c \in K_{\mathbb{R}}$ and standard deviation vector ς as the distribution obtained by independently sampling real numbers $(y)_{i \in [d]}$ with

$$\begin{cases} y_j \sim \mathcal{D}(0, \varsigma_j) & \text{for } j \in [d_{\mathbb{R}}] \\ y_{d_{\mathbb{R}}+j}, y_{d_{\mathbb{R}}+d_{\mathbb{C}}+j} \sim \mathcal{D}(0, \varsigma_{d_{\mathbb{R}}+j}) & \text{for } j \in [d_{\mathbb{C}}] \end{cases}$$

and then returning $c + y$ where $y \in K_{\mathbb{R}}$ is such that $\sigma_j(y) = y_j$ for $j \in [d_{\mathbb{R}}]$ and $\sigma_{d_{\mathbb{R}}+j}(y) = y_{d_{\mathbb{R}}+j} + iy_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}$ for $j \in [d_{\mathbb{C}}]$.

We define $\chi_{K_{\mathbb{R}}}$ as the distribution of $(\langle \mathbf{x}, \mathbf{x} \rangle_{K_{\mathbb{R}}})^{1/2}$ for $\mathbf{x} \in K_{\mathbb{R}}^2$ sampled according to $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^2$.

Note that for $r \in K_{\mathbb{R}}^+$, the distribution of $r \cdot x$ for $x \sim \mathcal{D}_{K_{\mathbb{R}}}(c, \varsigma)$ is $\mathcal{D}_{K_{\mathbb{R}}}(r \cdot c, (\sigma_i(r) \cdot \varsigma_i)_i)$. For a matrix $\mathbf{B} \in K_{\mathbb{R}}^{n \times n}$, we define $\det(\mathbf{B}) = \mathcal{N}(\det_{K_{\mathbb{R}}}(\mathbf{B}))$. We say that \mathbf{B} is orthogonal if $\mathbf{B}^{\dagger} \cdot \mathbf{B} = \mathbf{I}$, which implies that $\det(\mathbf{B}) = 1$. We let $\mathcal{O}_n(K_{\mathbb{R}})$ denote the set of orthogonal matrices. If a matrix $\mathbf{B} \in K_{\mathbb{R}}^{n \times n}$ has $K_{\mathbb{R}}$ -linearly independent columns (i.e., no non-trivial linear combination is zero), then it admits a QR-factorization $\mathbf{B} = \mathbf{Q}\mathbf{R}$ with $\mathbf{Q} \in \mathcal{O}_n(K_{\mathbb{R}})$ and $\mathbf{R} \in K_{\mathbb{R}}^{n \times n}$ upper triangular with diagonal elements in $K_{\mathbb{R}}^+$ (see, e.g., [LPSW19, Section 2.3]).

V.2.2 Rank-2 Modules with a Gap

The following result provides a lower bound on the probability that a rank-1 module $\mathbf{v} \cdot \mathcal{O}_K$ is primitive in a rank- k module M , when $\mathbf{v} \in M$ is sampled from a sufficiently wide Gaussian

distribution. Taking $M = \mathcal{O}_K^k$, this provides in particular a lower bound on the probability that k elements sampled independently of a Gaussian distribution in \mathcal{O}_K are relatively coprime. This result generalizes [SS13, Lemma 4.4], which proved the result for $k = 2$ and $M = \mathcal{O}_K^2$ (with a proof inspired from [Sit10]). The proof for the general case with rank- k modules is very similar to the special case $M = \mathcal{O}_K^2$, hence we postpone it to Appendix D.2.1. In this work, we will only use Lemma V.2.2 for modules of rank-2, however, for the sake of re-usability, we state and prove it for modules of arbitrary ranks.

Lemma V.2.2. *There exists an absolute polynomial P such that the following holds. For any $\delta \geq 0$, degree- d number field K , integer $k \geq 2$, rank- k module $M \subset K_{\mathbb{R}}^k$, if $\mathbf{c} \in \text{span}_{K_{\mathbb{R}}}(M)$ and $\varsigma > 0$ are such that $\|\mathbf{c}\| \leq \delta \cdot \varsigma$ and $\varsigma \geq \lambda_{kd}(M) \cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M)/\lambda_1(M))$, then it holds that*

$$\Pr_{\mathbf{v} \leftarrow D_{M, \varsigma, \mathbf{c}}}(\mathbf{v} \cdot \mathcal{O}_K \text{ is primitive in } M) \geq \frac{1}{4\zeta_K(k)},$$

where $\zeta_K(\cdot)$ is the Dedekind zeta function of K and the λ_i 's refer to the minima of the lattice $\Phi(M)$.

This can be used to show that we can use the QR-factorization to precisely describe rank-2 modules (see Appendix D.2 for a proof).

Lemma V.2.3. *Let M be a rank-2 module with gap $\gamma > 0$. Then M can be written as*

$$\frac{\mathcal{N}^{\frac{1}{2d}}(M)}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right),$$

where $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$, $r \in K_{\mathbb{R}}$, J_1 and J_2 are norm-1 ideals. We call this a QR-standard-form for M .

We note that there are multiple QR-standard forms for any module M , as units of \mathbb{C} can be transferred from the ideal coefficients to the matrix \mathbf{Q} . In the following section, we will be interested in modules with specific distributions expressed in terms of QR-standard forms. It will then be convenient to define a module by a (well-distributed) QR-standard form. Note that the modules we define this way have norm 1.

Definition V.2.4. *For any $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$, $\gamma > 0$, $r \in K_{\mathbb{R}}$ and norm-1 ideals J_1, J_2 , we define*

$$\text{QRSF-2-Mod}(\mathbf{Q}, \gamma, J_1, J_2, r) = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right).$$

We will use the following result on the first and last minimum of the dual of a rank-2 module with a gap. The proof is provided in Appendix D.2.

Lemma V.2.5. *Let M be a rank-2 module in $K_{\mathbb{R}}^2$ with gap $\gamma(M) \geq 1$. Then*

$$\begin{aligned} \lambda_{2d}(M^\vee) &\leq 2\sqrt{d} \cdot \gamma(M) \cdot \mathcal{N}(M)^{-\frac{1}{2d}} \\ \lambda_1(M^\vee)^{-1} &\leq 2d \cdot \gamma(M) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}}. \end{aligned}$$

V.3 New Tools on Module Lattices

In this section, we present new tools to manipulate module lattices. For the sake of re-usability, we describe them for modules of arbitrary ranks, but we will use them only in rank 2 in the reductions of the present work. The missing proofs of this section are available in Appendix D.3.

V.3.1 Module sparsification

An essential ingredient in the module randomization of Section V.5 is sparsification. In this subsection, we extend to modules the definition and some properties of sparsification over lattices [Kho06].

Definition V.3.1. *Let M a module, \mathfrak{p} a prime ideal, $\overline{\mathbf{b}^\vee} \in (M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and \mathbf{b}^\vee a lift of $\overline{\mathbf{b}^\vee}$ in M^\vee . The sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$ is the submodule*

$$M' = \{\mathbf{m} \in M, \langle \mathbf{b}^\vee, \mathbf{m} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}\}.$$

The submodule M' does not depend on the choice of the vector \mathbf{b}^\vee lifting $\overline{\mathbf{b}^\vee}$.

Note that $M \subseteq M' \subseteq \mathfrak{p}M$, implying that M' has the same rank as M . As showed by the following two lemmas, sparsification increases the module norm by a factor $\mathcal{N}(\mathfrak{p})$ and an arbitrary rank-1 submodule of M is not contained in M' (except with probability $\leq 1/\mathcal{N}(\mathfrak{p})$).

Lemma V.3.2. *Let M a module, \mathfrak{p} a prime ideal and $\overline{\mathbf{b}^\vee} \in (M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$. Let M' be the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$. Then $\mathcal{N}(M') = \mathcal{N}(\mathfrak{p}) \cdot \mathcal{N}(M)$.*

Lemma V.3.3. *Let M a rank- k module, \mathfrak{p} a prime ideal and $\mathbf{b}I$ a primitive rank-1 submodule of M . Let $\overline{\mathbf{b}^\vee}$ be uniformly distributed in $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and M' be the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$. Then $\mathbf{b}I \subseteq M'$ and, except with probability $1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^k$, we have $\mathbf{b}I \not\subseteq M'$.*

The following lemma states that a module sparsification can be efficiently computed. The algorithm generalizes the one for lattice sparsification, detailed, e.g., in [BSW16].

Lemma V.3.4. *There exists a polynomial-time algorithm taking as inputs an arbitrary pseudo-basis of $M \subset K_{\mathbb{R}}^k$, a prime ideal \mathfrak{p} and $\overline{\mathbf{b}^\vee} \in (M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and computing a pseudo-basis of the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$.*

V.3.2 Module rounding

In this section, we describe the `DualRound` algorithm that rounds a rank- k module contained in $K_{\mathbb{R}}^k$ into a module contained in \mathcal{O}_K^k (with a close geometry), in a way that does not depend on how the module in $K_{\mathbb{R}}^k$ was represented. We do that by sampling almost orthogonal vectors in the dual lattice, in a similar fashion to what was done in [BDPW20] in the context of ideal lattices. We believe that this technique of rounding via the dual might have other applications, especially in situations where one would like to have the analogue of an HNF basis for lattices with real coefficients.

`DualRound` is parameterized by a standard deviation parameter $\varsigma > 0$, a BKZ block-size $\beta \in \{2, \dots, kd\}$ and an error bound $\varepsilon > 0$. It starts by computing a short \mathbb{Z} -basis of C^\vee , by using a provable variant of the BKZ algorithm [Sch87, HPS11, GN08, ALNS20]. This offers different runtime-quality trade-offs. It then uses the discrete Gaussian sampler from Lemma II.1.12 with orthogonal center parameters \mathbf{t}_i and a well-chosen error parameter.

Lemma V.3.5. *Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank- k module $M \subset K_{\mathbb{R}}^k$. Let $\beta \in \{2, \dots, kd\}$, $\varepsilon > 0$, and ς be such that $\varsigma \geq (kd)^{kd/\beta+3/2} \cdot \lambda_{kd}(M^\vee)$. Algorithm `DualRound` runs in time polynomial in 2^β , $\log(\varsigma/\varepsilon)$ and the bitsize of its input. Further, on input (\mathbf{B}, \mathbb{I}) , `DualRound` $_{\varsigma, \beta, \varepsilon}$ outputs a matrix $\mathbf{Y} \in M_k(K_{\mathbb{R}})$ such that*

- $(\mathbf{Y} \cdot \mathbf{B}) \cdot \mathbb{I}$ is contained in \mathcal{O}_K^k ;

Algorithm V.3.1 Algorithm DualRound_{ϵ,β,ε}

Input: A pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank- k module $M \subset K_{\mathbb{R}}^k$.

- 1: Compute a \mathbb{Z} -basis of M^{\vee} ;
- 2: Run BKZ with block-size β on it to obtain a new \mathbb{Z} -basis \mathbf{C}^{\vee} of M^{\vee} ;
- 3: Set $R = \varepsilon^{-1} \cdot \sqrt{5kd}\zeta$;
- 4: For $i \in [k]$, set $\mathbf{t}_i = R \cdot \mathbf{e}_i$, where \mathbf{e}_i is the i -th canonical unit vector of $K_{\mathbb{R}}^k$;
- 5: For $i \in [k]$, sample $\mathbf{y}_i \leftarrow \widetilde{D}_{\mathbf{C}^{\vee}, \zeta, \mathbf{t}_i}$ using Lemma II.1.12 with error parameter 2^{-kd} ;
- 6: Return $\mathbf{Y} = (\mathbf{y}_1 | \dots | \mathbf{y}_k)^{\dagger}$.

- $\mathbf{Y} = R \cdot \mathbf{I}_k + \mathbf{E}$ for $R = \varepsilon^{-1} \sqrt{5kd}\zeta > 0$ and $\|e_{ij}\| \leq \varepsilon R$ for all $i, j \in [k]$.
- If $\mathbf{Y} = (\mathbf{y}_1 | \dots | \mathbf{y}_k)^{\dagger}$, it holds that $\text{SD}(\mathbf{y}_i, D_{M^{\vee}, \zeta, \mathbf{t}_i}) \leq 2^{-kd}$.

Moreover, if $(\mathbf{B}', \mathbb{I}')$ is another pseudo-basis of M and if \mathbf{Y}' is the output of DualRound given this pseudo-basis as input, then

$$\text{SD}(\mathbf{Y}, \mathbf{Y}') \leq 2^{-\Omega(kd)}.$$

Note that Lemma V.3.5 does not necessarily ensure that the matrix \mathbf{Y} is invertible, hence the module $\mathbf{Y} \cdot \mathbf{B} \cdot \mathbb{I}$ might not be of rank k . However, by choosing ε sufficiently small and using the second condition on \mathbf{Y} , one can make sure that \mathbf{Y} is indeed invertible. This is the purpose of Lemma V.3.6.

Lemma V.3.6. *Let $\mathbf{Y} \in K_{\mathbb{R}}^{k \times k}$ be such that $\mathbf{Y} = R \cdot \mathbf{I}_k + \mathbf{E}$ for some $R > 0$ and $\|e_{ij}\| \leq \varepsilon \cdot R$ for all $i, j \in [k]$. Assume that $\varepsilon \leq 1/(2k)$. Then \mathbf{Y} is invertible and we have $\mathbf{Y}^{-1} = R^{-1} \cdot \mathbf{I}_k + \mathbf{E}'$, with $\|e'_{ij}\| \leq (k+1) \cdot \varepsilon \cdot R^{-1}$ for all $i, j \in [k]$. Further, it holds that $\det(\mathbf{Y}) \in [(1 + (k+1)(k+2)\varepsilon)^{-d/2}, (1 + 3\varepsilon)^{d/2}] \cdot R^{kd}$.*

V.4 From mod-uSVP₂ to NTRU

In this section, we prove the following two results (the blue boxes highlight the differences between the two):

Theorem V.4.1. *Let K be a number field of degree d and let $\gamma^+ > 0$. There exists $q_0 = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ and an algorithm **uSVP-to-NTRU** such that the following holds. For any $q \geq q_0$, $\gamma_{\text{NTRU}} \geq \gamma'_{\text{NTRU}} > 1$, $\gamma_{\text{HSVP}} \geq \sqrt{d} \Delta_K^{1/(2d)}$, let*

$$\begin{aligned} \gamma_{\text{uSVP}} &= \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 16\sqrt{2} \cdot d^{3/2} \cdot \delta_K \\ \gamma'_{\text{uSVP}} &= \frac{\gamma'_{\text{NTRU}}}{\gamma_{\text{HSVP}}^{3/2} \cdot 150 \cdot d^{7/2} \cdot \delta_K^2}. \end{aligned}$$

For any distribution $\mathcal{D}_{\text{mod-uSVP}_2}$ over γ_{uSVP} -mod-uSVP₂ instances with gap $\leq \gamma^+$, let $\mathcal{D}_{\text{NTRU}}$ be the distribution **uSVP-to-NTRU**($\mathcal{D}_{\text{mod-uSVP}_2}, q, \gamma_{\text{HSVP}}$). We have four reductions

- from $(\mathcal{D}_{\text{mod-uSVP}_2}, \gamma_{\text{uSVP}})$ -mod-uSVP₂^{mod} to $(\mathcal{D}_{\text{NTRU}}, \gamma_{\text{NTRU}}, q)$ -NTRU^{mod};
- from γ_{uSVP} -wc-mod-uSVP₂^{mod} on modules with gap $\leq \gamma^+$ to $(\gamma_{\text{NTRU}}, q)$ -wc-NTRU^{mod};
- from $(\mathcal{D}_{\text{mod-uSVP}_2}, \gamma_{\text{uSVP}}, \gamma'_{\text{uSVP}})$ -mod-uSVP₂^{vec} to $(\mathcal{D}_{\text{NTRU}}, \gamma_{\text{NTRU}}, \gamma'_{\text{NTRU}}, q)$ -NTRU^{vec};

- from $(\gamma_{\text{uSVP}}, \gamma'_{\text{uSVP}})$ -wc-mod-uSVP₂^{vec} on modules with gap $\leq \gamma^+$ to $(\gamma_{\text{NTRU}}, \gamma'_{\text{NTRU}}, q)$ -wc-NTRU^{vec}.

Given access to an oracle solving γ_{HSVP} -id-HSVP, the four reductions run in time polynomial in their input size, in $\exp(\frac{d \log(d)}{\log(2q/q_0)})$ and in $\zeta_K(2)$.

Proof. See Appendix D.4.1. □

The outline of the reduction is given in Figure V.2. Note that the quantity $\zeta_K(2)$ may be exponential in d for some number fields (which may impact the running time of the reduction). In the case of power-of-two cyclotomic fields, it was proven in [SS13, Lemma 4.2] that $\zeta_K(2) = O(1)$. We also propose a version of Theorem V.4.1 whose running time does not depend on $\zeta_K(2)$, but with a higher approximation factor (still polynomial in d and $\Delta_K^{1/d}$).

Theorem V.4.2. *Let $\kappa > 0$ be the constant defined in Lemma D.5.2 and K a number field of degree d and let $\gamma^+ > 0$. There exists $q_0 = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ and an algorithm **uSVP-to-NTRU** such that the following holds. For any $q \geq q_0$, $\gamma_{\text{NTRU}} \geq \gamma'_{\text{NTRU}} > 1$, $\gamma_{\text{HSVP}} \geq \sqrt{d} \Delta_K^{1/(2d)}$, let*

$$\begin{aligned} \gamma_{\text{uSVP}} &= \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 16\sqrt{2} \cdot d^{\frac{3+\kappa}{2}} \cdot \Delta_K^{\frac{\kappa}{2d}} \cdot \delta_K \\ \gamma'_{\text{uSVP}} &= \frac{\gamma'_{\text{NTRU}}}{\gamma_{\text{HSVP}}^{3/2} \cdot 150 \cdot \Delta_K^{\frac{\kappa}{2d}} \cdot d^{\frac{7+3\kappa}{2}} \cdot \delta_K^2}. \end{aligned}$$

For any distribution $\mathcal{D}_{\text{mod-uSVP}_2}$ over γ_{uSVP} -mod-uSVP₂ instances with gap $\leq \gamma^+$, let $\mathcal{D}_{\text{NTRU}}$ be the distribution **uSVP-to-NTRU**($\mathcal{D}_{\text{mod-uSVP}_2}, q, \gamma_{\text{HSVP}}$). We have four reductions

- from $(\mathcal{D}_{\text{mod-uSVP}_2}, \gamma_{\text{uSVP}})$ -mod-uSVP₂^{mod} to $(\mathcal{D}_{\text{NTRU}}, \gamma_{\text{NTRU}}, q)$ -NTRU^{mod};
- from γ_{uSVP} -wc-mod-uSVP₂^{mod} on modules with gap $\leq \gamma^+$ to $(\gamma_{\text{NTRU}}, q)$ -wc-NTRU^{mod};
- from $(\mathcal{D}_{\text{mod-uSVP}_2}, \gamma_{\text{uSVP}}, \gamma'_{\text{uSVP}})$ -mod-uSVP₂^{vec} to $(\mathcal{D}_{\text{NTRU}}, \gamma_{\text{NTRU}}, \gamma'_{\text{NTRU}}, q)$ -NTRU^{vec};
- from $(\gamma_{\text{uSVP}}, \gamma'_{\text{uSVP}})$ -wc-mod-uSVP₂^{vec} on modules with gap $\leq \gamma^+$ to $(\gamma_{\text{NTRU}}, \gamma'_{\text{NTRU}}, q)$ -wc-NTRU^{vec}.

Given access to an oracle solving γ_{HSVP} -id-HSVP, the four reductions run in time polynomial in their input size, in $\exp(\frac{d \log(d)}{\log(2q/q_0)})$.

Proof. See Appendix D.5.2. □

We describe in the next subsection the steps to prove Theorem V.4.1. The proof of Theorem V.4.2 is postponed in Appendix D.5, as it is essentially the same. The missing proofs of this section are available in Appendix D.4.

V.4.1 Pre-conditioning the mod-uSVP₂ instance

In this section, we use algorithm **DualRound** to pre-process the input module and control its volume. In order to have the Hermite Normal Form of our integral module look like an NTRU instance, we slightly modify the geometry of our input module to make it have what we call the coprime property (see Definition V.4.3). Hence, we describe an algorithm, called **PreCond** (see Algorithm D.4.1), which combines all this and transform any mod-uSVP₂ instance (with a lower bounded gap) into a new mod-uSVP₂ instance with roughly the same geometry and with all the properties we will require in Section V.4.2.

Definition V.4.3 (Coprime property). *We say that a rank-2 module $M \subseteq \mathcal{O}_K^2$ has the coprime property if it holds that*

$$\{x \in \mathcal{O}_K \mid \exists y \in \mathcal{O}_K, (x, y)^T \in M\} = \mathcal{O}_K.$$

In other words, the module M has the coprime property if the ideal spanned by the first coordinate of all the vectors of M is equal to \mathcal{O}_K .

We note that having the coprime property is not very constraining. In fact, any module can be applied a small distortion in order to ensure the coprime property. This is formalized in Lemma V.4.4 below.

Lemma V.4.4. *Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank-2 module $M \subset K^2$ with gap $\gamma(M) \geq 1$. There exists some $V_0 > 0$ with $V_0^{1/(2d)} = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma(M))$ and an algorithm **PreCond** such that the following holds. Let $\beta \in \{2, \dots, 2d\}$ and $V > 0$ be such that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Then, on input (\mathbf{B}, \mathbb{I}) , V and β , algorithm **PreCond** outputs a matrix $\mathbf{Y} \in \text{GL}_2(K)$ such that*

- *if (\mathbf{B}, \mathbb{I}) is a γ_{uSVP} -mod-uSVP₂ instance, then $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a γ'_{uSVP} -mod-uSVP₂ instance for $\gamma'_{\text{uSVP}} = \gamma_{\text{uSVP}}/(2\sqrt{2})$;*
- *the rank-2 module $M' := \mathbf{Y}\mathbf{B} \cdot \mathbb{I}$ is contained in \mathcal{O}_K^2 ;*
- *$\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$;*
- *M' has the coprime property;*
- *$\mathbf{Y} = R \cdot \mathbf{I}_2 + \mathbf{E}$ for some $R = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)} > 0$ and $\|e_{ij}\| \leq R/5$ for all $1 \leq i, j \leq 2$.*

*Algorithm **PreCond** runs in expected time polynomial in its input bitsize, in 2^β and in $\zeta_K(2)$.*

Proof. See Appendix D.4.2. □

V.4.2 Transforming a mod-uSVP₂ instance into an NTRU instance

As the NTRU modules are free, the second step of our reduction finds a free module containing our mod-uSVP₂ instance and transforms it into an NTRU instance. For this purpose, we use the **BalanceIdeal** algorithm (cf Algorithm D.4.2) that takes as input any fractional ideal I and uses a γ_{HSVP} -id-HSVP oracle to output a balanced element x such that $\langle x \rangle$ contains I but is not much larger than it.

Lemma V.4.5. *There exists an algorithm **BalanceIdeal** that takes as input a fractional ideal $I \subset K$ and a parameter $\gamma_{\text{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, and outputs an element $x \in K$ such that $I \subseteq \langle x \rangle$ and $|\sigma_i(x)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$ for all $i \leq d$, where $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(I)^{-1/d}$.*

Moreover, given access to a γ_{HSVP} -id-HSVP oracle, it runs in polynomial time and makes one call to the γ_{HSVP} -id-HSVP oracle.

Proof. See Appendix D.4.3. □

We can now describe our algorithm transforming a mod-uSVP₂ instance into an NTRU instance: Algorithm V.4.1. The operations done by this algorithm are summarised in Figure V.2 and proven in Lemma V.4.7.

Algorithm V.4.1 Algorithm Conditioned-to-NTRU**Input:** A pseudo-basis $\mathbf{B}_1 \cdot \mathbb{I}$ of a rank-2 module in \mathcal{O}_K^2 and some parameters q and γ_{HSVP} **Output:** A basis \mathbf{B}_4 of a free rank-2 module and some auxiliary information \mathbf{aux}

- 1: Compute the HNF pseudo-basis
- $\mathbf{B}_2 \cdot \mathbb{J}$
- of the rank-2 module spanned by
- $\mathbf{B}_1 \cdot \mathbb{I}$

Let $a = \mathbf{B}_2[1, 0]$

$$\triangleright \# \mathbf{B}_2 = \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix}$$

- 2: Sample
- $b \leftarrow \text{BalanceIdeal}(J_2, \gamma_{\text{HSVP}})$

- 3: Compute
- $h = \lfloor a \cdot q/b \rfloor$

- 4:
- Return**
- $\mathbf{B}_4 = \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix}$
- and
- $\mathbf{aux} = (a, b, J_1, J_2)$

Module	Pseudo-basis	Short vector
M_1	$\begin{bmatrix} I_1 & I_2 \\ \left(\begin{array}{cc} b_{11} & b_{12} \\ b_{21} & b_{22} \end{array} \right) \end{bmatrix}$	$\mathbf{s}_1 = \begin{pmatrix} u \\ v \end{pmatrix}$
	$\begin{array}{c} \downarrow \text{Step 1} \\ \text{HNF} \\ \downarrow \end{array}$	
$M_2 = M_1$	$\begin{bmatrix} J_1 & J_2 \\ \left(\begin{array}{cc} 1 & 0 \\ a & 1 \end{array} \right) \end{bmatrix}$	$\mathbf{s}_2 = \mathbf{s}_1$
	$\begin{array}{c} \downarrow \text{Step 2} \\ \text{Principalization} \\ \downarrow \end{array}$	
$M_3 \supseteq M_2$	$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \left(\begin{array}{cc} 1 & 0 \\ a & b \end{array} \right) \end{bmatrix}$	$\mathbf{s}_3 = \mathbf{s}_2$
	$\begin{array}{c} \downarrow \text{Step 3} \\ \text{distorsion} \\ \text{+ rounding} \\ \downarrow \end{array}$	
M_4	$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \left(\begin{array}{cc} 1 & 0 \\ \lfloor a \cdot q/b \rfloor & q \end{array} \right) \end{bmatrix}$	$\mathbf{s}_4 = \begin{pmatrix} u \\ v \cdot q/b - u \cdot \{a \cdot q/b\} \end{pmatrix}$

Figure V.2: Outline of algorithm Conditioned-to-NTRU.

Lemma V.4.6. Let $\gamma_{\text{HSVP}} \geq \sqrt{d} \Delta_K^{1/(2d)}$, $q \in \mathbb{Z}_{>0}$ and (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank-2 module $M \subseteq \mathcal{O}_K^2$. Assume that we have access to a γ_{HSVP} -id-HSVP oracle. On input γ_{HSVP}, q and (\mathbf{B}, \mathbb{I}) , algorithm Conditioned-to-NTRU runs in polynomial time in the bitsize of its input and makes one call to the γ_{HSVP} -id-HSVP oracle.

Proof. See Appendix D.4.4. □

Lemma V.4.7. Let $\gamma_{\text{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, $\gamma_{\text{NTRU}} > 1$ and $q \in \mathbb{Z}_{>0}$ be some parameters. Define

$$V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$$

$$\text{and } \gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 8 \cdot d^{3/2} \cdot \delta_K.$$

Let (\mathbf{B}, \mathbb{I}) be any $\gamma_{\text{uSVP}}\text{-mod-uSVP}_2$ instance in \mathcal{O}_K^2 , with the coprime property and with norm in $[1/2^{2d} \cdot V, 2^{2d} \cdot V]$. Then on input $(\mathbf{B}, \mathbb{I}), \gamma_{\text{HSVP}}, q$, the algorithm `Conditioned-to-NTRU` outputs $(\mathbf{B}_4, \mathbf{aux})$ such that \mathbf{B}_4 is a $(\gamma_{\text{NTRU}}, q)$ -NTRU instance.

Proof. See Appendix D.4.5. □

The `aux` information output by algorithm `Conditioned-to-NTRU` will be used in the Algorithms D.4.3 and D.4.4 to lift any short vector / dense submodule from the NTRU instance back to the `mod-uSVP2` instance. The proofs of Lemmas V.4.6 and V.4.7 are available in Appendices D.4.4 and D.4.5 respectively.

V.4.3 Lifting back short vectors and dense submodules

In this section, we prove that using the auxiliary information `aux` produced by Algorithm `Conditioned-to-NTRU`, one can lift a short vector or a densest submodule from the output NTRU instance back to the input `mod-uSVP2` instance. The proofs of Lemmas V.4.8 and V.4.9 are available in Appendices D.4.6 and D.4.7 respectively.

Lemma V.4.8. *There exists an algorithm `LiftMod` such that the following holds. Let q, γ_{HSVP} and (\mathbf{B}, \mathbb{I}) be as in Lemma V.4.7. Let M_1 denote the rank-2 module generated by $(\mathbf{B}, \mathbb{I}), [\mathbf{C}, (a, b, J_1, J_2)] \leftarrow \text{Conditioned-to-NTRU}((\mathbf{B}, \mathbb{I}), q, \gamma_{\text{HSVP}})$ and let M_4 denote the rank-2 free module generated by \mathbf{C} .*

Let (\mathbf{v}, J) be a pseudo-basis of the densest rank-1 submodule of M_4 . Then, on input $a, b, (\mathbf{C}, \mathcal{O}_K^2)$ and (\mathbf{v}, J) , algorithm `LiftMod` outputs $\mathbf{w} \in K$ such that $\text{span}_K(\mathbf{w}) \cap M_1$ is the densest rank-1 submodule of M_1 .

Moreover, algorithm `LiftMod` runs in polynomial time.

Proof. See Appendix D.4.6. □

Lemma V.4.9. *There exists an algorithm `LiftVec` such that the following holds. Let q, γ_{HSVP} and (\mathbf{B}, \mathbb{I}) be as in Lemma V.4.7. Let M_1 denote the rank-2 module generated by $(\mathbf{B}, \mathbb{I}), [\mathbf{C}, \mathbf{aux}] \leftarrow \text{Conditioned-to-NTRU}((\mathbf{B}, \mathbb{I}), q, \gamma_{\text{HSVP}})$ and let M_4 denote the rank-2 free module generated by \mathbf{C} .*

Let $\mathbf{s} \in M_4$. Then, on input $\mathbf{aux}, \gamma_{\text{HSVP}}, (\mathbf{C}, \mathcal{O}_K^2)$ and \mathbf{s} , algorithm `LiftVec` outputs a vector $\mathbf{t} \in M$ such that $\|\mathbf{t}\| \leq \|\mathbf{s}\| \cdot 68 \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$.

If given access to a $\gamma_{\text{HSVP-id-HSVP}}$ oracle, algorithm `LiftVec` runs in polynomial time and makes 1 call to the oracle.

Proof. See Appendix D.4.7. □

Combining all the results of this section, one can prove Theorem V.4.1.

V.5 Randomization of Rank-2 Modules with Gaps

A rank-2 module with a gap can, by Lemma V.2.3 and the fact that densest submodules are primitive, be written as $M = \mathbf{u} \cdot J_1 + \mathbf{v} \cdot J_2$ where $\mathbf{u} \cdot J_1$ is a densest rank-1 submodule of M . Informally, the goal of this section is to randomize $\mathbf{u}, \mathbf{v}, J_1, J_2$ without changing the gap too much. The missing proofs of this section are available in Appendix D.6.

We first describe the average-case distribution we are considering. Note that the gap parameter γ' is itself a random variable.

Definition V.5.1. Let $\gamma > 0$ and $B > 2$. We define the distribution $D_{B,\gamma}^{\text{module}}$ over rank-2 and norm-1 modules by:

$$D_{B,\gamma}^{\text{module}} = \text{QRSF-2-Mod}(\mathbf{Q}, \gamma', I_1, I_2, r),$$

where

- the matrix \mathbf{Q} is uniform in $\mathcal{O}_2(K_{\mathbb{R}})$;
- the gap parameter γ' is set as $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)} / B^{1/(2d)}$ with $(a, c) \in K_{\mathbb{R}}^2$ distributed as $\chi_{K_{\mathbb{R}}} \times \mathcal{D}(0, 1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$;
- the ideals I_1, I_2 are uniform in \mathcal{I}_1 (the set of norm-1 ideals);
- the element r is uniform in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I_1 I_2^{-1}$.

We now state the main theorem of this section, which can be viewed as a worst-case to average-case reduction for rank-2 modules with a gap.

Theorem V.5.2 (Assuming ERH). For all $B \geq (d^d \Delta_k)^{\Omega(1)}$ and $\gamma \geq B^{1/(2d)}$ there exists a procedure Randomize_B that runs in time polynomial in $\log B$ and the bitsize of its input, and such that on input a pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 and norm-1 module M of gap γ outputs a pair $((\mathbf{B}', \mathbb{I}'), \mathbf{aux})$ such that

- the pseudo-basis $(\mathbf{B}', \mathbb{I}')$ spans a rank-2 and norm-1 module M' ;
- any event that holds for $D_{B,\gamma}^{\text{module}}$ with probability $\varepsilon \geq 2^{-o(d)}$ also holds for M' with probability $\Omega(\varepsilon^4)$ over the internal randomness of Randomize_B .

Further, there exists a deterministic algorithm Recover that runs in time polynomial in the bitsize of its input such that for M' as above, if U' is a densest rank-1 submodule of M' , then $\text{Recover}(U', \mathbf{aux})$ returns the densest rank-1 submodule of M , with probability $1 - 2^{-\Omega(d)}$ over the randomness of Randomize_B .

We emphasize that the theorem does not state that the output distribution of Randomize_B is $D_{B,\gamma}^{\text{module}}$, but only that they are close in the sense that any event that holds with sufficient probability for $D_{B,\gamma}^{\text{module}}$ also holds for the output distribution of Randomize_B with a polynomially related probability.

Randomize_B is described in Algorithm V.5.6. It consists of two main steps: a coefficient randomization (described in Section V.5.1), whose purpose is to randomize the ideal coefficients; and a geometric randomization (described in Section V.5.2), whose purpose is to randomize the pseudo-basis matrix. Section V.5.3 compares the distribution that would ideally be returned by the composition of the coefficient and geometric randomizations, with the distribution of the pseudo-basis in Definition V.5.1. Finally, we complete the proof of Theorem V.5.2 in Section V.5.4.

V.5.1 Coefficient randomization

In the coefficient randomization step, our aim is to randomize the ideal coefficients of a good pseudo-basis (i.e., whose first pair corresponds to the densest rank-1 submodule), given an arbitrary pseudo-basis of a rank-2 module. One may multiply the whole pseudo-basis by a random ideal, but this only randomizes the pair of ideal coefficients. More precisely, this leaves the ratio of the ideal coefficients unchanged. To decouple the ideal coefficients, we use module sparsification, as described in Section V.3. This first step towards coefficient randomization is formally

Algorithm V.5.1 Partial Coefficient Randomization: `Partial-CRB`

Input: A pseudo-basis of a rank-2 module M .

- 1: Sample \mathfrak{p} uniformly among prime ideals of norms $\leq B$;
- 2: Sample $\bar{\mathbf{b}}^\vee$ uniformly in $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$;
- 3: Return a pseudo-basis of the sparsification of M by $(\bar{\mathbf{b}}^\vee, \mathfrak{p})$ along with \mathfrak{p} .

described in Algorithm V.5.1. Steps 1 and 3 are respectively performed using Lemmas II.2.11 and V.3.4.

Theorem V.5.3 (Assuming ERH). *Let $B \geq (\log \Delta_K)^{\Omega(1)}$. The runtime of `Partial-CRB` is polynomial in $\log B$ and the bitsize of its input. Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank-2 module M , and let $(J_1, \mathbf{u}), (J_2, \mathbf{v})$ be an arbitrary pseudo-basis of M . Let M' be the rank-2 module spanned by the pseudo-basis output by `Partial-CRB` when given (\mathbf{B}, \mathbb{I}) as input, let $\bar{\mathbf{b}}^\vee$ be the element of $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ sampled in `Partial-CRB` and let \mathbf{b}^\vee be a lift of $\bar{\mathbf{b}}^\vee$ in M^\vee . Then, with probability $1 - (1/B)^{\Omega(1)}$, we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p}J_1^{-1}$. In that case, there exists $x \in J_1J_2^{-1}$ such that*

$$M' = \mathbf{u} \cdot \mathfrak{p}J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2.$$

Assume further that $\gamma(M) \geq B^{1/(2d)}$ and that $\mathbf{u} \cdot J_1$ is the densest rank-1 submodule of M . Then, still when $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p}J_1^{-1}$, we have that $\gamma(M') = \gamma(M)/\mathcal{N}(\mathfrak{p})^{1/(2d)} > 1$ and $\mathbf{u} \cdot \mathfrak{p}J_1$ is the densest rank-1 submodule of M' .

The result follows from Lemmas V.5.4 and V.5.5, whose proofs are postponed to Appendix D.6.

Lemma V.5.4. *Borrowing the notations of Theorem V.5.3, we have*

$$\mathbf{u} \cdot \mathfrak{p}J_1 \subset M' \quad \text{and} \quad \mathbf{u} \cdot J_1 \not\subset M',$$

with probability $1 - (1/B)^{\Omega(1)}$ over the choices of \mathfrak{p} and $\bar{\mathbf{b}}^\vee$.

Lemma V.5.5. *Borrowing the notations of Theorem V.5.3 and assuming that $\mathbf{u} \cdot J_1 \not\subset M'$, there exists $x \in J_1J_2^{-1}$ such that $(\mathbf{v} + x\mathbf{u}) \cdot J_2 \subset M'$.*

We now describe the coefficient randomization. Ideally, we would have access to a pseudo-basis $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ of the module M under scope, for which the densest rank-1 submodule is $\mathbf{u} \cdot J_1$. We would multiply J_1 by a random ideal and J_2 by another random ideal. Unfortunately, given only access to an arbitrary pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of M , this seems difficult to achieve obliviously. Instead, we use algorithm `Ideal-Sample` (Algorithm II.2.1) to obtain a uniform norm-1 ideal I , and multiply M by it. This will obviously multiply J_1 and J_2 by I . As this distribution is invariant by ideal multiplication, the ideal $J_2I/\mathcal{N}(J_2)^{1/d}$ will be uniform among norm-1 ideals. It remains to obliviously randomize the first ideal independently of the second one. For this purpose, we use `Partial-CR` (Algorithm V.5.1), which has the effect of obliviously multiplying the first ideal with a random prime ideal \mathfrak{p} while leaving the second one unchanged (with overwhelming probability). Note that multiplying by a random prime ideal is the main component of the ideal randomization algorithm `Ideal-Sample`. In a sense, this ‘‘almost’’ randomizes J_1 .

Algorithm V.5.2 describes the process on the input basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$. The corresponding randomization performed on the hidden pseudo-basis $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ is described in Algorithm V.5.3. Note that there is no need for Algorithm V.5.3 to be efficient as its sole purpose is to describe the behavior of Algorithm V.5.2 on the hidden pseudo-basis.

Algorithm V.5.2 Real Coefficient Randomization: $\text{Real-CR}_{B,B'}$

Input: A pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of a module $M \subset K_{\mathbb{R}}^2$.

- 1: Let $((I'_1, \mathbf{b}'_1), (I'_2, \mathbf{b}'_2)), \mathfrak{p}$ be the output of Partial-CR_B on input $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$;
- 2: Sample \mathfrak{q} using $\text{Ideal-Sample}_{B'}$;
- 3: Let $\mathbf{b}''_i = \mathbf{b}'_i / \mathcal{N}(\mathfrak{p})^{1/(2d)}$ for $i \in [2]$;
- 4: Return $((qI'_1, \mathbf{b}''_1), (qI'_2, \mathbf{b}''_2)), \mathfrak{p}, \mathfrak{q}$.

Algorithm V.5.3 Ideal Coefficient Randomization: Ideal-CR_B

Input: $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}}), \gamma > 1, J_1, J_2$ ideals of norm 1, $r \in K_{\mathbb{R}}$;

- 1: Let $M = \text{QRSF-2-Mod}(\mathbf{Q}, \gamma, J_1, J_2, r)$;
- 2: Let $\mathbf{u} = 1/\gamma \cdot \mathbf{Q} \cdot (1, 0)^T$ and $\mathbf{v} = \gamma \cdot \mathbf{Q} \cdot (r, 1)^T$;
- 3: Sample \mathfrak{p} uniformly among prime ideals of norms $\leq B$;
- 4: Sample \mathbf{b}^{\vee} in M^{\vee} , uniform in $M^{\vee}/\mathfrak{p}M^{\vee}$ conditioned on $\langle \mathbf{b}^{\vee}, \mathbf{u} \rangle_{K_{\mathbb{R}}} \notin \mathfrak{p}J_1^{-1}$;
- 5: Find $x \in J_1 J_2^{-1}$ such that $\langle \mathbf{b}^{\vee}, \mathbf{v} + x \cdot \mathbf{u} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}J_2^{-1}$;
- 6: Sample J uniformly among norm-1 ideals;
- 7: Return $(\mathbf{Q}, \gamma/\mathcal{N}(\mathfrak{p})^{1/(2d)}, J_1 J_2^{-1} J \mathfrak{p} / \mathcal{N}^{1/d}(\mathfrak{p}), J, r + x)$.

In Theorem V.5.6, we show that the resulting distributions on the output modules are statistically close, and describe the evolution of the densest rank-1 submodule.

Theorem V.5.6 (Assuming ERH). *Assume that $B' \geq (d^d \Delta_K)^{\Omega(1)}$ and $B \geq (\log \Delta_K)^{\Omega(1)}$. The runtime of $\text{Real-CR}_{B,B'}$ is polynomial in $\log(BB')$ and the bitsize of its input.*

Let $M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) \subset K_{\mathbb{R}}^2$ a module with norm 1, in QR-standard form.

Then the distribution of the module output by $\text{Real-CR}_{B,B'}$ on input an arbitrary pseudo-basis of M is within statistical distance $(1/B)^{\Omega(1)} + 2^{-d}$ of $\text{QRSF-2-Mod}(\text{Ideal-CR}_B(\mathbf{Q}, \gamma, J_1, J_2, r))$.

Assume further that $\gamma \geq B^{1/(2d)}$ and let U denote the densest rank-1 submodule of M . Let $(M', \mathfrak{p}, \mathfrak{q})$ be the output of $\text{Real-CR}_{B,B'}$ on input M . Then, with probability $1 - (1/B)^{\Omega(1)}$, we have that $\gamma(M') = \gamma(M)/\mathcal{N}(\mathfrak{p})^{1/(2d)} > 1$ and the densest rank-1 submodule of M' is

$$\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}} \cdot U \cdot \mathfrak{q} \frac{\mathfrak{p}}{\mathcal{N}^{1/d}(\mathfrak{p})}.$$

V.5.2 Geometric randomization

In the geometric module randomization, we will use a distribution D_{distort} over $K_{\mathbb{R}}^{2 \times 2}$ whose purpose is to distort the geometric relationship between the densest rank-1 submodule and the complementing rank-1 submodule of the rank-2 module under scope. We define D_{distort} as $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^{2 \times 2}$ conditioned on the event that $|\det(\sigma_i(\mathbf{D}))| > 1/d$ holds for all $i \in [d]$.

The following lemmas describe useful properties of the distribution D_{distort} .

Lemma V.5.7. *The following properties hold.*

- The distribution D_{distort} can be sampled from in time polynomial in d .
- The distribution D_{distort} is invariant by multiplication on the left and the right by matrices in $\mathcal{O}_2(K_{\mathbb{R}})$.

Lemma V.5.8. *Let D be the distribution over $K_{\mathbb{R}}^{2 \times 2}$ of*

$$\mathbf{Q} \cdot \begin{pmatrix} a & b \\ 0 & c \end{pmatrix}$$

where $\mathbf{Q} \leftarrow \mathcal{U}(\mathcal{O}_2(K_{\mathbb{R}}))$, $a \leftarrow \chi_{K_{\mathbb{R}}}$ and $b, c \leftarrow \mathcal{D}_{K_{\mathbb{R}}}(0, 1)$, conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$. Then $D = D_{\text{distort}}$.

Let $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ be a pseudo-basis of a rank-2 module M . Assume that $\mathbf{u} \cdot J_1$ is the densest rank-1 submodule, but that we have access to this pseudo-basis only indirectly, via an arbitrary pseudo-basis of M . Write

$$(\mathbf{u}|\mathbf{v}) = \mathbf{Q} \cdot \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix},$$

for some $r \in K_{\mathbb{R}}$. The purpose of the geometric randomization is to map r to some r' that is uniform modulo $J_1 J_2^{-1}$, while at the same time not distorting the module M too much, so that the randomized M still has a gap and its rank-1 densest submodule is related to $\mathbf{u} \cdot J_1$. For this purpose, we multiply M on the left by a matrix sampled from D_{distort} . For the analysis, it is convenient to take it Gaussian, and to avoid a potentially large distortion, we avoid matrix samples with small determinant. This corresponds to algorithm **Real-GR** (Algorithm V.5.4). The effect on the hidden pseudo-basis $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ is described in algorithm **Ideal-GR** (Algorithm V.5.5). In Theorem V.5.9, we show that the resulting module distributions are identical, and describe the evolution of the densest rank-1 sublattice.

Algorithm V.5.4 Real Geometric Randomization: **Real-GR**

Input: A pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of a norm-1 module $M \subset K_{\mathbb{R}}^2$.

- 1: Sample $\mathbf{D} \leftarrow D_{\text{distort}}$ (using Lemma V.5.7);
 - 2: $(\mathbf{b}'_1|\mathbf{b}'_2) \leftarrow \det(\mathbf{D})^{-1/(2d)} \cdot \mathbf{D} \cdot (\mathbf{b}_1|\mathbf{b}_2)$;
 - 3: Return $((I_1, \mathbf{b}'_1), (I_2, \mathbf{b}'_2)), \mathbf{D}$.
-

Algorithm V.5.5 Ideal Geometric Randomization: **Ideal-GR**

Input: $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$, $\gamma > 1$, J_1, J_2 ideals of norm 1, $r \in K_{\mathbb{R}}$;

- 1: Sample $a \leftarrow \chi_{K_{\mathbb{R}}}$ and $c \leftarrow \mathcal{D}(0, 1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$;
 - 2: Sample $b \leftarrow \mathcal{D}(0, 1)$;
 - 3: Sample $\mathbf{Q}' \leftarrow \mathcal{U}(\mathcal{O}_2(K_{\mathbb{R}}))$;
 - 4: Set $J'_1 = a/\mathcal{N}^{1/d}(a) \cdot J_1$ and $J'_2 = c/\mathcal{N}^{1/d}(c) \cdot J_2$;
 - 5: Set $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)}$;
 - 6: Set $r' = (b + ar)/c$;
 - 7: Return $(\mathbf{Q}', \gamma', J'_1, J'_2, r')$.
-

Theorem V.5.9. *Algorithm **Real-GR** runs in polynomial time. Let*

$$M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) \subset K_{\mathbb{R}}^2$$

*a module with norm 1, in QR-standard-form. Let M' be the module spanned by the output of **Real-GR** on input an arbitrary pseudo-basis of M . Then the distribution of M' is identical to the distribution $\text{QRSF-2-Mod}(\text{Ideal-GR}(\mathbf{Q}, \gamma, J_1, J_2, r))$.*

Further, if $\gamma > d$ and U is the densest rank-1 submodule of M , then, with probability $1 - 2^{-\Omega(d)}$, we have $\gamma(M') > 1$ and the densest rank-1 submodule of M' is $\det(\mathbf{D})^{-1/(2d)} \cdot \mathbf{D} \cdot U$, where \mathbf{D} is the Gaussian matrix sampled during the execution of **Real-GR**.

V.5.3 On the Ideal-GR \circ Ideal-CR distribution

We define a few probability distributions over the inputs of **QRSF-2-Mod**, which we will use to show that the operations performed on the available arbitrary pseudo-basis randomize the rank-2 module, so that the input module is “forgotten” in the output module distribution while at the same time controlling the evolution of the densest rank-1 submodule.

Definition V.5.10. Let $B \geq 2$ and $\gamma > 0$. We consider the following random variables, which are assumed independent (unless stated otherwise).

- \mathbf{Q} uniform in $\mathcal{O}_2(K_{\mathbb{R}})$;
- $b \in K_{\mathbb{R}}$ distributed as $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)$;
- $(a, c) \in K_{\mathbb{R}}^2$ distributed as $\chi_{K_{\mathbb{R}}} \times \mathcal{D}_{K_{\mathbb{R}}}(0, 1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$; we define $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)} / B^{1/(2d)}$;
- \mathfrak{p} uniform among prime ideals of norms $\leq B$;
- I_1, I_2, J uniform in \mathcal{I}_1 (the set of norm-1 ideals);
- $\zeta \in E$ sampled from the centered normal law of standard deviation $d^{-3/2}$, conditioned on $\|\zeta\| \leq 1/d$;
- u uniform in $\{x \in K_{\mathbb{R}}, \forall i \in [d] : |\sigma_i(x)| = 1\}$;
- r' uniform in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I_1 I_2^{-1}$.

Let $J_1, J_2 \in \mathcal{I}_1$ and $r \in K_{\mathbb{R}}$ arbitrary. Let x be as in Step 5 of **Ideal-CR_B**, when given as input $(\mathbf{Q}, \gamma, J_1, J_2, r)$ and with the variable \mathfrak{p} of **Ideal-CR_B** being the random variable above. In order to simplify the notations, we define the random variable:

$$I(J_1, J_2) = \mathcal{N}^{\frac{1}{d}} \left(\frac{c}{a} \right) \cdot \frac{au}{c \text{Exp}(\zeta)} \cdot J_1 J_2^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{1/d}(\mathfrak{p})} \in \mathcal{I}_1.$$

Let $r''(J_1, J_2)$ be uniformly distributed in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I(J_1, J_2) \cdot J^{-1}$.

$$D_{B,\gamma}^{\text{rand}} = D_{B,\gamma}^{(1)} \xrightarrow{\text{RD}_2=O(1)} D_{B,\gamma}^{(2)} \xrightarrow{\text{RD}_2=O(1)} D_{B,\gamma}^{(3)} \xleftarrow{\text{SD}=2^{-\Omega(d)}} D_{B,\gamma}^{(4)} \xleftarrow{\text{SD}=2^{-\Omega(d)}} D_{B,\gamma}^{\text{target}}$$

Figure V.3: The relations between the distributions of Definition V.5.10, proved in Lemmas D.6.1, D.6.2, D.6.3, D.6.4 and D.6.6. Here $D \xrightarrow{\text{RD}_2=O(1)} D'$ means $\text{RD}_2(D' \parallel D) = O(1)$ and $D \xrightarrow{\text{SD}=2^{-\Omega(d)}} D'$ means $\text{SD}(D, D') = 2^{-\Omega(d)}$.

We define the following distributions of the form $(\tilde{\mathbf{Q}}, \tilde{\gamma}, \tilde{I}_1, \tilde{I}_2, \tilde{r})$, where the random variables \tilde{r} is defined modulo $\tilde{\gamma}^{-2} \cdot \tilde{I}_1 \cdot \tilde{I}_2^{-1}$:

$$\begin{aligned} D_{B,\gamma}^{\text{rand}} &: \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \frac{a}{\mathcal{N}^{1/d}(a)} J_1 J_2^{-1} J \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, \frac{c}{\mathcal{N}^{1/d}(c)} \cdot J, \frac{b+a(r+x)}{c} \right), \\ D_{B,\gamma}^{(1)} &: \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{au}{c} \cdot J_1 J_2^{-1} J \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, J, u \frac{b+a(r+x)}{c} \right), \\ D_{B,\gamma}^{(2)} &: \left(\mathbf{Q}, \gamma \cdot \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, I(J_1, J_2), J, u \frac{b+a(r+x)}{c \text{Exp}(\zeta)} \right), \\ D_{B,\gamma}^{(3)} &: \left(\mathbf{Q}, \gamma', I(J_1, J_2), J, \frac{B^{\frac{1}{d}}}{\mathcal{N}^{1/d}(\mathbf{p})} \cdot u \frac{b+a(r+x)}{c \text{Exp}(\zeta)} \right), \\ D_{B,\gamma}^{(4)} &: \left(\mathbf{Q}, \gamma', I(J_1, J_2), J, r''(J_1, J_2) \right), \\ D_{B,\gamma}^{\text{target}} &: \left(\mathbf{Q}, \gamma', I_1, I_2, r' \right). \end{aligned}$$

Note that $D_{B,\gamma}^{\text{rand}}$ is the distribution obtained by composing **Ideal-CR_B** (Algorithm V.5.3) and **Ideal-GR** (Algorithm V.5.5), on an input of the form $(\mathbf{Q}_0, \gamma, J_1, J_2, r)$ with (γ, J_1, J_2, r) as above and $\mathbf{Q}_0 \in \mathcal{O}_2(K_{\mathbb{R}})$ arbitrary. These algorithms significantly randomize the QR-standard form, but it still depends on (J_1, J_2, r) . On the other hand, the distribution $D_{B,\gamma}^{\text{target}}$ is independent of (J_1, J_2, r) . Our goal is to show that these two distributions are similar, in the sense that any event that holds with some probability $\varepsilon \geq 2^{-o(d)}$ for one holds with probability $\varepsilon^{O(1)}$ for the other one.

For this purpose, we consider the intermediate (hybrid) distributions of Definition V.5.10. To help the reader, we use two colours in the definition of the successive distributions. The entries of the tuples that are in red are those that change compared to the previous distribution. The variables with blue background are those that depend on (J_1, J_2, r) . The relations between the distributions of Definition V.5.10 are pictorially summarized in Figure V.3. The lemmas formally stating these relations and their proofs are provided in Appendix D.6. Some of the relations require $B \geq (d^d \Delta_K)^{\Omega(1)}$ or $\gamma \geq d^{1/4} \cdot \Delta_K^{1/(2d)}$.

V.5.4 Full module randomization

The full randomization algorithm **Randomize_B** (Algorithm V.5.6) is the composition of algorithms **Real-CR** and **Real-GR**.

Let $((\mathbf{B}', \mathbb{I}'), \mathbf{aux})$ be an output of **Randomize_B**, and U' be a rank-1 submodule of the module spanned by $(\mathbf{B}', \mathbb{I}')$. We define:

$$\text{Recover}(U', \mathbf{aux} = (\mathbf{p}, \mathbf{q}, D)) = (\mathcal{N}(\mathbf{p}) \cdot \det(\mathbf{D}))^{\frac{1}{2d}} \cdot \mathbf{D}^{-1} \cdot U' \cdot \mathbf{q}^{-1} \mathbf{p}^{-1}.$$

With these choices of algorithms **Randomize_B** and **Recover**, we can finally prove Theorem V.5.2. For this purpose, we show that the module distribution that is output from the ran-

Algorithm V.5.6 (Real) Full Randomization: `RandomizeB`

Input: A pseudo-basis (\mathbf{B}, \mathbb{I}) of a norm-1 module $M \subset K_{\mathbb{R}}^2$.

- 1: Apply `Real-CRB, (d^d \Delta_K)^{\Omega(1)}` to (\mathbf{B}, \mathbb{I}) and let $((\mathbf{B}^\circ, \mathbb{I}^\circ), \mathbf{p}, \mathbf{q})$ be the output;
- 2: Apply `Real-GR` to $(\mathbf{B}^\circ, \mathbb{I}^\circ)$ and let $((\mathbf{B}', \mathbb{I}'), \mathbf{D})$ be the output;
- 3: Return $((\mathbf{B}', \mathbb{I}'), \mathbf{aux})$ with $\mathbf{aux} = (\mathbf{p}, \mathbf{q}, \mathbf{D})$.

domination algorithm (on an arbitrary input) and the distribution $D_{B, \gamma}^{\text{module}}$ from Definition V.5.1 are close in the mixed “SD plus RD” sense of Figure V.3. The full proof is available in Appendix D.6.1.

V.6 Random Self-Reducibility of Module uSVP

The main result of this section is the worst-case to average-case reduction for $\text{mod-uSVP}_2^{\text{mod}}$ of Theorem V.6.1.

Theorem V.6.1 (Assuming ERH). *There exist $\gamma_0 = (d\Delta_K^{1/d})^{O(1)}$ and $(D_\gamma^{\text{mod-uSVP}_2})_{\gamma \geq \gamma_0}$ a family of distributions such that the following properties hold for any $\gamma \geq \gamma_0$:*

- if $\gamma \leq (2^d \Delta_K^{1/d})^{O(1)}$, then $D_\gamma^{\text{mod-uSVP}_2}$ can be sampled from in time polynomial in $\log \Delta_K$;
- with probability $1 - 2^{-\Omega(d)}$, a sample from $D_\gamma^{\text{mod-uSVP}_2}$ is a pseudo-basis of a rank-2 module $M \subseteq \mathcal{O}_K^2$ with gap $\gamma(M) \geq \gamma \cdot \sqrt{d} \Delta_K^{1/(2d)}$; in particular, these are γ -mod-uSVP₂ instances;
- there exists a Karp reduction from γ' -wc-mod-uSVP₂^{mod} to $(D_\gamma^{\text{mod-uSVP}_2}, \gamma)$ -mod-uSVP₂^{mod}, with $\gamma' = \gamma \cdot (d \cdot \Delta_K^{1/d})^{O(1)}$; the reduction runs in time polynomial in $\log \Delta_K$ and the input bitsize.

Note that the restriction on γ for the first condition is very mild, as in this parameter range, $\text{mod-uSVP}_2^{\text{mod}}$ can be solved in polynomial time using the LLL algorithm [LLL82]. We now proceed in two steps. We first define and study the distribution $D^{\text{mod-uSVP}_2}$, and then prove Theorem V.6.1.

V.6.1 A distribution over uSVP instances

Let $\gamma > 1$. The distribution $D_\gamma^{\text{mod-uSVP}_2}$ is defined as follows:

- sample a module from $D_{B, \gamma'}^{\text{module}}$ along with a pseudo-basis (\mathbf{B}, \mathbb{I}) , with $B = (d^d \Delta_K)^{O(1)}$ and $\gamma' = 2\gamma \cdot \sqrt{d} \Delta_K^{1/(2d)} \cdot \sqrt{d} B^{1/d}$ (see Definition V.5.1) and using `Ideal-Sample` to sample from \mathcal{I}_1 ;
- call `DualRound\varsigma, \beta, \varepsilon` (\mathbf{B}, \mathbb{I}) with $\varsigma = (2^d \Delta_K^{1/d})^{O(1)}$, $\beta = 2$ and $\varepsilon = 1/(2d)^{3/2}$, and let \mathbf{Y} denote the output;
- return `HNF` $(\mathbf{Y} \cdot \mathbf{B}, \mathbb{I})$.

The first two statements of Theorem V.6.1 are implied by the following lemmas, whose proofs can be found in Appendix D.7.

Lemma V.6.2. *A sample M from $D_{B,\gamma'}^{\text{module}}$ has gap $\gamma(M) \geq \gamma'/(\sqrt{d}B^{1/d})$, with probability $1 - 2^{-\Omega(d)}$.*

Using the latter result and Lemma V.2.5, we obtain that the assumptions of Lemma V.3.5 are satisfied. This implies that the above sampling algorithm runs in time polynomial in $\log \Delta_K$. By Lemmas V.3.5 and V.3.6, the output is a pseudo-basis of a rank-2 module in \mathcal{O}_K^2 .

Lemma V.6.3. *Let $\gamma > 2$. Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank-2 module M with gap γ . Let \mathbf{Y} denote the output of $\text{DualRound}_{\varsigma,\beta,\varepsilon}(\mathbf{B}, \mathbb{I})$ with $\varsigma = \gamma \cdot (2d)^{2d+3}$, $\beta = 2$ and $\varepsilon = 1/(2d)^{3/2}$. Then the module spanned by $(\mathbf{Y} \cdot \mathbf{B}, \mathbb{I})$ has gap $\geq \gamma/2$.*

The definition of $D_\gamma^{\text{mod-uSVP}_2}$ and Lemmas V.6.2 and V.6.3 implies that the modules whose pseudo-basis are sampled from $D_\gamma^{\text{mod-uSVP}_2}$ have gap $\geq \gamma \cdot \sqrt{d}\Delta_K^{1/(2d)}$, and hence are γ -mod-uSVP₂ instances with overwhelming probability.

V.6.2 Reducing worst-case instances to $D^{\text{mod-uSVP}_2}$ instances

We first introduce intermediate problems, that will allow us to split the reduction into several steps.

Definition V.6.4. *Let $\gamma > 1$. A γ -mod-uSVP ^{\mathcal{N}} instance consists in a pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 module $M \subset K^2$ such that $\gamma(M) \geq \gamma$.*

Let \mathcal{D} a distribution over γ -mod-uSVP ^{\mathcal{N}} instances. The (D, γ) -mod-uSVP ^{\mathcal{N}, mod} problem asks, given as input a sample (\mathbf{B}, \mathbb{I}) from \mathcal{D} , to recover a densest rank-1 submodule of the module spanned by (\mathbf{B}, \mathbb{I}) .

The variant γ -wc-mod-uSVP ^{\mathcal{N}, mod} asks to solve this problem for any γ -mod-uSVP ^{\mathcal{N}} instance.

The γ^\approx -wc-mod-uSVP ^{\mathcal{N}, mod} variant is the restriction of γ -wc-mod-uSVP ^{\mathcal{N}, mod} to the instances of γ -mod-uSVP ^{\mathcal{N}} whose spanned modules M satisfy $\gamma(M) \in [\gamma, \gamma \cdot (1 + 1/d)]$.

Note that worst-case wc-mod-uSVP ^{\mathcal{N}, mod} reduces to wc-mod-uSVP ^{\mathcal{N}, mod} as the existence of a short non-zero vector implies the one of a dense rank-1 module. Similarly, mod-uSVP ^{\mathcal{N}, mod} reduces to mod-uSVP ^{mod} with a loss of a $(\sqrt{d}\Delta_K^{1/d})$ factor in the parameters, thanks to Minkowski's theorem. To prove the third statement of Theorem V.6.1, it hence suffices to reduce wc-mod-uSVP ^{\mathcal{N}, mod} to mod-uSVP ^{\mathcal{N}, mod} for distribution $D_\gamma^{\text{mod-uSVP}_2}$. The result follows from Lemmas V.6.5 and V.6.7.

The first lemma states that to solve γ -wc-mod-uSVP ^{\mathcal{N}, mod} (in which the gap is only bounded from below), then it suffices to solve γ^\approx -wc-mod-uSVP ^{\mathcal{N}, mod} (in which the gap is almost known). It relies on sparsification.

Lemma V.6.5 (Assuming ERH). *Let $\gamma, \gamma' > 1$ satisfying $\gamma' \geq 2 \log(\Delta_K)^{O(1/d)} \cdot \gamma$. Then the problem γ' -wc-mod-uSVP ^{\mathcal{N}, mod} reduces to γ^\approx -wc-mod-uSVP ^{\mathcal{N}, mod} . The reduction runs in time polynomial in $(\log \Delta_K)^{O(1)}$ and its input bitsize and succeeds with probability $\Omega(1/(d^2 + \log \Delta_K))$.*

Using the Rényi divergence, it is possible to relate the success probability of an algorithm towards solving mod-uSVP ^{\mathcal{N}, mod} for samples from $D_\gamma^{\text{mod-uSVP}_2}$ with the same probability for $D_{\gamma'}^{\text{mod-uSVP}_2}$, when γ and γ' are sufficiently close.

Lemma V.6.6. *Let $\gamma, \gamma', \gamma'' > 1$ with $\gamma' \in \gamma \cdot [1, 1 + 1/d]$ and $\gamma'' = \gamma/(d\Delta_K^{1/d})^{O(1)}$. Then any algorithm that solves $(D_\gamma^{\text{mod-uSVP}_2}, \gamma'')$ -mod-uSVP ^{\mathcal{N}, mod} with probability ε also solves $(D_{\gamma'}^{\text{mod-uSVP}_2}, \gamma'')$ -mod-uSVP ^{\mathcal{N}, mod} with probability $\Omega(\varepsilon^2)$.*

Equipped with the latter result, we are now able to state the worst-case to average case component of the reduction.

Lemma V.6.7 (Assuming ERH). *Let $\gamma, \gamma', \gamma'' > 1$ with $\gamma' = \gamma \cdot (d\Delta_K^{1/d})^{O(1)}$ and $\gamma'' = \gamma / (d\Delta_K^{1/d})^{O(1)}$. Then there is a reduction from γ^\approx -wc-mod-uSVP $_2^{\mathcal{N}, \text{mod}}$ to $(D_{\gamma'}^{\text{mod-uSVP}_2}, \gamma'')$ -mod-uSVP $_2^{\mathcal{N}, \text{mod}}$. The reduction runs in time polynomial in $\log \Delta_K$ and the input bitsize, and if the $(D_{\gamma'}^{\text{mod-uSVP}_2}, \gamma'')$ -mod-uSVP $_2^{\mathcal{N}, \text{mod}}$ oracle succeeds with probability $\varepsilon \geq 2^{-o(d)}$, then the reduction succeeds with probability $\varepsilon^{O(1)}$.*

Chapter VI

Conclusion and Perspectives

VI.1 Summary of Contributions

We give a summary of the contributions of this manuscript in Figure VI.1. An arrow from A to B means that A reduces to B , dashed arrows indicate quantum reductions and blue arrows are contributions of this work. An arrow with a star means that the reduction needs a id-HSVP oracle.




Figure VI.1: Summary of the reductions proven in this manuscript.

Overall, in this thesis, we worked in the direction of giving a broader understanding of structured lattice problems. We worked on low-rank modules, i.e., with rank 1 or 2.

On id-HSVP. We gave a general reduction from id-HSVP on ideals to id-HSVP on their inverses. More precisely, we showed that (up to solving id-HSVP on uniform small-norm integral ideals), solving id-HSVP on a set of ideals and solving it on their inverses is computationally the same. This implies, in particular, that solving SVP on random small prime ideals is enough to solve it for any ideal of a given number field (using the reduction of [Gen09]). This result can be seen in two ways: the “attacker way” which is that in order to solve id-HSVP, only focusing on prime ideals of small norm is enough, or the “protocol-designer way”, that such “algebraically natural” distribution of ideals yields hard SVP instances.

This result should not be seen, in our opinion, as a security statement, since cryptographic protocols based on id-HSVP are not really the ones that are used in practice now, and (to our

knowledge) no attack on id-HSVP has been used to break protocols whose security relies on related problems, such as Ring-SIS.

The reduction of [PS21] and our result yield a new distribution over NTRU instances with polynomial modulus whose security is based on id-HSVP in the worst-case, namely the distribution yielded by sampling a random prime ideal and running the reduction of [PS21]. This distribution is not suited for cryptographic purposes yet, since a short vector in the random ideal should be sampled along with it, in order to be used as a private key.

On mod-uSVP₂ and NTRU. We showed that the problems NTRU and mod-uSVP₂ are computationally equivalent for some range of parameters. This indicates that the biggest particularity of NTRU modules compared to other more generic modules are their gap: NTRU modules are representative of generic mod-uSVP₂ modules. This gives an indication about the fact that the NTRU problem is not “ad hoc”, in the sense that it is closely related to other more generic module-lattice security problems.

We proposed a worst-case to average-case self-reduction for mod-uSVP₂ using id-HSVP oracle calls. The average-case distribution over mod-uSVP₂ can efficiently be sampled. This, along with the previous reduction, gives a new distribution over NTRU instances.

Understanding the gap between rank-1 and rank-2 modules. Taking a step back, working on mod-uSVP₂ can be seen as working on modSVP instances whose difficulty lies between the rank-1 and the rank-2 cases. As said in introduction, there seems to be a big gap of difficulty between rank-1 and rank-2 modSVP. When the gap of a rank-2 module M is larger than $2^{O(d)}$ (where d is the degree of the underlying number field), the LLL algorithm allows finding the densest submodule N of M , and finding a short vector in M reduces to finding a short vector in N , which can be done with a call to a id-HSVP oracle. In the other direction, 1-mod-uSVP₂ is exactly modSVP₂. The γ -mod-uSVP₂ problem then gives a difficulty gradient between id-HSVP and modSVP₂.

VI.2 Perspective and open problems

Understanding and improvement of the reductions of this work. The reductions presented Figure VI.1 do not compose, since the distribution yielded by sampling random ideals and applying [PS21] has no reason to match with the distribution obtained from taking a mod-uSVP₂ instance from $D_{B,\gamma}^{\text{module}}$ and applying the mod-uSVP₂ to NTRU reduction. In particular, in order to work in polynomial-time, the reduction of Theorem V.4.1 needs q to be exponential in the degree of the field. A possible improvement on our results would be to study discrepancies between the two distributions. A first step could be to try to extend the reduction of Theorem V.4.1 to smaller values of q (e.g., $q = \text{poly}(d)$).

A better understanding of the difference between those distributions could also lead to removing the id-HSVP oracle call used in the reduction from the average-case mod-uSVP₂ to average-case NTRU.

Find how to efficiently sample hard NTRU instances. As said before, in order to use the NTRU distribution described in Corollary IV.6.2 ($D_{\text{NTRU}}^{A,q}$) in a cryptographic context, one would have to sample the associated prime ideal along with a small vector in it. Currently, the only way to do it is to use a factoring oracle (and hence a quantum computer), which is not of interest in the post-quantum cryptographic setting, where the protocols should not require a quantum computer. An improvement compared to our work would be to sample such

distribution without having to factor integers, which could be done in two different ways: either by sampling a prime ideal along with a small vector in it without factoring oracle (said otherwise, implementing a function such as `SampleWithTrap` without quantum computers); or by relating the distribution $D_{\text{NTRU}}^{A,q}$ and the more classical NTRU distribution used in cryptography (where h is computed from f and g that are sampled beforehand). Maybe ways of sampling f and g could be found such that the distribution of $h = g/f \bmod q$ would relate to the distribution described in Corollary IV.6.2.

Another direction in sampling hard NTRU instances would be to understand how to randomize an NTRU instance h in a reversible way. This would lead to a worst-case to average-case reduction for NTRU without having to rely on intermediate problems.

Study other distributions for modSVP instances. The distribution $D_{B,\gamma}^{\text{module}}$ of Definition V.5.1 is a natural distribution over the set of mod-uSVP₂ instances with gap approximately γ , in the sense that almost all elements composing the mod-uSVP₂ instances are sampled from a uniform distribution. This kind of distribution could be generalized to distributions over other kinds of module lattices.

In the same way that the set of all rank- n lattices is isomorphic to the set $\text{GL}_n(\mathbb{R})/\text{GL}_n(\mathbb{Z})$, the set of module-lattices of rank k and norm 1 is isomorphic to the set

$$\{(\mathbf{B}, \mathbb{I}), \mathbf{B} \in \text{GL}_k(K_{\mathbb{R}}), \mathbb{I} \in (\text{IdLat}_K^0)^k, \det(\mathbf{B}) = 1\} / \sim,$$

where $(\mathbf{B}, \mathbb{I}) \sim (\mathbf{B}', \mathbb{I}')$ if they span the same module. Even if it can be proven that this set is not compact, it is locally compact. It would be interesting to study the link between (rounding of) Haar distributions on compact subsets of it and more computationally friendly distributions such as Module-SIS modules [LS15]. A first step could be to study the link between the distribution $D_{B,\gamma}^{\text{module}}$ and the Haar distribution over the set of module lattices of gap γ .

Generalization to mod-NTRU. The problem mod-NTRU _{m,k} has been first introduced in 2019 [CKKS19] as a generalization of NTRU: instead of taking $h = g/f \bmod q$, the public key is a matrix $\mathbf{H} \in \mathcal{O}_K^{m \times k}$ satisfying $\mathbf{H} = \mathbf{F}^{-1} \cdot \mathbf{G} \bmod q$ for $\mathbf{F} \in \mathcal{O}_K^{m \times m}$ and $\mathbf{G} \in \mathcal{O}_K^{m \times k}$ small-norm matrices. It has a natural interpretation in terms of modules lattices of rank $m+k$, but its link standard problems about them has not been widely studied yet. It would be interesting to study the place of mod-NTRU _{m,k} relatively to other module lattices problems. The first reduction to generalize would be the one of [PS21]. We think that it might be generalized to a reduction from mod-SIVP _{m} to mod-NTRU _{m,k} (the analogue of id-HSVP for high rank modules is mod-SIVP). The second reduction to generalize would be the one presented in this manuscript, namely a reduction from mod-uSVP₂ to generalizations of mod-NTRU. It should be noted that, in contrast with mod-uSVP₂, the module associated with mod-NTRU _{m,k} has m short vectors inside of it. A more subtle control over the parameters of the module is to be expected.

Other computational problems relying on module lattices. In 2022, a new family of lattice-based computational problems is proposed by Ducas and van Woerden [DW22], the Lattice Isomorphism Problem, which asks to decide whether two bases span the same lattice up to rotation. Soon after, a module version is proposed to be used in practice in the context of signatures [DPPW22]. This module version (in the rank-2) was cryptanalysed recently [MPPW24] in the case of real number fields, and its analysis is still an active research question.

It would be interesting to relate LIP (respectively Module-LIP) to more standard lattices (respectively module-lattices) problems.

Better bounds on the error of $N_K(\cdot)$. This last direction has far fewer cryptographic implications than the others, but nevertheless raises my curiosity. The bound on the error on the approximation $N_K(X) \sim \rho_K \cdot X$ depends on $X^{1-\eta}$ with $\eta = \Theta(1/\ln(d))$. This choice of η appears in the proof in order to balance the bound of Theorem III.1.8. It would be interesting to see if η can be taken to go to zero slower (or even $\eta = \Omega(1)$), or to see if one can reach better bounds by constraining the considered field (Galois, or even Abelian). It would also be interesting to characterize families of fields where using the heavy machinery of Chapter III is not necessary. For example to find fields where good bounds on ρ_K or $\zeta_K(2)$ are known (for example, it is known that $\zeta_K(2)$ is bounded when K is a power-of-two cyclotomic field [SS13, Lemma 4.2]).

Bibliography

- [AB09] S. Arora and B. Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- [ABD16] M. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions. In *CRYPTO*, 2016.
- [ABD⁺19] R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. Crystals-kyber algorithm specifications and supporting documentation. *NIST PQC Round*, 2019.
- [AD17] M. R. Albrecht and A. Deo. Large Modulus Ring-LWE \geq Module-LWE. In *ASIACRYPT*, 2017.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems. *STOC*, 1996.
- [Ajt98] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, 1998.
- [AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, New York, NY, USA, 2001.
- [Ale09] AlexanderAIUS. Graphen.jpg, 2009. [<https://commons.wikimedia.org/wiki/File:Graphen.jpg>; accessed 20-08-2024].
- [ALNS20] D. Aggarwal, J. Li, P. Q. Nguyen, and N. Stephens-Davidowitz. Slide reduction, revisited - filling the gaps in SVP approximation. In *CRYPTO*, 2020.
- [Apo98] T. M. Apostol. *Introduction to analytic number theory*. 1998.
- [AR05] D. Aharonov and O. Regev. Lattice problems in $\text{NP} \cap \text{coNP}$. *Journal of the ACM*, 2005.
- [Bab86] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 1986.
- [Bac90] E. Bach. Explicit bounds for primality testing and related problems. 1990.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Math Ann*, 1993.
- [BCLV17] D. Bernstein, C. Chuengsatiansup, T. Lange, and C. v. Vredendaal. NTRU prime: reducing attack surface at low cost. *SAC*, 2017.

- [BDPW20] K. Boer, L. Ducas, A. Pellet-Mary, and B. Wesolowski. Random self-reducibility of Ideal-SVP via Arakelov random walks. In *CRYPTO*, 2020.
- [BEP22] K. Boudgoust, G. E., and A. Pellet-Mary. Some easy instances of Ideal-SVP and implications on the partial Vandermonde knapsack problem. In *CRYPTO*, 2022.
- [BL94] J. A. Buchmann and H. W. Lenstra. Computing maximal orders and factoring over \mathbb{Z}_p . *Preprint*, 1994.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, 2013.
- [Boe22] K. Boer. *Random Walks on Arakelov Class Groups*. PhD thesis, Leiden University, 2022. Available on request from the author.
- [BS96] E. Bach and J. O. Shallit. *Algorithmic Number Theory: Efficient Algorithms*. MIT Press, 1996.
- [BST⁺20] M. Bhargava, A. Shankar, T. Taniguchi, F. Thorne, J. Tsimerman, and Y. Zhao. Bounds on 2-torsion in class groups of number fields and integral points on elliptic curves. *Journal of the AMS*, 2020.
- [BSW16] S. Bai, D. Stehlé, and W. Wen. Improved reduction from the bounded distance decoding problem to the unique shortest vector problem in lattices. In *ICALP*, 2016.
- [CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT*, 2016.
- [CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *EUROCRYPT*, 2017.
- [CDW21] R. Cramer, L. Ducas, and B. Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *J ACM*, 2021.
- [CJL16] J. H. Cheon, J. Jeong, and C. Lee. An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero. *LMS Journal of Computation and Mathematics*, 2016.
- [CKKS19] J. Cheon, D. Kim, T. Kim, and Y. Son. A new trapdoor over module-NTRU lattice and its application to ID -based encryption. Cryptology ePrint Archive, Paper 2019/1468, 2019. URL <https://eprint.iacr.org/2019/1468>.
- [Coh93] H. Cohen. *A course in computational algebraic number theory*. Springer, 1993.
- [Coh00] H. Cohen. *Advanced Topics in Computational Number Theory*. Springer, 2000.
- [CS97] D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *EUROCRYPT*, 1997.
- [DPPW22] L. Ducas, E. W. Postlethwaite, L. N. Pulles, and W. v. Woerden. Hawk: Module LIP makes lattice signatures fast, compact and simple. (2022/1155), 2022. URL <https://eprint.iacr.org/2022/1155>. Publication info: Preprint.
- [DW22] L. Ducas and W. v. Woerden. On the lattice isomorphism problem, quadratic forms, remarkable lattices, and cryptography. In *EUROCRYPT*, 2022.

- [Emd] P. v. Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. *Technical Report, Department of Mathematics, University of Amsterdam*.
- [FP85] U. Fincke and M. Pohst. Improved methods for calculating vectors of short length in a lattice , including a complexity analysis. *Mathematics of Computation*, 1985.
- [FPS22] J. Felderhoff, A. Pellet-Mary, and D. Stehlé. On module unique-SVP and NTRU. In *ASIACRYPT*, 2022.
- [FPSW23] J. Felderhoff, A. Pellet-Mary , D. Stehlé, and B. Wesolowski. Ideal-SVP is hard for small-norm uniform prime ideals. In *TCC*, 2023.
- [Fri89] E. Friedman. Analytic formulas for the regulator of a number field. *Inventiones mathematicae*, 1989.
- [FS10] C. Fieker and D. Stehlé. Short bases of lattices over number fields. In *Algorithmic Number Theory*, 2010.
- [Gal12] S. D. Galbraith. *Mathematics of public key cryptography*. Cambridge University Press, 2012.
- [GAL13] M. Gil, F. Alajaji, and T. Linder. Rényi divergence measures for commonly used univariate continuous distributions. *Inform Sciences*, 2013.
- [Gee14] Geek3. mplwp_lambert_w_branches.svg, 2014. [https://commons.wikimedia.org/wiki/File:Mplwp_lambert_W_branches.svg; accessed 25-07-2024].
- [Gen09] C. Gentry. *A Fully Homomorphic Encryption Scheme*. PhD thesis, Stanford University, 2009.
- [Gen10] C. Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, 2010.
- [GMSS99] O. Goldreich, D. Micciancio, S. Safra, and J.-P. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Information Processing Letters*, 1999.
- [GN08] N. Gama and P. Q. Nguyen. Finding short lattice vectors within Mordell’s inequality. In *STOC*, 2008.
- [GNR10] N. Gama, P. Q. Nguyen, and O. Regev. *Lattice Enumeration Using Extreme Pruning*. 2010.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
- [HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In *ANTS*, 1998.
- [HPS11] G. Hanrot, X. Pujol, and D. Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *CRYPTO*, 2011.
- [HS07] G. Hanrot and D. Stehlé. Improved analysis of kannan’s shortest lattice vector algorithm. In *CRYPTO 2007*, 2007.

- [Kan87] R. Kannan. Minkowski's convex body theorem and integer programming. *Mathematics of Operations Research*, 1987.
- [KF17] P. Kirchner and P.-A. Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In *EUROCRYPT*, 2017.
- [Kho06] S. Khot. Hardness of approximating the shortest vector problem in high ℓ_p norms. *Journal of Computer and System Sciences*, 2006.
- [Lan13] S. Lang. *Algebraic number theory*. Springer, 2013.
- [LDK⁺20] V. Lyubashevsky, L. Ducas, E. Kiltz, T. Lepoint, P. Schwabe, G. Seiler, D. Stehlé, and S. Bai. Crystals-dilithium. *Algorithm Specifications and Supporting Documentation*, 2020.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math Ann*, 1982.
- [LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP*, 2006.
- [Lou00] S. Louboutin. Explicit bounds for residues of Dedekind zeta functions, values of L-functions at $s=1$, and relative class numbers. *Journal of Number Theory*, 2000.
- [LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, 2010.
- [LPSW19] C. Lee, A. Pellet-Mary, D. Stehlé, and A. Wallet. An LLL algorithm for module lattices. In *ASIACRYPT*, 2019.
- [LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Design Code and Cryptography*, 2015.
- [Mai00] C. Maire. On infinite unramified extensions. *Pacific Journal of Mathematics*, 2000.
- [MG02] D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Springer, 2002.
- [Mic02] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *FOCS*, 2002.
- [Mic08] D. Micciancio. Efficient reductions among lattice problems. *SODA '08*, 2008.
- [Mic18] D. Micciancio. On the hardness of learning with errors with binary secrets. *Theory of Computing*, 2018.
- [MP13] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *CRYPTO 2013*, 2013.
- [MPPW24] G. Mureau, A. Pellet-Mary, G. Pliatsok, and A. Wallet. Cryptanalysis of Rank-2 Module-LIP in totally real number fields. In *EUROCRYPT*, 2024.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *Foundations of Computer Science*, 2004.

- [MR07] D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 2007.
- [MV13] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. *SIAM Journal on Computing*, 2013.
- [Nam] N. Nam. Lattice challenge - SVP challenge. [<https://latticechallenge.org/svp-challenge/>; accessed 06-09-2024].
- [Neu13] J. Neukirch. *Algebraic number theory*. Springer, 2013.
- [NIST] I. T. L. Computer Security Division. Selected algorithms 2022 - post-quantum cryptography, 2022. URL <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [Pei16] C. Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 2016.
- [PHS19] A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-SVP in ideal lattices with pre-processing. In *EUROCRYPT*, 2019.
- [PML21] C. Porter, A. Mendelsohn, and C. Ling. Subfield algorithms for Ideal- and Module-SVP based on the decomposition group. *IACR Cryptol. ePrint Arch.*, 2021.
- [PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.
- [PRS17] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *STOC*, 2017.
- [PS21] A. Pellet-Mary and D. Stehlé. On the hardness of the NTRU problem. In *ASIACRYPT*, 2021.
- [PT21] D. Platt and T. Trudgian. The Riemann hypothesis is true up to $3 \cdot 10^{12}$. *Bulletin of the London Mathematical Society*, 2021.
- [PXWC21] Y. Pan, J. Xu, N. Wadleigh, and Q. Cheng. On the ideal shortest vector problem over random rational primes. In *EUROCRYPT*, 2021.
- [Rad59] H. Rademacher. On the Phragmén-Lindelöf theorem and some applications. *Mathematische Zeitschrift*, 1959.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
- [Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 1987.
- [SE94] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical Programming*, 1994.
- [Sho94] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *FOCS*, 1994.

- [Sim23] A. Simonič. Estimates for L-functions in the critical strip under GRH with effective applications. *Mediterranean Journal of Mathematics*, 2023.
- [Sit10] B. D. Sittinger. The probability that random algebraic integers are relatively r -prime. *J Number Theory*, 2010.
- [SS13] D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure as standard worst-case problems over ideal lattices. In *EUROCRYPT*, 2013.
- [SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, 2009.
- [Ste15] N. Stephens-Davidowitz. Dimension-preserving reductions between lattice problems. Available at <http://noahsd.com/latticeproblems.pdf>, 2015.
- [Ten95] G. Tenenbaum. *Introduction à la théorie analytique et probabiliste des nombres : Cours et exercices*. Dunod, 1995.
- [Web08] H. Weber. Lehrbuch der algebra, vol. II. *Vieweg und Sohn, Braunschweig*, 1908.

Appendix A

Appendices of Chapter II

A.1 Missing Proofs

A.1.1 Proof of Lemma II.2.12

Let $\zeta' \in E$ sampled from the centered normal law with standard deviation $d^{-3/2}$, $z' = \exp(\zeta')$. We use the notations from [BDPW20] and instantiate [BDPW20, Theorem 3.3] with $\varsigma = d^{-3/2}$, $\varepsilon = 2^{-d}$, $N = 1$ and

$$k = \frac{\Theta(d \log d) + \log(\text{Vol}(\text{Pic}_K^0))}{\log d}.$$

By the log-unit lattice smoothing analysis from [BDPW20, Appendix B.1], the condition on N in [BDPW20, Theorem 3.3] is satisfied. Now, note that the bound on $\text{Vol}(\text{Pic}_K^0)$ in [BDPW20, Lemma 2.3] implies that $k \leq O(d + \log \Delta_K / \log d)$. Therefore, our lower bound on B implies the one in [BDPW20, Theorem 3.3]. By [BDPW20, Theorem 3.3], we deduce that the distribution of the projection of $z' \cdot \mathfrak{p} / \mathcal{N}^{1/d}(\mathfrak{p})$ into Pic_K^0 is within 2^{-d} statistical distance from $\mathcal{U}(\text{Pic}_K^0)$, implying by Gaussian tail-bounds that the distribution of $\text{Exp}(\zeta) \cdot \mathfrak{p} / \mathcal{N}^{1/d}(\mathfrak{p})$ is within $2^{-\Omega(d)}$ statistical distance from $\mathcal{U}(\text{Pic}_K^0)$. The proof can be completed by using [BDPW20, Lemma 2.7]. \square

A.1.2 Equivalence of the conditions in Definition II.3.3

Assume that N is maximal for the inclusion. Let $m = \text{rank}(M)$ and $k = \text{rank}(N)$ and write $N = \sum_{i \in [k]} \mathbf{c}_i J_i$. By [FS10, Theorem 4], there exists a pseudo-basis $(\mathbf{b}_i, I_i)_{i \in [m]}$ of M such that $\text{span}_{i \in [k]}(\mathbf{b}_i I_i) = \text{span}_{i \in [k]}(\mathbf{c}_i J_i) = \text{span}(N)$. By maximality of N this implies that $N = \sum_{i \in [k]} \mathbf{b}_i \cdot I_i$. Taking $N' = \sum_{i > k} \mathbf{b}_i \cdot I_i$ allows to conclude that $M = N + N'$ and $\text{rank}(M) = \text{rank}(N) + \text{rank}(N')$.

Now, assume that there is a module N' with $M = N + N'$ and $\text{rank}(M) = \text{rank}(N) + \text{rank}(N')$. As $N \subseteq M$, we have $N \subseteq M \cap \text{span}_K(N)$. Further, by the rank equality, we must have $N' \cap \text{span}_K(N) = \{\mathbf{0}\}$. Then we have

$$N \subseteq M \cap \text{span}_K(N) = (N + N') \cap \text{span}_K(N) = N \cap \text{span}_K(N) \subseteq N.$$

Finally, assume that $N = M \cap \text{span}_K(N)$. Let P with $\text{rank}(P) = \text{rank}(N)$ and $N \subseteq P$. We have that $\text{span}_K(N) \subseteq \text{span}_K(P)$, and hence $\text{span}_K(N) = \text{span}_K(P)$ by equality of the dimensions. Then we have

$$N \subseteq P \subseteq M \cap \text{span}_K(P) = M \cap \text{span}_K(N) = N.$$

This completes the equivalency proof. \square

A.1.3 Proof of Lemma II.3.4

Let $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ be a pseudo-basis of M and write $N = \mathbf{s}_1 J_1$ and $N' = \mathbf{s}_2 J_2$. There exists $\mathbf{Z} \in K^{2 \times 2}$ such that $\mathbf{S} = \mathbf{BZ}$. Assume by contradiction that $\text{span}_K(N') \neq \text{span}_K(N)$. In that case, the matrix \mathbf{Z} has rank 2, the vectors \mathbf{s}_1 and \mathbf{s}_2 are $K_{\mathbb{R}}$ -linearly independent and $M' = \mathbf{s}_1 J_1 + \mathbf{s}_2 J_2$ is a submodule of M . By using a QR-factorization $\mathbf{S} = \mathbf{QR}$, one sees that $\det(\mathbf{S}) = \mathcal{N}(r_{11})\mathcal{N}(r_{22})$ and $\mathcal{N}(M') \leq \mathcal{N}(N)\mathcal{N}(N')$. We hence obtain:

$$\mathcal{N}(M) \leq \mathcal{N}(M') \leq \mathcal{N}(N)\mathcal{N}(N') < \frac{\sqrt{\mathcal{N}(M)}}{\gamma^d} \left(\gamma^d \sqrt{\mathcal{N}(M)} \right) = \mathcal{N}(M),$$

which gives a contradiction. We thus have that $\text{span}_K(N') = \text{span}_K(N)$. Definition II.3.3 allows us to conclude that $N' \subseteq N$.

Assume now that $\gamma > 1$. Then the first statement implies that the densest rank-1 submodule N is unique. Let $\mathbf{b} \in M$ with $0 < \|\mathbf{b}\| < \gamma \cdot \mathcal{N}(M)^{1/(2d)}$. Then $\mathbf{b}\mathcal{O}_K$ is a rank-1 submodule of M and

$$\mathcal{N}(\mathbf{b}\mathcal{O}_K) \leq \|\mathbf{b}\|^d < \gamma^d \cdot \sqrt{\mathcal{N}(M)}.$$

By the above, we must have $\mathbf{b}\mathcal{O}_K \subseteq N$, which is equivalent to $\mathbf{b} \in N$. \square

A.1.4 Proof of Lemma II.3.5

Let k denote the rank of M . By Minkowski's theorem, there exists a non-zero vector in M of ℓ_2 -norm $\leq \sqrt{kd} \cdot (\det M)^{1/(kd)}$. By considering the rank-1 module that it spans, we obtain that $\lambda_1^{\mathcal{N}}(M) \leq (kd)^{d/2} \cdot (\det M)^{1/k}$. Now, by using Minkowski's theorem again, we obtain that all rank-1 submodules of norm $\leq (kd)^{d/2} \cdot (\det M)^{1/k}$ contain a non-zero vector of M of ℓ_2 -norm $\leq \sqrt{kd} \cdot \Delta_K^{1/(2d)} \cdot (\det M)^{1/(kd)}$. By discreteness of M , the non-zero vectors $\{\mathbf{s}_i\}_{i \geq 1}$ of M with ℓ_2 -norm $\leq \sqrt{kd} \cdot \Delta_K^{1/(2d)} \cdot (\det M)^{1/(kd)}$ form a finite set. Now, we can consider all the maximal rank-1 submodules of M containing at least one of these vectors. By Condition 3 of Definition II.3.3, two maximal rank-1 submodules of M containing the same vector \mathbf{s}_i must be equal, hence there are only finitely many such submodules. This allows us to conclude that the infimum corresponding to $\lambda_1^{\mathcal{N}}(M)$ is over a finite set and must be reached. \square

A.1.5 Proof of Lemma II.1.12

The algorithm from Lemma II.1.12 is obtained by running the algorithm from Lemma II.1.11 until the output \mathbf{v} satisfies $\|\mathbf{v} - \mathbf{u}\| < \varsigma \cdot \sqrt{\ln(1/\varepsilon) + 4n}$. From Corollary II.1.10, this event happens with probability at least $1 - \varepsilon \geq 1/2$, hence the algorithm resamples at most twice on average, and the output distribution is within statistical distance $\leq \varepsilon$ from $D_{L, \varsigma, \mathbf{u}}$ (the distribution before rejection). Finally, note that $\sqrt{\ln(2n+4)/\pi} \leq \sqrt{n}$ for all $n \geq 1$, hence we can indeed apply Lemma II.1.11, and we conclude that the expected running time of the algorithm is polynomial. \square

Appendix B

Appendices of Chapter III

B.1 Analysis proofs

The equation $y \cdot e^y = x$ has a solution for any $x \in (-1/e, \infty)$. These solutions are given by the branches of the Lambert-W function (see Fig. B.1), which has two branches for $x \in (-1/e, 0)$, satisfying $W_{-1}(x) \in [2 \ln(-x), \ln(-x)]$ and $W_0(x) \in [-1, 0]$ for $x \in (-1/e, 0)$. For $x \geq 0$, there is only one solution $W(x) := W_0(x)$ satisfying $W(x) \in [\ln(1+x)/2, \ln(1+x)]$. We have that

$$\ln(-W_i(x)) = \ln(-x) - W_i(x). \quad (\text{B.1})$$

for any $x \in (-1/e, 0)$ and $i \in \{-1, 0\}$.

B.1.1 Proof of Lemma III.1.4

We bound the function

$$f(x) = \frac{(\ln \ln x)^d}{x^\alpha}$$




Figure B.1: Lambert-W function (figure borrowed from [Gee14]).

for $x \in (e, \infty)$. Note that f is not defined if $x < e$. Let $g(x) = d \ln \ln(x) - \alpha x$. The maximum of g is reached for $x_0 = \exp(W(d/\alpha))$. Let us bound $g(x_0)$ using the bounds on $W(x)$ for positive x .

$$\begin{aligned} g(x_0) &= d \ln \left(W \left(\frac{d}{\alpha} \right) \right) - \alpha \exp \left(W \left(\frac{d}{\alpha} \right) \right) \\ &\leq d \ln \left(\ln \left(1 + \frac{d}{\alpha} \right) \right). \end{aligned}$$

This last bound and the fact that $f = \exp \circ g \circ \ln$ allows us to conclude. \square

B.1.2 Comparisons between $\exp(\ln(x)^\alpha \ln \ln(x))$ and x^ε

Lemma B.1.1. *Let $d \geq 2$ be an integer and $\alpha(d) \in (0, 1/4)$ satisfying $\alpha(d) = \Theta(1/\ln(d))$. It holds that,*

$$\exp \left((\ln x)^{\alpha(d)} \cdot \ln \ln x \right) \leq \tilde{B}_d \cdot x^{\frac{\alpha(d)}{96d}}$$

for any $x > e$, and for some \tilde{B}_d satisfying

$$\ln \tilde{B}_d \ll \ln(d).$$

Proof. Let $\varepsilon(d) = \eta(d)/(96d)$. For the sake of readability, we will omit the dependence in d for $\eta(d)$ and $\varepsilon(d)$ and just write η and ε . To show the inequality of Lemma B.1.1, we bound the function

$$f(x) = \frac{\exp((\ln x)^\alpha \cdot \ln \ln(x))}{x^\varepsilon} = \exp(\ln(x)^\alpha \cdot \ln \ln(x) - \varepsilon \cdot \ln(x))$$

for $x \in (0, \infty)$. Let $g(x) = x^\alpha \cdot \ln(x) - \varepsilon \cdot x$. The derivative of this function is

$$g'(x) = x^{\alpha-1} + \alpha x^{\alpha-1} \cdot \ln(x) - \varepsilon,$$

whose zeros are (see Lemma B.1.2 below):

$$x_i = - \left(\frac{\alpha}{1-\alpha} \cdot \frac{1}{\varepsilon} \cdot \left(W_i \left(-\exp \left(-\frac{1-\alpha}{\alpha} \right) \cdot \varepsilon \cdot \frac{1-\alpha}{\alpha} \right) \right) \right)^{\frac{1}{1-\alpha}}$$

for $i \in \{-1, 0\}$. We have $x_0 < x_{-1}$. The function g tends to $-\infty$ when x tends to $+\infty$ and to 0 when x tends to 0, which implies that its maximum is reached at $x = x_{-1}$. Note that

$$x_{-1}^{\alpha-1} + x_{-1}^{\alpha-1} \cdot \alpha \cdot \ln(x_{-1}) = \varepsilon$$

is equivalent to

$$x_{-1}^{\alpha-1} \ln(x_{-1}) - \varepsilon = \frac{1}{\alpha} \cdot ((1-\alpha)\varepsilon - x_{-1}^{\alpha-1}).$$

For the sake of readability, we let $\alpha' = \alpha/(1-\alpha) = \Theta(1/\ln(d))$. Note that $\alpha = \alpha'/(1+\alpha')$ and that $1/(1-\alpha) = 1+\alpha'$. We bound $g(x_{-1})$ using the previous equality and the fact that $W_{-1}(x) \geq 2 \ln(-x)$ for $x \in (-1/e, 0)$:

$$\begin{aligned} g(x_{-1}) &= x_{-1} \cdot (x_{-1}^{\alpha-1} \ln x_{-1} - \varepsilon) = \frac{x_{-1}}{\alpha} \cdot ((1-\alpha)\varepsilon - x_{-1}^{\alpha-1}) \\ &\leq \frac{\varepsilon}{\alpha'} \cdot x_{-1} = -\frac{\varepsilon}{\alpha'} \cdot \left(\frac{\alpha'}{\varepsilon} \cdot W_{-1} \left(-\exp \left(-\frac{1}{\alpha'} \right) \cdot \frac{\varepsilon}{\alpha'} \right) \right)^{1+\alpha'} \\ &\leq 2^{1+\alpha'} \cdot \left(\frac{\alpha'}{\varepsilon} \right)^{\alpha'} \cdot \left(\frac{1}{\alpha'} + \ln \left(\frac{\alpha'}{\varepsilon} \right) \right)^{1+\alpha'} \\ &\ll d^{\alpha'} \cdot (\ln(d) + \ln(d))^{1+\alpha'} \\ &\ll \ln(d). \end{aligned}$$

This last bound and the fact that $f(x) = \exp \circ g \circ \ln$ allow us to conclude. \square

Lemma B.1.2. *Let $\alpha \in (0, 1/2)$ and $\varepsilon \in (0, 1)$. Let*

$$h(x) = x^{\alpha-1} + \alpha \frac{\ln(x)}{x^{1-\alpha}} - \varepsilon,$$

defined for $x \in (0, \infty)$. The function h is zero exactly at $x = x_i$ for $i \in \{-1, 0\}$ with

$$x_i = - \left(\frac{\alpha}{1-\alpha} \cdot \frac{1}{\varepsilon} \cdot \left(W_i \left(-\exp \left(-\frac{1-\alpha}{\alpha} \right) \cdot \varepsilon \cdot \frac{1-\alpha}{\alpha} \right) \right) \right)^{\frac{1}{1-\alpha}},$$

where W_{-1} and W_0 are the branches of the Lambert-W function. We have $x_0 \leq x_{-1}$.

Proof. Let $\alpha' = \alpha/(1-\alpha) \in (0, 1)$. By the change of variable $\tilde{x} = x^{\alpha-1}$, proving this result is equivalent to proving that the function

$$\tilde{h}(x) = x - \alpha' \cdot x \ln(x) - \varepsilon,$$

defined for any $x \in (0, \infty)$, has two zeros \tilde{x}_0 and \tilde{x}_{-1} equal to

$$\tilde{x}_i = -\frac{\varepsilon}{\alpha'} \cdot \left(W_i \left(-\exp \left(-\frac{1}{\alpha'} \right) \cdot \frac{\varepsilon}{\alpha'} \right) \right)^{-1}.$$

For the sake of readability, we define $y = -\varepsilon \cdot \exp(-1/\alpha')/\alpha'$. It can be checked that $y \in (-1/e, 0)$, so that $x_{-1}, x_0, \tilde{x}_{-1}$ and \tilde{x}_0 are well-defined. Also, for any $x \in (-1/e, 0)$ we have $W_{-1}(x) < W_0(x)$, which gives that $\tilde{x}_{-1} < \tilde{x}_0$, and hence $x_0 < x_{-1}$.

First, we prove that \tilde{h} has exactly two zeros. We have that

$$\tilde{h}'(x) = 1 - \alpha' - \alpha' \ln(x),$$

which is positive on $(0, \exp((1-\alpha')/\alpha'))$ and negative on $(\exp((1-\alpha')/\alpha'), \infty)$. The maximum of \tilde{h} is then equal to

$$\begin{aligned} \tilde{h} \left(\exp \left(\frac{1-\alpha'}{\alpha'} \right) \right) &= \alpha' \cdot \exp \left(\frac{1-\alpha'}{\alpha'} \right) - \varepsilon \\ &\geq \alpha' \left(1 + \frac{1-\alpha'}{\alpha'} \right) - \varepsilon \\ &= 1 - \varepsilon > 0. \end{aligned}$$

Now, since \tilde{h} tends to $-\varepsilon$ when x tends to 0 and to $-\infty$ when x tends to $+\infty$, the intermediate value theorem implies that \tilde{h} has exactly two zeros, one on $(0, \exp((1-\alpha')/\alpha'))$ and one on $(\exp((1-\alpha')/\alpha'), \infty)$. We now show that those zeros are \tilde{x}_0 and \tilde{x}_{-1} . Note that $\ln(-y) = \ln(\varepsilon/\alpha') - 1/\alpha'$. For $i \in \{-1, 0\}$, we have that

$$\begin{aligned} \tilde{h}(\tilde{x}_i) &= -\varepsilon - \frac{\varepsilon}{\alpha' \cdot W_i(y)} \cdot \left(1 - \alpha' \ln \left(\frac{\varepsilon}{\alpha'} \right) + \alpha' \ln(-W_i(y)) \right) \\ &= -\varepsilon - \frac{\varepsilon}{\alpha' \cdot W_i(y)} \cdot \left(1 - \alpha' \ln \left(\frac{\varepsilon}{\alpha'} \right) + \alpha' \ln \left(\frac{\varepsilon}{\alpha'} \right) - 1 - \alpha' W_i(y) \right) \\ &= -\varepsilon - \frac{\varepsilon}{\alpha' \cdot W_i(y)} \cdot (-\alpha' W_i(y)) = 0, \end{aligned}$$

where the second equality holds thanks to Eq. (B.1). This completes the proof. \square

B.2 Proof of Theorem III.1.8

Classical properties of the complex logarithm give that $|\zeta_K(s)| \leq \exp(|\ln \zeta_K(s)|)$. Note that if σ satisfies the assumption of Theorem III.1.8, then we have $2(1 - \sigma) \leq 4\eta = 1/(4 \ln(d))$. The two previous facts, along with Theorem B.3.1 and Lemma B.1.1, imply that for any $s = \sigma + it$ with σ and t satisfying the conditions of Theorem III.1.8, we have

$$\begin{aligned} |\zeta_K(s)| &\leq \exp\left(6d \cdot \ln(c_K|t|)^{2(1-\sigma)} \cdot \ln \ln(c_K|t|)\right) \\ &\leq \exp\left(6d \cdot \ln(c_K|t|)^{4\eta} \cdot \ln \ln(c_K|t|)\right) \\ &\leq \left(\tilde{B}_d \cdot (c_K|t|)^{\frac{\eta}{24d}}\right)^{6d} \\ &= \tilde{B}_d^{6d} \cdot c_K^{\eta/4} \cdot |t|^{\eta/4}. \end{aligned}$$

We set $B(K) = \tilde{B}_d^{6d} \cdot c_K^{\eta/4}$. The value of c_K and the bound on $\ln \tilde{B}_d$ allow us to conclude. \square

B.2.1 Proof of Lemma III.1.6

The class number formula gives that

$$\rho_K = \frac{2^{d_{\mathbb{R}}} \cdot (2\pi)^{d_{\mathbb{C}}} \cdot R_K \cdot |\text{Cl}_K|}{|\mu_K| \cdot \sqrt{|\Delta_K|}}$$

where R_K is the regulator, μ_K is the set of roots of unity of K and Cl_K the class group of K . By [Fri89, Theorem B] we have that $R_K/|\mu_K| \geq 0.08$. It also holds that $2^{d_{\mathbb{R}}} \cdot (2\pi)^{d_{\mathbb{C}}} \geq 2^d \geq 8$ and $|\text{Cl}_K| \geq 1$, hence the claimed result. \square

B.3 Bounds for $|\ln \zeta_K|$

Theorem B.3.1 (Assuming ERH). *Let $s = \sigma + it$. Assume that $|t| \geq T_0$ and $\sigma \in I_t$. Then*

$$|\ln \zeta_K(s)| \leq 6d \cdot (\ln(c_K|t|))^{2(1-\sigma)} \cdot \ln \ln(c_K|t|)$$

where $\ln \zeta_K(s)$ denotes the complex logarithm of $\zeta_K(s)$ ¹.

Proof. This theorem is a subset of [Sim23, Corollary 3], with modified constants. We now prove that our choice of constants is sound. It is stated in [Sim23, Corollary 3] that for $T'_0 = 9650 + 10^3 \log \log(c_K)$, $s = \sigma + it$ satisfying $|t| \geq T'_0$ and $\sigma \in I_t$, assuming the ERH (their condition is more precise, but implied by the ERH), it holds that

$$|\ln \zeta_K(s)| \leq 5.44d \cdot (b_1 \cdot \ln(c_K|t|))^{2(1-\sigma)} \cdot \ln \ln(c_K|t|),$$

where $b_1 \in (0.949, 0.95)$. Since $T_0 \geq T'_0$, the theorem is valid for $|t| \geq T_0$. One can check that $5.44 \cdot b_1^{2(1-\sigma)} \leq 6$ for any $\sigma \in I_t$. \square

¹It is well-defined for $\sigma \in I_t$, by the ERH.

Appendix C

Appendices of Chapter IV

C.1 Proof of Lemma IV.2.3

The `SampleWithTrap` algorithm is given below, as Algorithm C.1.1. It relies on an ideal-factoring oracle which can be implemented either in quantum polynomial time or in classical sub-exponential time. We prove the following statement, which can be viewed as a reformulation of Lemma IV.2.3. (Recall that factoring ideals reduces in polynomial time to factoring integers.)

Algorithm C.1.1 `SampleWithTrap`_{A,B}

Input: Integers $2 \leq A \leq B$, a real $\delta \in (0, 1]$ and a basis \mathbf{B}_I of a non-zero ideal I .

Oracle: \mathcal{F} for factoring integral ideals.

Output: (\mathfrak{p}, w) with $\mathfrak{p} \in \mathcal{P}_{A,B}$, and $w \in I\mathfrak{p}$.

- 1: Set $\varepsilon = \delta/(8B)$.
 - 2: Set $M = \sqrt{4 + \ln(3/\varepsilon)}/d$.
 - 3: Set $\varsigma = \max(\sqrt{d} \cdot \|\mathbf{B}_I^*\|, \Delta_K^{1/d} \cdot B^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \sqrt{\ln(3/\varepsilon)})$.
 - 4: Set $\mathbf{u} = M\varsigma \cdot \mathbf{1}$ with $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^d$.
 - 5: Set $k_{\max} = d \cdot \log_A(2M \cdot \sqrt{d} \cdot \mathcal{N}(I)^{-1/d})$.

 - 6: **repeat**
 - 7: Sample $w \leftarrow \tilde{D}_{\mathbf{B}_I, \varsigma, \mathbf{u}}$ using Lemma II.1.12 with error bound ε .
 - 8: Compute $\mathfrak{a} = I^{-1} \cdot (w)$.
 - 9: Factor \mathfrak{a} using \mathcal{F} and let \mathcal{S} be the set of distinct factors of \mathfrak{a} in $\mathcal{P}_{A,B}$.
 - 10: **until** $\mathcal{S} \neq \emptyset$.
 - 11: Sample \mathfrak{p} uniformly in \mathcal{S} .
 - 12: With probability $1 - \frac{|\mathcal{S}| \cdot \mathcal{N}(\mathfrak{p})}{k_{\max} \cdot B}$, go to Step 6.

 - 13: Return (\mathfrak{p}, w)
-

Lemma C.1.1. *Let \mathcal{F} be an ideal-factoring oracle. Given as inputs two integers $2 \leq A < B$, a real $\delta \in (0, 1]$ and the basis \mathbf{B}_I of a non-zero ideal I , `SampleWithTrap` outputs (\mathfrak{p}, w) such that*

- *the distribution of \mathfrak{p} is at statistical distance δ from the uniform distribution on $\mathcal{P}_{A,B}$;*

- the element w belongs to $I \cdot \mathfrak{p} \setminus \{0\}$ and satisfies $\|w\| \leq 2s \cdot \sqrt{4d + \ln(24B/\delta)}$, where

$$\varsigma = \max\left(\sqrt{d} \cdot \|\mathbf{B}_I^*\|, \Delta_K^{1/d} \cdot B^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \sqrt{\ln(24B/\delta)}\right).$$

Furthermore, `SampleWithTrap` runs in expected time polynomial in $B/|\mathcal{P}_{A,B}|$, B/A , $\log \Delta_K$, $\log B$, $\log(1/\delta)$ and in the size of its input.

Proof. We first analyze the running time of `SampleWithTrap` and then its correctness.

Running time. Observe that every step of the algorithm can be performed in polynomial time. For Step 7, we use Lemma II.1.12, whose assumptions are indeed satisfied. We further observe that at Step 12, the rejection probability is always between 0 and 1, hence we can indeed reject with this probability. Note that we have $B \geq \mathcal{N}(\mathfrak{p})$ since $\mathfrak{p} \in \mathcal{P}_{A,B}$. Also, we have $|\mathcal{S}| \leq \log_A \mathcal{N}(\mathfrak{a})$. It hence suffices to show that for any non-zero ideal \mathfrak{a} computed at Step 8, we have $\log_A \mathcal{N}(\mathfrak{a}) \leq k_{\max}$. From Lemma II.1.12, we know that $\|w - \mathbf{u}\| < \varsigma \cdot \sqrt{\ln(3/\varepsilon) + 4d}$. As $\|\mathbf{u}\| = \sqrt{d} \cdot M \cdot \varsigma = \varsigma \cdot \sqrt{\ln(3/\varepsilon) + 4d}$, we have $\|w\| \leq 2\|\mathbf{u}\| = 2M \cdot \sqrt{d} \cdot \varsigma$, which in turn implies that $\mathcal{N}(w) \leq \|w\|^d \leq (2M \cdot \sqrt{d} \cdot \varsigma)^d$. Hence, we conclude that $\mathcal{N}(\mathfrak{a}) \leq (2M \cdot \sqrt{d} \cdot \varsigma)^d \cdot \mathcal{N}(I)^{-1} = A^{k_{\max}}$. This shows that $(|\mathcal{S}| \cdot \mathcal{N}(\mathfrak{p})) / (k_{\max} \cdot B)$ belongs to $[0, 1]$, as desired.

We now study the probability of exiting the outer loop, from Step 6 to Step 12. It is bounded from below by $A/(k_{\max}B)$ (since we have $|\mathcal{S}| \geq 1$ when we exit the inner loop). Hence, the expected number of iterations of this loop is at most $k_{\max} \cdot B/A$. Since $A \geq 2$, then k_{\max} is polynomial in $d = \text{poly}(\log \Delta_K, \log(\varsigma), \log(M)$ and $\log \mathcal{N}(I^{-1})$). From the definition of ς , one can check that k_{\max} is polynomial in $\log \Delta_K$, $\log B$, $\log \log(1/\delta)$ and the size of the input.

It remains to bound from below the probability of exiting the inner loop, from Step 6 to Step 10. The proof of this statement is an adaptation of the proof of [Gen09, Lemma 15.2.3]. This probability can be written as:

$$\Pr_{w \leftarrow \tilde{D}_{\mathbf{B}_I, \varsigma, \mathbf{u}}} (\exists \mathfrak{p} \in \mathcal{P}_{A,B} : \mathfrak{p} \text{ divides } I^{-1} \cdot (w)) = \sum_{w \in I} \mathbf{1}_W(w) \cdot \tilde{D}_{\mathbf{B}_I, \varsigma, \mathbf{u}}(w) \quad (\text{C.1})$$

where $W = \cup_{\mathfrak{p} \in \mathcal{P}_{A,B}} I \cdot \mathfrak{p}$ and $\mathbf{1}_W(\cdot)$ is the indicator function of W . For any $w \in I \setminus \{0\}$, we have

$$\mathbf{1}_W(w) \geq \frac{1}{\ln(\mathcal{N}(w \cdot I^{-1}))} \cdot \sum_{\substack{\mathfrak{p} \in \mathcal{P}_{A,B} \\ \mathfrak{p} | w \cdot I^{-1}}} \ln(\mathcal{N}(\mathfrak{p})).$$

Indeed, either $w \notin W$ and the sum on the right is empty, or $w \in W$ and the sum on the right is bounded from above by 1 (since the norm of the product of all the primes dividing $w \cdot I^{-1}$ is at most the norm of $w \cdot I^{-1}$ when $w \cdot I^{-1}$ is non-zero). Moreover, we have already seen that the algebraic norm of $\mathfrak{a} = w \cdot I^{-1}$ is at most $(2M \cdot \sqrt{d} \cdot \varsigma)^d \cdot \mathcal{N}(I^{-1})$, and by assumption we know that $\mathcal{N}(\mathfrak{p}) \geq A$ for all $\mathfrak{p} \in \mathcal{P}_{A,B}$. Hence, letting $\mathbf{1}_{I \cdot \mathfrak{p}}(\cdot)$ be the indicator function of $I \cdot \mathfrak{p}$, it holds that

$$\begin{aligned} \mathbf{1}_W(w) &\geq \frac{\ln(A)}{\ln((2M \cdot \sqrt{d} \cdot \varsigma)^d \cdot \mathcal{N}(I^{-1}))} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbf{1}_{I \cdot \mathfrak{p}}(w) \\ &= \frac{1}{k_{\max}} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbf{1}_{I \cdot \mathfrak{p}}(w), \end{aligned}$$

where k_{\max} is defined as in Step 5.

Before returning to (C.1), note that $\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(0) = 0$. Indeed, we have seen that $\|w - \mathbf{u}\| < M \cdot \sqrt{d} \cdot \varsigma$ and, by construction, we have $\|\mathbf{u}\| = M \cdot \sqrt{d} \cdot \varsigma$. Using the above, we obtain:

$$\begin{aligned} \sum_{w \in I} \mathbf{1}_W(w) \cdot \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(w) &= \sum_{w \in I \setminus \{0\}} \mathbf{1}_W(w) \cdot \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(w) \\ &\geq \frac{1}{k_{\max}} \sum_{w \in I \setminus \{0\}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbf{1}_{I \cdot \mathfrak{p}}(w) \cdot \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(w) \\ &= \frac{1}{k_{\max}} \sum_{w \in I} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbf{1}_{I \cdot \mathfrak{p}}(w) \cdot \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(w) \\ &= \frac{1}{k_{\max}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p}). \end{aligned}$$

From Lemma II.1.12, we know that $\text{SD}(\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}, D_{I,\varsigma,\mathbf{u}}) \leq \varepsilon$. Hence, it holds that $\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p}) \geq D_{I,\varsigma,\mathbf{u}}(I \cdot \mathfrak{p}) - \varepsilon$. Moreover, observe that by definition of ς , it holds that for any $\mathfrak{p} \in \mathcal{P}_{A,B}$ we have $\varsigma \geq \mathcal{N}(I \cdot \mathfrak{p})^{1/d} \cdot \Delta_K^{1/d} \cdot \sqrt{\ln(3/\varepsilon)}$. Hence, we can apply Corollary II.2.6 and we obtain that $D_{I,\varsigma,\mathbf{u}}(\mathfrak{p} \cdot I) \geq (1 - \varepsilon) \cdot \mathcal{N}(\mathfrak{p})^{-1} \geq 1/(2B)$. By choice of ε , we finally obtain

$$\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p}) \geq \frac{1}{2B} - \varepsilon \geq \frac{1}{4B}.$$

Plugging this back in our lower bound on the probability to exit the inner loop, we have

$$\Pr_{w \leftarrow \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}} (\exists \mathfrak{p} \in \mathcal{P}_{A,B} : \mathfrak{p} \text{ divides } I^{-1} \cdot (w)) \geq \frac{1}{k_{\max}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \frac{1}{4B} = \frac{|\mathcal{P}_{A,B}|}{4B \cdot k_{\max}}.$$

The expected number of iterations of the inner loop is then $\leq 4k_{\max} \cdot B/|\mathcal{P}_{A,B}|$.

Correctness. Let (\mathfrak{p}, w) be the output of `SampleWithTrap` on input $(A, B, \delta, \mathbf{B}_I)$. By construction, we have $\mathfrak{p} \in \mathcal{P}_{A,B}$. Further, as $w \in I$ and $\mathfrak{p} | I^{-1} \cdot (w)$, we have that $w \in I \cdot \mathfrak{p}$. The bound on $\|w\|$ comes from the fact that $\|w\| \leq 2\|\mathbf{u}\| = 2M \cdot \sqrt{d} \cdot \varsigma$. It remains to prove that the distribution \mathcal{D} of the ideal \mathfrak{p} is within statistical distance $\leq \delta/2$ from uniform over $\mathcal{P}_{A,B}$.

Let us fix $\mathfrak{p} \in \mathcal{P}_{A,B}$ and compute $\mathcal{D}(\mathfrak{p})$. First, we compute the probability that \mathfrak{p} is chosen at Step 11 of the algorithm. The distribution of the element w when exiting of the inner loop is $\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}$ conditioned on $w \in W = \cup_{\mathfrak{q} \in \mathcal{P}_{A,B}} I \cdot \mathfrak{q}$ (which is equivalent to $S \neq \emptyset$). Moreover, the ideal \mathfrak{p} belongs to S if and only if $w \in I \cdot \mathfrak{q}$. So the probability that \mathfrak{p} belongs to S in Step 11 is

$$\Pr(\mathfrak{p} \in S \text{ in Step 11}) = \frac{\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p})}{\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(W)}.$$

Note that the quantity $\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(W)$ is a fixed and independent of \mathfrak{p} (and non-zero, since the algorithm terminates). In the rest of the computation, we will write it p_0 . After running Step 11, we obtain

$$\Pr(\mathfrak{p} \text{ is chosen in Step 11}) = \frac{\tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p})}{|\mathcal{S}| \cdot p_0}.$$

By Lemma II.1.12, Corollary II.2.6 and the choice of ς , we know that

$$\begin{aligned} \tilde{D}_{\mathbf{B},\varsigma,\mathbf{u}}(I \cdot \mathfrak{p}) &\in [1 - \varepsilon, 1 + \varepsilon] \cdot \mathcal{N}(\mathfrak{p})^{-1} + [-\varepsilon, \varepsilon] \\ &\subseteq [1 - \delta/4, 1 + \delta/4] \cdot \mathcal{N}(\mathfrak{p})^{-1}, \end{aligned}$$

where in the last inequality we used the fact that $\varepsilon = \delta/(8 \cdot B) \leq \delta/8 \cdot \mathcal{N}(\mathbf{p})^{-1}$. Combining this with the equation above, we obtain that

$$\Pr(\mathbf{p} \text{ is chosen in Step 11}) \in \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{\mathcal{N}(\mathbf{p}) \cdot |\mathcal{S}| \cdot p_0}.$$

Finally, because of the rejection sampling in Step 12, we have

$$\begin{aligned} \Pr(\mathbf{p} \text{ is selected after Step 12}) &\in \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{\mathcal{N}(\mathbf{p}) \cdot |\mathcal{S}| \cdot p_0} \cdot \frac{|\mathcal{S}| \cdot \mathcal{N}(\mathbf{p})}{k_{\max} \cdot B} \cdot \frac{1}{p'_0} \\ &= \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{k_{\max} \cdot B \cdot p_0 \cdot p'_0}, \end{aligned}$$

where p'_0 is the probability (over the random choice of w , the random choice of \mathbf{p} and the rejection probability of Step 12) that one exists the outer loop.

Overall, we have just proven that there exists some quantity C such that for any $\mathbf{p} \in \mathcal{P}_{A,B}$, it holds that $\mathcal{D}(\mathbf{p}) \in [1 - \delta/4, 1 + \delta/4] \cdot C$. Since $\sum_{\mathbf{p} \in \mathcal{P}_{A,B}} \mathcal{D}(\mathbf{p}) = 1$, it must be that $C \in \left[\frac{1}{1+\delta/4}, \frac{1}{1-\delta/4}\right] \cdot \frac{1}{|\mathcal{P}_{A,B}|}$. It implies that for all $\mathbf{p} \in \mathcal{P}_{A,B}$,

$$\left| \mathcal{D}(\mathbf{p}) - \frac{1}{|\mathcal{P}_{A,B}|} \right| \leq \max\left(1 - \frac{1-\delta/4}{1+\delta/4}, \frac{1+\delta/4}{1-\delta/4} - 1\right) \cdot \frac{1}{|\mathcal{P}_{A,B}|} \leq \frac{\delta}{|\mathcal{P}_{A,B}|}.$$

The statistical distance between \mathcal{D} and the uniform distribution satisfies

$$\begin{aligned} \text{SD}(\mathcal{D}, \mathcal{U}(\mathcal{P}_{A,B})) &= \frac{1}{2} \cdot \sum_{\mathbf{p} \in \mathcal{P}_{A,B}} \left| \mathcal{D}(\mathbf{p}) - \frac{1}{|\mathcal{P}_{A,B}|} \right| \\ &\leq \frac{\delta}{2} \cdot \sum_{\mathbf{p} \in \mathcal{P}_{A,B}} \frac{1}{|\mathcal{P}_{A,B}|} = \frac{\delta}{2}. \end{aligned}$$

This completes the proof. □

C.2 Proof of Theorem IV.2.4

In this section, we provide a proof of Gentry's reduction for SVP, as stated in Theorem IV.2.4. The proof is similar to the one provided in Gentry's thesis [Gen09], but we instantiate it directly with the shortest vector problem, instead of the variant of the bounded distance decoding problem used in [Gen09].

C.2.1 Balanced-ideal-HSVP

In the proof, we will make use of balanced elements (as defined in Definition IV.2.1). We introduce the problem ideal-balanced-HSVP, and give a (folklore) proof that this problem is equivalent to id-HSVP (up to a polynomial loss in the approximation factor). The balanced version of id-HSVP will be more convenient to use in the following proof.

Definition C.2.1. *Let $\eta > 1$ and $\gamma \geq 1$. The problem $\text{id-BHSVP}_\gamma^\eta$ asks, given as input a fractional ideal I , to find a non-zero element $x \in I$ such that $\|x\| \leq \gamma \cdot \text{Vol}(I)^{1/d}$ and x is η -balanced. The problem $\text{inv-BHSVP}_\gamma^\eta$ is the problem $\text{id-BHSVP}_\gamma^\eta$ restricted to inverses of integral lattices.*

Algorithm C.2.1 BalanceElement

Input: The HNF of a fractional ideal I , an element $x \in I$ and $M > 0$.

Output: $y \in I$.

- 1: Let $\varsigma = \sqrt{d} \cdot \delta_K \cdot \|x\|_\infty$.
- 2: Let $\mathbf{B}_I = \text{ReduceIdeal}(I, x)$.
- 3: Let $\mathbf{t} = \varsigma \sqrt{d}(M+1)/2 \cdot \mathbf{1}$ with $\mathbf{1} = (1, \dots, 1) \in K_{\mathbb{R}}$.
- 4: Run Babai's nearest plane algorithm on $(\mathbf{B}_I, \mathbf{t})$; let $y \in I$ be the output.
- 5: Return y .

We describe in Algorithm C.2.1 a polynomial-time reduction from id-BHSVP to id-HSVP, which relies on Babai's nearest plane algorithm [Bab86].

Lemma C.2.2. *Algorithm C.2.1 runs in polynomial time. On input I, x, M with $x \in I \setminus \{0\}$, it outputs $y \in I \setminus \{0\}$ that is $(1 + 2/M)$ -balanced and satisfies*

$$\|y\| \leq (1 + M/2) \cdot d^{3/2} \cdot \delta_K \cdot \|x\|_\infty.$$

Proof. The running time follows directly from the description of the algorithm. Let y be the output of Algorithm C.2.1 on input I, x and M . We have, by property of the nearest plane algorithm (see [Bab86, Theorem 3.1]), that there exist μ_1, \dots, μ_d in $[-1/2, 1/2]$ such that

$$\|y - \mathbf{t}\|_\infty \leq \|y - \mathbf{t}\| \leq \left(\sum_i \mu_i^2 \cdot \|\mathbf{b}_i^*\|^2 \right)^{1/2}.$$

By Lemma II.2.15, we have $\|\mathbf{b}_i^*\| \leq \delta_K \cdot \|x\| \leq \varsigma$. We hence obtain that $\|y - \mathbf{t}\|_\infty \leq \sqrt{d}\varsigma/2$. As a result, we have $|y_i| \in [\varsigma\sqrt{d}M/2, \varsigma\sqrt{d}(M+2)/2]$ for all i .

We then have $\mathcal{N}^{1/d}(y) \geq \varsigma\sqrt{d}M/2 \geq M/(M+2)\|y\|_\infty$. The same holds for y^{-1} , which gives that y is $(1 + 2/M)$ -balanced. Finally, the inequality $\|y\| \leq \sqrt{d} \cdot \|y\|_\infty$ gives the desired bound on the norm of y . \square

Corollary C.2.3. *For any $\gamma \geq 1$ and $\eta > 1$, there is a Karp polynomial time reduction from id-BHSVP $_{\gamma'}$ to id-HSVP $_{\gamma}$, where $\gamma' = \gamma \cdot \delta_K \cdot d^{3/2} \cdot \eta/(\eta - 1)$.*

Note that the converse reduction holds without any parameter loss, by definition of id-BHSVP.

Proof. Let I be an instance of id-BHSVP $_{\gamma'}$, and assume we have an oracle \mathcal{O} for id-HSVP $_{\gamma}$. Let x be the output of \mathcal{O} on I . We let $M = 2/(\eta - 1)$ and return $y = \text{BalanceElement}(I, x, M)$. The fact that y is a valid id-BHSVP $_{\gamma'}$ solution follows from the definition of M and Lemma C.2.2. \square

Corollary C.2.4. *For any constant $\eta > 1$, id-BHSVP $_{\gamma_{\text{easy}}(\eta)}$ can be solved in polynomial time for $\gamma_{\text{easy}}(\eta) = \delta_K \cdot d^{3/2} \cdot (\eta/(\eta - 1)) \cdot 2^d$.*

Proof. The result follows from Corollary C.2.3, by using the LLL algorithm to solve id-HSVP $_{\gamma}$. \square

C.2.2 Finding a non-trivial solution to inv-HSVP using a $\mathcal{P}_{A,B}^{-1}$ -avg-HSVP oracle

The reduction from Theorem IV.2.4 is an iterative reduction, which proceeds by iteratively improving an existing solution with the usage of an oracle solving $\mathcal{P}_{A,B}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$. In this subsection, we focus on the main ingredient of one iteration of the reduction, the `SampleSmall` algorithm, presented in Algorithm C.2.2. The objective of this algorithm is, given as input \mathfrak{b}^{-1} the inverse of a prime ideal, to find a non-trivial short non-zero vector in \mathfrak{b}^{-1} . Indeed, since \mathfrak{b} is integral, we know that $\mathcal{O}_K \subseteq \mathfrak{b}^{-1}$, so the short non-zero vectors of \mathcal{O}_K give trivial solutions to short non-zero vectors in \mathfrak{b}^{-1} . The objective of the `SampleSmall` algorithm is to find slightly shorter vectors than those trivial short vectors lying in \mathcal{O}_K (which will exist if the norm of \mathfrak{b} is large enough). In particular, we would like to obtain $x \in \mathfrak{b}^{-1} \setminus \{0\}$ with $\|x\| < 1$, so that multiplying by x decrease the euclidean norm. This will be used in the reduction to iteratively decrease the norm of a short non-zero vector found in our input ideal I .

Algorithm C.2.2 `SampleSmall` $_{A,B}$

Input: A basis of an integral ideal \mathfrak{b} .

Oracles: \mathcal{O} for $\mathcal{P}_{A,B}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$, \mathcal{F} for factoring integral ideals.

Output: $x \in \mathfrak{b}^{-1}$ or $x = \perp$.

- 1: Compute a basis \mathbf{B} of \mathfrak{b}^{-1} with $\|\mathbf{B}^*\| \leq \delta_K$ (using `InvertIdeal`).
 - 2: Set (\mathfrak{p}, w) be the output of `SampleWithTrap` $_{A,B}$ on input $(A, B, 2^{-(d+1)}, \mathbf{B})$ (this relies on \mathcal{F}).
 - 3: Set $v = \mathcal{O}(\mathfrak{p}^{-1})$.
 - 4: If $v \neq \perp$, then return $v \cdot w$.
 - 5: Else, return \perp .
-

Theorem C.2.5. *Let $\gamma_{\text{avg}} \geq 1$ and $3 \leq A < B$ satisfying $B/|\mathcal{P}_{A,B}|, B/A \leq \text{poly}(\log \Delta_K)$. Let \mathcal{O} be an oracle solving $\mathcal{P}_{A,B}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ with success probability $\delta \geq 2^{-d}$ and let \mathcal{F} be an ideal-factoring oracle.*

On input a non-zero integral ideal \mathfrak{b} and given access to \mathcal{O} and \mathcal{F} , Algorithm `SampleSmall` $_{A,B}$ runs in expected time $\text{poly}(\log \Delta_K, \log B, \log \mathcal{N}(\mathfrak{b}))$, and performs only one call to \mathcal{O} and possibly multiple calls to \mathcal{F} for integral ideals of norm $\text{poly}(\log \Delta_K, \log B, \log \mathcal{N}(\mathfrak{b}))$ bits. It outputs $x \neq \perp$ with probability $\geq \delta/2$ and, when this is the case, it holds that $x \in \mathfrak{b}^{-1} \setminus \{0\}$ and

$$\|x\| \leq \frac{10\gamma_{\text{avg}}(d + \ln B) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, (B \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1}))^{1/d}\right).$$

Proof. We first focus on the running time of the algorithm. Every step can be performed in polynomial time. For Step 1, we use Lemma II.2.15 and the fact that \mathfrak{b} is integral. For Step 2, we use Lemma IV.2.3. Note that Step 3 is not inside a loop, hence the call to \mathcal{O} is performed only once.

Let us now prove that the algorithm returns an element $x \neq \perp$ with probability at least $\delta/2$. Note that by Lemma IV.2.3, the distribution \mathcal{D} of the ideal \mathfrak{p} given as input to \mathcal{O} is within statistical distance $\leq 2^{-(d+1)} \leq \delta/2$ from uniform over $\mathcal{P}_{A,B}$ (here we used the lower bound $\delta \geq 2^{-d}$). Since we know that \mathcal{O} has success probability δ when its input \mathfrak{p}^{-1} is distributed uniformly in $\mathcal{P}_{A,B}^{-1}$, this proves that the probability that \mathcal{O} succeeds in solving id-HSVP $_{\gamma_{\text{avg}}}$ in Step 3 of the algorithm is at least $\delta - \delta/2 \geq \delta/2$, as desired.

Finally, let us prove the upper bound on $\|x\|$ when $x \neq \perp$. In this case, we have $x = v \cdot w$ and use the upper bounds on v and on w (from Lemma IV.2.3) to obtain

$$\begin{aligned} \|x\| &\leq \|v\| \cdot \|w\| \\ &\leq \gamma_{\text{avg}} \cdot \text{Vol}(\mathfrak{b}^{-1})^{1/d} \cdot 2(5d + \ln B + \ln(48)) \cdot \max\left(\delta_K, (\Delta_K \cdot B \cdot \mathcal{N}(\mathfrak{b}^{-1}))^{1/d}\right) \\ &\leq \frac{10\gamma_{\text{avg}}(d + \ln B) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, (B \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1}))^{1/d}\right), \end{aligned}$$

where we used the fact that $B \geq 3$. This completes the proof. \square

For simplicity, we will use the following corollary, where we use the Extended Riemann Hypothesis in order to estimate the number of prime ideals in the set $\mathcal{P}_{A,B}$ and simplify the conditions.

Corollary C.2.6 (Assuming ERH). *Let $\gamma_{\text{avg}} \geq 1$ and $3 \leq A \leq (\Delta_K)^{d^{O(1)}}$. Let \mathcal{O} be an oracle solving $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ with success probability $\delta \in (0, 1]$ and let \mathcal{F} be an oracle factoring integral ideals. Let $\varepsilon \in (0, 1)$ and assume that*

$$A^{1/d} \geq 10 \cdot \gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}.$$

Then there exists an algorithm \mathcal{A} that takes as input any integral ideal \mathfrak{b} with $\mathcal{N}(\mathfrak{b}) \geq 4A$ and outputs $x \in \mathfrak{b}^{-1} \setminus \{0\}$ such that $\|x\| \leq \varepsilon$. If given access to \mathcal{O} and \mathcal{F} , algorithm \mathcal{A} runs in expected time $\text{poly}(\log \Delta_K, \log(\mathcal{N}(\mathfrak{b})), 1/\delta)$ and calls \mathcal{F} on ideals of norm at most $\text{poly}(\log \Delta_K, \log(\mathcal{N}(\mathfrak{b})))$ bits.

Proof. Algorithm \mathcal{A} consists in repeatedly running `SampleSmall` $_{A,B}$ with $B = 4A$, on input \mathfrak{b} , until it outputs $x \neq \perp$. Let us prove that A and B satisfy the constraints required in Theorem C.2.5. If $A \leq \text{poly}(\log \Delta_K)$, then $4A/|\mathcal{P}_{A,4A}|$ is polynomial. Else, the ERH implies that

$$\frac{4A}{|\mathcal{P}_{A,B}|} \leq O(\ln A) \leq \text{poly}(\log \Delta_K).$$

The claim on the running time of algorithm \mathcal{A} and the fact that $x \in \mathfrak{b}^{-1} \setminus 0$ follow from Theorem C.2.5. Note that Theorem C.2.5 is guaranteed to work only if the success probability δ of \mathcal{O} is at least 2^{-d} . If the success probability is smaller than this quantity, algorithm \mathcal{A} simply runs an SVP solver on ideal \mathfrak{b}^{-1} and returns a shortest non-zero vector. This shortest non-zero vector will have Euclidean norm $\leq \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot A^{1/d} \leq \varepsilon$ by assumption on A , and the call to the SVP solver has a running time $2^{O(d)} = \text{poly}(1/\delta)$.

We now bound $\|x\|$. From Theorem C.2.5 and by choice of B , we know that

$$\|x\| \leq \frac{10\gamma_{\text{avg}}(d + \ln(4A)) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, (4A \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1}))^{1/d}\right).$$

Since $\mathcal{N}(\mathfrak{b}) \geq 4A$ and $\delta_K \geq \lambda_d(\mathcal{O}_K) \geq \Delta_K^{1/(2d)}$, it holds that $(4A \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1}))^{1/d} \leq \delta_K \cdot \Delta_K^{1/(2d)}$, and hence

$$\|x\| \leq \frac{10\gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K}{A^{1/d}} \leq \varepsilon.$$

The last inequality follows from the assumption on A and ε . \square

C.2.3 Iterating the reduction

In order to prove Theorem IV.2.4, we are going to use the id-BHSVP problem. Recall that the id-BHSVP problem is equivalent to the id-HSVP problem, up to some polynomial loss, so we can safely replace id-HSVP by id-BHSVP, which will make our reductions easier to prove. The lemma below states that if we have an oracle solving $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ and an algorithm solving inv-BHSVP $_{\gamma'}^{\eta}$, then we can create an algorithm solving inv-BHSVP $_{\gamma'}^{\eta'}$ where γ' is slightly smaller than γ and η' is slightly larger than η (i.e., we can find smaller less balanced vectors in our ideals). This corresponds to one iteration of the full reduction.

For the whole subsection, we fix $\gamma_{\text{avg}} \geq 1$, $\varepsilon \in (0, 1)$ and $3 \leq A \leq (\Delta_K)^{d^{O(1)}}$ satisfying:

$$A^{1/d} \geq 10 \cdot \gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}.$$

Lemma C.2.7 (Assuming ERH). *Let $\gamma_{\min} = (4A)^{1/d}/(\varepsilon \cdot \Delta_K^{1/(2d)})$, $\gamma > \gamma_{\min}$ and $\eta \in (1, \gamma/\gamma_{\min}]$.*

$$\text{inv-BHSVP}_{\gamma'}^{\eta'} \text{ reduces to inv-BHSVP}_{\gamma}^{\eta} \text{ and } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}$$

for $\eta' = \eta \cdot (1 + 1/d)$ and $\gamma' = 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon \cdot \gamma$. If given access to an oracle \mathcal{F} factoring integral ideals, the expected running time of the reduction is polynomial in $\log \Delta_K$, $\log \gamma$, $1/\delta$ and the size of its input, where δ is the success probability of the oracle for $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$. Moreover, the oracle \mathcal{F} is called on ideals whose norms have a bit-size $\text{poly}(\log \Delta_K, \log \gamma)$.

Proof. Assume that we are given $I = \mathfrak{b}^{-1}$ the inverse of an integral ideal. Let x be the output of the inv-BHSVP $_{\gamma}^{\eta}$ oracle on input I . As $\eta' \geq \eta$, the element x is η' -balanced. If $\|x\|_{\infty} \leq \varepsilon \cdot \gamma \cdot \text{Vol}(I)^{1/d}$, then it is a solution for inv-BHSVP $_{\sqrt{d} \cdot \varepsilon \cdot \gamma}^{\eta'}$ and we can output it. Else, we have

$$|\mathcal{N}(x)| \geq \eta^{-d} \cdot \|x\|_{\infty}^d \geq \eta^{-d} \cdot \varepsilon^d \cdot \gamma^d \cdot \Delta_K^{1/2} \cdot \mathcal{N}(I).$$

Now we set $\mathfrak{b} = (x) \cdot I^{-1}$. This ideal is the inverse of an integral ideal, and by the previous inequality and the condition on η we have

$$\mathcal{N}(\mathfrak{b}) = \frac{\mathcal{N}(x)}{\mathcal{N}(I)} \geq \frac{\varepsilon^d \cdot \gamma^d \cdot \Delta_K^{1/2}}{\eta^d} \geq 4A.$$

This last inequality, and the definition of A meet the conditions of Corollary C.2.6, we then can make a call to $\text{SampleSmall}_{A,4A}(\mathfrak{b})$ and denote by y its output. The element y satisfies $\|y\|_{\infty} \leq \varepsilon$ and $y \in \mathfrak{b}^{-1} \setminus \{0\}$.

We now denote $y' = \text{BalanceElement}(\mathfrak{b}^{-1}, y, 2d)$. By Lemma C.2.2, we have that $y' \in \mathfrak{b}^{-1} \setminus 0$ is $(1 + 1/d)$ -balanced and that

$$\|y'\| \leq (1 + d) \cdot d^{3/2} \cdot \delta_K \cdot \varepsilon \leq 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon$$

We then return $y' \cdot x$. We have that $y' \cdot x \in I$, and since x is η -balanced and y' is $(1 + 1/d)$ -balanced, then xy' is η' -balanced and

$$\|x \cdot y'\| \leq \|y'\| \cdot \|x\| \leq 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon \cdot \gamma \cdot \text{Vol}(I)^{1/d} = \gamma' \cdot \text{Vol}(I)^{1/d}.$$

The running time of the algorithm comes from the running time of $\text{SampleSmall}_{A,4A}(\mathfrak{b})$ and the running time of $\text{BalanceElement}(\mathfrak{b}^{-1}, y, 2d)$. The former is polynomial in $\log \Delta_K$, $\log \mathcal{N}(\mathfrak{b})$ and $1/\delta$ and requires factoring ideals of norm at most $\text{poly}(\log \Delta_K, \log \mathcal{N}(\mathfrak{b}))$ bits. The latter has a running time polynomial in $\log \Delta_K$ and $\log \mathcal{N}(\mathfrak{b})$. Observe that $\mathcal{N}(\mathfrak{b}) = |\mathcal{N}(x)|/\mathcal{N}(I) \leq \|x\|^d/\mathcal{N}(I) \leq \gamma^d \cdot \sqrt{\Delta_K}$. The result follows. \square

We will now iterate Lemma C.2.7, instantiated with $\varepsilon = 1/2 \cdot (2 \cdot d^{5/2} \cdot \delta_K)^{-1}$. This choice of ε ensures that $\gamma' = \gamma/2$, i.e., the approximation factor is divided by 2 at every iteration of the reduction (at the cost of slightly less balanced elements). We will iterate this reduction step until we obtain a reduction from $\text{inv-BHSVP}_{\gamma'}^{\eta'}$ with an approximation factor γ' as small as possible, to $\text{inv-BHSVP}_{\gamma}^{\eta}$ with γ so large that it can be solved in polynomial time using the LLL algorithm. Hence, the only oracle that will remain for the reduction to work is the one solving $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ (and the one factoring ideals, which can be quantumly efficiently instantiated).

Lemma C.2.8. *Let $\gamma_{\text{avg}} \geq 1$, $3 \leq A \leq \Delta_K^{d^{O(1)}}$ satisfying*

$$A^{1/d} \geq \gamma_{\text{avg}} \cdot 40 \cdot d^{5/2} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K^2$$

and

$$\gamma_{\min} = \frac{4 \cdot d^{5/2} \cdot \delta_K \cdot (4A)^{1/d}}{\Delta_K^{1/(2d)}}.$$

There exists a reduction

$$\text{from } \text{inv-BHSVP}_{2e\gamma_{\min}}^{2e} \text{ to } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

Given access to an ideal-factoring oracle \mathcal{F} , the expected running time of this reduction is polynomial in its input bit-size, in $\log \Delta_K$ and in $1/\delta$, where $\delta \in (0, 1]$ is the success probability of the $\mathcal{P}_{A,4A}^{-1}$ -avg-id-HSVP $_{\gamma_{\text{avg}}}$ oracle. Moreover, the reduction calls \mathcal{F} on integral ideals whose algebraic norms have bit-size $\text{poly}(\log \Delta_K)$.

Proof. Let $\varepsilon = (4d^{5/2} \cdot \delta_K)^{-1}$. Define $\gamma_0 = \gamma_{\min} \cdot 2e \cdot 2^d$, $\eta_0 = 2$, and for any $k \in \{1, \dots, d\}$ $\gamma_k = \gamma_0 \cdot 2^{-k}$ and $\eta_k = \eta_0 \cdot (1 + 1/d)^k$. Observe that, for any k , we have that $\gamma_k > \gamma_{\min}$ and $\eta_k \in (1, \gamma_k/\gamma_{\min}]$. Moreover, if we let $\varepsilon = (4d^{5/2} \cdot \delta_K)^{-1}$, then our choice of γ_{\min} coincide with the definition of γ_{\min} in Lemma C.2.7, and our choice of A satisfies the constraint $A^{1/d} \geq 10 \cdot \gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}$.

We can then apply Lemma C.2.7 and we get, for any $0 \leq k < d$, that

$$\text{inv-BHSVP}_{\gamma_{k+1}}^{\eta_{k+1}} \leq \text{inv-BHSVP}_{\gamma_k}^{\eta_k} + \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

By combining the reduction, we then have that

$$\text{inv-BHSVP}_{\gamma_d}^{\eta_d} \leq \text{inv-BHSVP}_{\gamma_0}^{\eta_0} + \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

Now, from the definition of γ_{\min} and the lower bound on $A^{1/d}$, one can check that $\gamma_0 \geq \delta_K \cdot d^{3/2} \cdot \left(\frac{\eta_0}{\eta_0 - 1}\right) \cdot 2^d$. Hence, by Corollary C.2.4 we have that $\text{inv-BHSVP}_{\gamma_0}^{\eta_0}$ can be solved in polynomial time.

Regarding the running time, our reduction consists in d consecutive reductions. Lemma C.2.7 implies that the k -th reduction has a running time polynomial in $\log \Delta_K$, $\log \gamma_k$ and $1/\delta$. Since for every k we have that $\log \gamma_k \leq \log \gamma_0 = \text{poly}(\log \Delta_K)$, we conclude that the total running time of the reduction is polynomial in $\log \Delta_K$ and $1/\delta$. The same argument also shows that the ideal-factoring oracle is only called on integral ideals whose norm have a bit-size $\text{poly}(\Delta_K)$. \square

We are now ready to prove our main theorem of this section. To do so, we instantiate Lemma C.2.8 with an appropriate value of A , and combine the reduction with the ones from Appendix C.2.1 showing that inv-BHSVP is equivalent to id-HSVP (up to polynomial losses).

of Theorem IV.2.4. Let $C_{1,K}$ be minimal such that $C_{1,K} \geq 40 \cdot d^{5/2} \cdot (d + d^2 + \ln(4C_{1,K}^d)) \cdot \Delta_K^{1/d} \cdot \delta_K^2$. Then $C_{1,K} = \text{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$. Moreover, using the fact that $\gamma_{\text{avg}} \leq 2^d$, one can check that

$$(\gamma_{\text{avg}}^d \cdot C_{1,K}^d)^{1/d} \geq \gamma_{\text{avg}} \cdot 40 \cdot d^{5/2} \cdot (d + \ln(4 \cdot \gamma_{\text{avg}}^d \cdot C_{1,K}^d)) \cdot \Delta_K^{1/d} \cdot \delta_K^2.$$

This inequality also holds for any $A \geq \gamma_{\text{avg}}^d \cdot C_{1,K}^d$. Hence, any such A with $A \leq \Delta_K^{d^{O(1)}}$ satisfies the conditions of Lemma C.2.8. Now let

$$C_{2,K} = 2e \cdot \frac{4 \cdot d^{5/2} \cdot \delta_K \cdot 4^{1/d}}{\Delta_K^{1/(2d)}} = \text{poly}(\log \Delta_K, \delta_K).$$

We set $\gamma = A^{1/d} \cdot C_{2,K}$ and observe that $\gamma \geq 2e \cdot \gamma_{\text{min}}$ for γ_{min} as in Lemma C.2.8. Then by the Lemmas C.2.8 and IV.2.2 we have:

$$\text{id-HSVP}_\gamma \leq \text{inv-HSVP}_\gamma \leq \text{inv-BHSVP}_\gamma^{2e} \leq \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}},$$

where the second reduction comes from the definition of id-BHSVP (a solution to id-BHSVP $_\gamma^\eta$ in any fractional ideal I is by definition also a solution of id-HSVP $_\gamma$ in L). This completes the proof. \square

Appendix D

Appendices of Chapter V

D.1 Properties of the Rényi Divergence

We will use the following result that bounds the Rényi divergence between two zero-centered normal distributions over $K_{\mathbb{R}}$. It follows from standard divergence bounds on Gaussians such as in [GAL13, Table 2] (note that in this work the Rényi divergence is the logarithm of ours).

Lemma D.1.1. *Let $a, b \in K_{\mathbb{R}}^+$. Let $\mathbf{a} = (\sigma_i(a))_{i \in [d_{\mathbb{R}} + d_{\mathbb{C}}]}$ and $\mathbf{b} = (\sigma_i(b))_{i \in [d_{\mathbb{R}} + d_{\mathbb{C}}]}$. If $2b_i - a_i > 0$ for all $i \in [d_{\mathbb{R}} + d_{\mathbb{C}}]$, then we have*

$$\text{RD}_2(\mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{a}) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{b})) \leq \mathcal{N} \left(\frac{b^2}{a(2b - a)} \right)^{\frac{1}{2}}.$$

We will also use the following technical lemma on the Rényi divergence of a product of random variables.

Lemma D.1.2. *Let X, Y be independent random variables in \mathbb{R} with respective probability distributions D_X, D_Y . Assume that D_X is non-zero over \mathbb{R} (whereas Y can even be discrete). Then*

$$\text{RD}_2(X \cdot Y \parallel X) \leq \left(\mathbb{E}_{y \sim D_Y} (\text{RD}_2(X \cdot y \parallel X))^{\frac{1}{2}} \right)^2.$$

Proof. Let D' be the distribution probability of $X \cdot Y$. We have, for all $t \in \mathbb{R}$:

$$D'(t) = \int_y D_Y(y) D_X\left(\frac{t}{y}\right) dy.$$

This implies that:

$$\begin{aligned} \text{RD}_2(X \cdot Y \parallel X) &= \int_t \frac{1}{D_X(t)} \left(\int_y D_Y(y) D_X\left(\frac{t}{y}\right) dy \right)^2 dt \\ &= \int_{y_1, y_2} D_Y(y_1) D_Y(y_2) \int_t \frac{D_X\left(\frac{t}{y_1}\right) D_X\left(\frac{t}{y_2}\right)}{D_X(t)} dt dy_1 dy_2. \end{aligned}$$

By the Cauchy-Schwartz inequality, we have

$$\begin{aligned} \left(\int_t \frac{D_X\left(\frac{t}{y_1}\right) D_X\left(\frac{t}{y_2}\right)}{D_X(t)} dt \right)^2 &\leq \int_t \frac{\left(D_X\left(\frac{t}{y_1}\right)\right)^2}{D_X(t)} dt \cdot \int_t \frac{\left(D_X\left(\frac{t}{y_2}\right)\right)^2}{D_X(t)} dt \\ &= \text{RD}_2(X \cdot y_1 \parallel X) \cdot \text{RD}_2(X \cdot y_2 \parallel X). \end{aligned}$$

Overall, we obtain that

$$\text{RD}_2(X \cdot Y \parallel X) \leq \left(\int_y D_Y(y) \left(\text{RD}_2(X \cdot y \parallel X) \right)^{\frac{1}{2}} dy \right)^2,$$

which completes the proof. \square

D.2 Missing proofs from Section V.2

D.2.1 Proof of Lemma V.2.2

We first recall the lemma statement.

Lemma V.2.2. *There exists an absolute polynomial P such that the following holds. For any $\delta \geq 0$, degree- d number field K , integer $k \geq 2$, rank- k module $M \subset K_{\mathbb{R}}^k$, if $\mathbf{c} \in \text{span}_{K_{\mathbb{R}}}(M)$ and $\varsigma > 0$ are such that $\|\mathbf{c}\| \leq \delta \cdot \varsigma$ and $\varsigma \geq \lambda_{kd}(M) \cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M)/\lambda_1(M))$, then it holds that*

$$\Pr_{\mathbf{v} \leftarrow D_{M, \varsigma, \mathbf{c}}}(\mathbf{v} \cdot \mathcal{O}_K \text{ is primitive in } M) \geq \frac{1}{4\zeta_K(k)},$$

where $\zeta_K(\cdot)$ is the Dedekind zeta function of K and the λ_i 's refer to the minima of the lattice $\Phi(M)$.

Before proving the lemma, we recall some facts regarding the Dedekind zeta function (see, e.g., [Neu13, Chapter 7] for more details). First, let us define the Möbius function of a field K . It is defined over integral ideals of \mathcal{O}_K by

$$\mu_K \left(\prod_{i=1}^r \mathfrak{p}_i^{e_i} \right) := \begin{cases} 1 & \text{if } r = 0 \\ (-1)^r & \text{if } e_1 = \dots = e_r = 1 \\ 0 & \text{otherwise} \end{cases}$$

where the \mathfrak{p}_i 's are distinct prime ideals. For any $s > 1$, the two following equations holds, where the sums are over integral ideals of \mathcal{O}_K :

$$\zeta_K(s) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^s}$$

$$\zeta_K(s)^{-1} = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{\mu_K(\mathfrak{a})}{\mathcal{N}(\mathfrak{a})^s}.$$

The Dedekind zeta function is well-defined for any $s > 1$ (i.e., the sums above are absolutely converging for $s > 1$).

Lemma D.2.1. *Let $H(N) := |\{\mathfrak{a} \subseteq \mathcal{O}_K \text{ ideal} \mid \mathcal{N}(\mathfrak{a}) \leq N\}|$, for $N \geq 0$. For any $s > 1$ and number field K , it holds that $H(N) \leq \zeta_K(s) \cdot N^s$.*

Proof. This follows from $\zeta_K(s) \geq \sum_{\mathfrak{a}, \mathcal{N}(\mathfrak{a}) \leq N} \frac{1}{\mathcal{N}(\mathfrak{a})^s} \geq H(N)/N^s$. \square

Lemma D.2.2. *For any $s \geq 3/2$ and degree- d number field K , it holds that $\zeta_K(s) \leq 2^{2d}$.*

Proof. We use the Euler product form of the Dedekind zeta function

$$\begin{aligned}
\zeta_K(s) &= \prod_{\substack{\mathfrak{p} \subset \mathcal{O}_K \\ \mathfrak{p} \text{ prime}}} \frac{1}{1 - \mathcal{N}(\mathfrak{p})^{-s}} \\
&= \prod_{\substack{p \in \mathbb{Z} \\ p \text{ prime}}} \prod_{\substack{\mathfrak{p} | p\mathcal{O}_K \\ \mathfrak{p} \text{ prime}}} \frac{1}{1 - \mathcal{N}(\mathfrak{p})^{-s}} \\
&\leq \prod_{\substack{p \in \mathbb{Z} \\ p \text{ prime}}} \prod_{\substack{\mathfrak{p} | p\mathcal{O}_K \\ \mathfrak{p} \text{ prime}}} \frac{1}{1 - p^{-s}} \\
&\leq \prod_{\substack{p \in \mathbb{Z} \\ p \text{ prime}}} \left(\frac{1}{1 - p^{-s}} \right)^d \\
&\leq \zeta_{\mathbb{Q}}(s)^d \leq \zeta_{\mathbb{Q}}(3/2)^d \leq 4^d,
\end{aligned}$$

where we used the fact that $\zeta_{\mathbb{Q}}(3/2) \approx 2.6 \leq 4$. \square

Finally, we will use the following notations and facts regarding Gaussian distributions. We let $\rho_{\varsigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \varsigma^2)$. We also write $D_{M, \varsigma, \mathbf{c}}(\mathbf{x})$ the probability that the Gaussian distribution $D_{M, \varsigma, \mathbf{c}}$ outputs the vector \mathbf{x} , i.e., $D_{M, \varsigma, \mathbf{c}}(\mathbf{x}) = \rho_{\varsigma, \mathbf{c}}(\mathbf{x}) / \rho_{\varsigma, \mathbf{c}}(M)$:

- From [Ban93, Lemma 1.5], we know that for any rank- n lattice L and $c > 1/\sqrt{2\pi}$, we have $\rho_{\varsigma, \mathbf{c}}(\{\mathbf{v} \in L \mid \|\mathbf{v} - \mathbf{c}\| > c \cdot \sqrt{n} \cdot \varsigma\}) \leq 2C^n \cdot \rho_{\varsigma}(L)$, where $C = c \cdot \sqrt{2\pi}e \cdot e^{-\pi \cdot c^2} < 1$.
- From [MR07, Lemma 3.3], we know that for any $\varepsilon > 0$ and rank- n lattice L , the smoothing parameter of L satisfies $\eta_{\varepsilon}(L) \leq \sqrt{\ln(2n(1 + 1/\varepsilon))} / \pi \cdot \lambda_n(L)$.
- Finally, from the proof of [MR07, Lemma 4.4], we know that if $\varsigma \geq \eta_{\varepsilon}(L)$, then it holds that $\rho_{\varsigma, \mathbf{c}}(L) \in [1 - \varepsilon, 1 + \varepsilon] \cdot \varsigma^n / \det(L)$.

Lemma V.2.2. We follow the same proof structure as in [SS13, Lemma 4.4]. Let us fix some number field K of degree d , integer $k \geq 2$, rank- k module $M \subset K_{\mathbb{R}}^k$ and real number $\delta \geq 0$. Let

$$\varsigma_0 = 2^{45} \cdot \Delta_K^{7/(2d)} \cdot k^8 \cdot d^4 \cdot (k^2 \cdot d^2 + \delta^3) \cdot \lambda_1(M)^{-3} \cdot \lambda_{kd}(M)^4.$$

Observe that $\varsigma_0 = \lambda_{kd}(M) \cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M) / \lambda_1(M))$ for some absolute polynomial P . We will prove that the lemma holds for this polynomial P .

Let us then fix some ς and $\mathbf{c} \in \text{span}_{K_{\mathbb{R}}}(M)$ such that $\|\mathbf{c}\| \leq \delta \cdot \varsigma$ and $\varsigma \geq \varsigma_0$. We define the following quantities.

$$\begin{aligned}
\varepsilon &= 2^{-2kd-5} \\
B_1 &= \frac{\varsigma^d}{\sqrt{\Delta_K} \cdot \lambda_{kd}(M)^d \cdot \max(\varepsilon^{-1/k}, (2 \ln(2/\varepsilon))^d)} \\
B_2 &= (2 \cdot \sqrt{kd} \cdot \varsigma + \|\mathbf{c}\|)^d \cdot \sqrt{\Delta_K} \cdot \lambda_1(M)^{-d}.
\end{aligned}$$

With these notations, we are ready to bound from below the probability that a Gaussian element in M is primitive. To do so, observe that for $\mathbf{v} \in M$, the rank-1 module $\mathbf{v}\mathcal{O}_K$ is not primitive in M if and only if there exists some prime ideal \mathfrak{p} such that $\mathbf{v} \in \mathfrak{p} \cdot M$. Indeed, let

us define $I = \{x \in K \mid x \cdot \mathbf{v} \in \mathcal{O}_K\}$. One can check that I is a fractional ideal with $\mathcal{O}_K \subseteq I$. Moreover, by definition of a primitive submodule, we have $I = \mathcal{O}_K$ if and only if $\mathbf{v} \cdot \mathcal{O}_K$ is a primitive submodule of M . Let $\mathfrak{a} = I^{-1}$, which is an integral ideal. By definition of I and \mathfrak{a} , we have that $\mathbf{v} \in \mathfrak{a} \cdot M$. If $\mathbf{v} \cdot \mathcal{O}_K$ is not primitive, then $\mathfrak{a} \neq \mathcal{O}_K$ so there exists $\mathfrak{p} \mid \mathfrak{a}$, and it holds that $\mathbf{v} \in \mathfrak{p} \cdot M$. Reciprocally, if $\mathbf{v} \in \mathfrak{p} \cdot M$ for some prime ideal \mathfrak{p} , then $\mathbf{v} \cdot \mathfrak{p}^{-1} \subset M$, so $I \neq \mathcal{O}_K$ and $\mathbf{v} \cdot \mathcal{O}_K$ is not primitive in M .

From this observation, we can rewrite

$$\begin{aligned} \Pr_{\mathbf{v} \leftarrow D_{M,\varsigma,c}} \left(\mathbf{v} \cdot \mathcal{O}_K \text{ is primitive in } M \right) &= D_{M,\varsigma,c} \left(M \setminus \bigcup_{\mathfrak{p} \text{ prime}} \mathfrak{p}M \right) \\ &\geq D_{M,\varsigma,c}^T \left(M \setminus \bigcup_{\mathfrak{p} \text{ prime}} \mathfrak{p}M \right), \end{aligned}$$

with $D_{M,\varsigma,c}^T$ being the truncated Gaussian function, defined as $D_{M,\varsigma,c}^T(\mathbf{v}) = D_{M,\varsigma,c}(\mathbf{v})$ if $\mathbf{v} \neq \mathbf{0}$ and $\|\mathbf{v} - \mathbf{c}\| \leq 2 \cdot \sqrt{kd}$, and $D_{M,\varsigma,c}^T(\mathbf{v}) = 0$ otherwise (note that $D_{M,\varsigma,c}^T$ does not sum to 1 and is not a probability distribution).

Let us then focus on $p := D_{M,\varsigma,c}^T(M \setminus \bigcup_{\mathfrak{p} \text{ prime}} \mathfrak{p}M)$. Observe that for any distinct prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_t$, it holds that $\bigcap_{i \leq t} (\mathfrak{p}_i \cdot M) = (\prod_{i \leq t} \mathfrak{p}_i) \cdot M$. Hence, from the inclusion-exclusion principle, we obtain

$$p = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \mu_K(\mathfrak{a}) \cdot D_{M,\varsigma,c}^T(\mathfrak{a} \cdot M) = \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_2}} \mu_K(\mathfrak{a}) \cdot D_{M,\varsigma,c}^T(\mathfrak{a} \cdot M),$$

where the sums are above the integral ideals $\mathfrak{a} \subseteq \mathcal{O}_K$ and B_2 was defined at the start of the proof. The second equality comes from the fact that if $\mathcal{N}(\mathfrak{a}) > B_2$, then $\mathfrak{a} \cdot M$ does not contain any non-zero vector shorter than $2\sqrt{kd} \cdot \varsigma + \|\mathbf{c}\|$ (since otherwise $M = \mathfrak{a}^{-1} \cdot (\mathfrak{a} \cdot M)$ would contain a non-zero vector smaller than $(2\sqrt{kd} \cdot \varsigma + \|\mathbf{c}\|) \cdot \mathcal{N}(\mathfrak{a})^{-1/d} \cdot \Delta_K^{1/(2d)} < \lambda_1(M)$, contradicting the definition of $\lambda_1(M)$). This implies that $\mathfrak{a} \cdot M$ does not contain any non-zero vector in the ball $\{\mathbf{v} \mid \|\mathbf{v} - \mathbf{c}\| \leq 2\sqrt{kd} \cdot \varsigma\}$, hence $D_{M,\varsigma,c}^T(\mathfrak{a} \cdot M) = 0$.

Combining this with the equation recalled before the proof relating the Dedekind zeta function and the Möbius function, we obtain

$$\begin{aligned} |p - \zeta_K(k)^{-1}| &\leq \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \left| D_{M,\varsigma,c}^T(\mathfrak{a} \cdot M) - \mathcal{N}(\mathfrak{a})^{-k} \right| \\ &\quad + \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2}} D_{M,\varsigma,c}^T(\mathfrak{a} \cdot M) + \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) > B_1}} \mathcal{N}(\mathfrak{a})^{-k} \end{aligned}$$

Note that by definition of B_1 and B_2 , it holds that $B_1 \leq B_2$. We will bound each one of the three sums above by $\zeta_K(k)^{-1}/4$, which will prove the result.

Let us start with the first sum. Let \mathfrak{a} be an integral ideal with $\mathcal{N}(\mathfrak{a}) \leq B_1$. We know that $\lambda_{kd}(\mathfrak{a} \cdot M) \leq \lambda_1^\infty(\mathfrak{a}) \cdot \lambda_{kd}(M)$ (since multiplying kd linearly independent vectors from M by a shortest vector of \mathfrak{a} provides kd linearly independent vectors of $\mathfrak{a}M$). Moreover, since $\mathcal{N}(\mathfrak{a}) \leq B_1$, we know that $\lambda_1^\infty(\mathfrak{a}) \leq B_1^{1/d} \cdot \Delta_K^{1/(2d)}$. By definition of B_1 and ε , we have

$$\varsigma \geq B_1^{\frac{1}{d}} \cdot \Delta_K^{\frac{1}{2d}} \cdot \lambda_{kd}(M) \cdot 2\ln(2/\varepsilon) \geq \lambda_{kd}(\mathfrak{a}M) \cdot 2\ln(2/\varepsilon) \geq \eta_\varepsilon(\mathfrak{a}M).$$

Since $\varsigma \geq \eta_\varepsilon(\mathfrak{a}M)$, we know that $\rho_{\varsigma,c}(\mathfrak{a} \cdot M) \in [1 - \varepsilon, 1 + \varepsilon] \cdot \varsigma^{kd} / \det(\mathfrak{a}M)$. From [Ban93, Lemma 1.5] (recalled above) with $c = 2$, we also know that $\rho_{\varsigma,c}(\{\mathbf{v} \in \mathfrak{a}M \mid \|\mathbf{v} - \mathbf{c}\| > 2 \cdot \sqrt{kd} \cdot \varsigma\}) \leq \varepsilon \cdot \rho_\varsigma(\mathfrak{a}M) \leq 2\varepsilon \cdot \varsigma^{kd} / \det(\mathfrak{a}M)$.

Recall also that by definition of B_1 , we have $\varsigma \geq \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d} \cdot \lambda_{kd}(M) \cdot \varepsilon^{-1/(kd)}$. Hence, we obtain that $\rho_{\varsigma, \mathfrak{c}}(\mathbf{0}) \leq \varepsilon \cdot \varsigma^{kd} / \det(\mathfrak{a}M)$. Combining everything, this implies that

$$\rho_{\varsigma, \mathfrak{c}}\left(\mathfrak{a} \cdot M \setminus \{\mathbf{v} \mid \mathbf{v} = \mathbf{0} \text{ or } \|\mathbf{v} - \mathfrak{c}\| > 2 \cdot \sqrt{kd} \cdot \varsigma\}\right) \in [1 - 4\varepsilon, 1 + \varepsilon] \cdot \frac{\varsigma^{kd}}{\det(\mathfrak{a}M)}.$$

By definition of $D_{M, \varsigma, \mathfrak{c}}^T$, this implies that

$$D_{M, \varsigma, \mathfrak{c}}^T(\mathfrak{a} \cdot M) \in \left[\frac{1 - 4\varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon}\right] \cdot \frac{\det(M)}{\det(\mathfrak{a}M)} \subset [1 - 5\varepsilon, 1 + 4\varepsilon] \cdot \frac{1}{\mathcal{N}(\mathfrak{a})^k},$$

where we used the identities $\det(M) = \Delta_K^{k/2} \cdot \mathcal{N}(M)$ and $\det(\mathfrak{a}M) = \Delta_K^{k/2} \cdot \mathcal{N}(\mathfrak{a}M)$. This concludes the upper bound on the first sum

$$\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \left| D_{M, \varsigma, \mathfrak{c}}^T(\mathfrak{a} \cdot M) - \mathcal{N}(\mathfrak{a})^{-k} \right| \leq 5\varepsilon \cdot \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \frac{1}{\mathcal{N}(\mathfrak{a})^k} \leq 5\varepsilon \zeta_K(k) \leq \frac{1}{4\zeta_K(k)},$$

by definition of ε , and using Lemma D.2.2 to assert that $\zeta_K(k)^2 \leq 2^{4d} \leq 2^{2kd}$.

Let us now consider the second sum. Let \mathfrak{a} be an integral ideal with $B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2$. Let us define $I = \lceil (\mathcal{N}(\mathfrak{a})/B_1)^{1/d} \rceil^{-1} \cdot \mathfrak{a}$. This is a fractional ideal with $\mathcal{N}(I) \in [1/2^d, 1] \cdot B_1$. Moreover, we have $\mathfrak{a} \subseteq I$, hence $D_{M, \varsigma, \mathfrak{c}}^T(\mathfrak{a} \cdot M) \leq D_{M, \varsigma, \mathfrak{c}}^T(I \cdot M)$, where we let $D_{M, \varsigma, \mathfrak{c}}^T(\mathbf{v}) := \rho_{\varsigma, \mathfrak{c}}(\mathbf{v}) / \rho_{\varsigma, \mathfrak{c}}(M)$ for any $\mathbf{v} \in K_{\mathbb{R}}$, even those not in M (note that since I is a fractional ideal, then $I \cdot M$ needs not be contained in M). Observe that everything we did above when $\mathcal{N}(\mathfrak{a}) \leq B_1$ can be adapted to a fractional ideal I of norm $\mathcal{N}(I) \leq B_1$. Hence, we have (using the analysis for the first sum):

$$D_{M, \varsigma, \mathfrak{c}}^T(I \cdot M) \leq (1 + 4\varepsilon) \cdot \frac{1}{\mathcal{N}(I)^k} \leq \frac{2^{kd+1}}{B_1^k}.$$

We can hence bound the second sum from above by

$$\begin{aligned} \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2}} D_{M, \varsigma, \mathfrak{c}}^T(\mathfrak{a} \cdot M) &\leq \frac{|\{\mathfrak{a} \subseteq \mathcal{O}_K \mid \mathcal{N}(\mathfrak{a}) \leq B_2\}| \cdot 2^{kd+1}}{B_1^k} \\ &\leq \frac{\zeta_K(s_0) \cdot B_2^{s_0} \cdot 2^{kd+1}}{B_1^k} \\ &\leq B_1^{-(k-s_0)} \cdot \left(\zeta_K(s_0) \cdot 2^{kd+1} \cdot (B_2/B_1)^{s_0} \right), \end{aligned}$$

where we used Lemma D.2.1 for the second inequality, with some $s_0 \in (1, k)$. Let us choose $s_0 = \max(3/2, k/2)$. This choice of s_0 ensures that $s_0 \in [3/2, k)$ and that $s_0/(k-s_0) \leq 3$ and $k/(k-s_0) \leq 4$ for any $k \geq 2$. Using Lemma D.2.2, the definitions of s_0 , B_1 , B_2 and the lower bound on ς , one can check that this is $\leq 1/4 \cdot \zeta_K(k)^{-1}$.

We are finally left with the last sum. Recall that $H(N)$ denotes the number of integral ideals of norm $\leq N$. With this notation, we can rewrite

$$\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) > B_1}} \mathcal{N}(\mathfrak{a})^{-k} = \sum_{N > B_1} \frac{H(N) - H(N-1)}{N^k} \leq \sum_{N > B_1} H(N) \cdot \left(\frac{1}{N^k} - \frac{1}{(N+1)^k} \right).$$

Let us prove that the last sum above is absolutely converging (in order to prove that our transformation was valid). Using Lemma D.2.1 with $s = 1.5$, we know that $H_N \leq \zeta_K(1.5) \cdot N^{1.5}$, where the quantity $\zeta_K(1.5)$ depends on the number field but is fixed when N tends to infinity. Hence, the quantity inside the sum is bounded by $O(N^{1.5} \cdot N^{k-1}/N^{2k}) = O(1/N^{k-0.5}) = O(1/N^{1.5})$ since $k \geq 2$. We conclude that the sum is converging absolutely as desired.

Let us now compute an upper bound on this sum. Since $N \geq B_1 \geq k$, we know that $(N+1)^k - N^k \leq k^2 \cdot N^{k-1}$. Applying Lemma D.2.1 again with $s = s_0 \in (1, k)$, we obtain

$$\begin{aligned} \sum_{N > B_1} H(N) \cdot \left(\frac{1}{N^k} - \frac{1}{(N+1)^k} \right) &\leq \zeta_K(s_0) \cdot \sum_{N > B_1} \frac{k^2 \cdot N^{s_0+k-1}}{N^{2k}} \\ &\leq \zeta_K(s_0) \cdot k^2 \cdot \int_{\lfloor B_1 \rfloor}^{+\infty} x^{-(k+1-s_0)} dx \\ &= \zeta_K(s_0) \cdot k^2 \cdot (k-s_0)^{-1} \cdot \lfloor B_1 \rfloor^{-(k-s_0)} \end{aligned}$$

Using $s_0 = \max(3/2, k/2)$ again, the definitions of B_1 , the lower bound on ς and Lemma D.2.2, one can check that $\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) > B_1}} \mathcal{N}(\mathfrak{a})^{-k} \leq 1/4 \cdot \zeta_K(k)^{-1}$ as desired. \square

D.2.2 Proof of Lemma V.2.3

Let N be a densest rank-1 submodule of M . By Definition II.3.3, there exists a rank-1 submodule N' such that $M = N + N'$. Equivalently, we obtain a pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of M such that $N = \mathbf{b}_1 I_1$. Wlog, we may assume that $\mathcal{N}(I_1) = \mathcal{N}(I_2) = 1$, by multiplying \mathbf{b}_i by $\mathcal{N}^{-1/d}(I_i)$ for $i \in [2]$. Let $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$ and $\mathbf{R} \in K_{\mathbb{R}}^{2 \times 2}$ upper triangular such that $\mathbf{B} = \mathbf{Q}\mathbf{R}$. Let \mathbf{D} be the diagonal matrix with $d_1 = r_{11}/\mathcal{N}^{1/d}(r_{11})$ and $d_2 = r_{22}/\mathcal{N}^{1/d}(r_{22})$ as diagonal coefficients. Let $J_1 = d_1 I_1$, $J_2 = d_2 I_2$ and

$$\mathbf{B}' = \frac{\mathcal{N}^{\frac{1}{2d}}(M)}{\gamma} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot \left(\frac{\gamma}{\mathcal{N}^{\frac{1}{2d}}(M)} \mathbf{D}^{-1} \right).$$

It now suffices to prove that $((J_1, \mathbf{b}'_1), (J_2, \mathbf{b}'_2))$ is a pseudo-basis of M of the desired form, i.e., to check that $\mathbf{R}' = \mathbf{R} \cdot (\gamma/\mathcal{N}^{1/(2d)}(M)) \mathbf{D}^{-1}$ has diagonal coefficients equal to 1 and γ . We have $r'_{ii} = \mathcal{N}^{1/d}(r_{ii})(\gamma/\mathcal{N}^{1/(2d)}(M))$ for $i \in [2]$, by construction. The fact that $\mathcal{N}(r_{11}) = \lambda_1^{\mathcal{N}}(M)$ gives that $r'_{11} = 1$. The equality $\mathcal{N}(M) = \det(\mathbf{B}')$ provides the result. \square

D.2.3 Proof of Lemma V.2.5

From Banaszczyk's transference theorem (Theorem II.1.5), we know that $1 \leq \lambda_{2d}(M^{\vee}) \cdot \lambda_1(M) \leq 2d$. We also know that $\lambda_1(M)^d \geq \sqrt{d} \cdot \lambda_1^{\mathcal{N}}(M)$, since for any vector $\mathbf{v} \in K_{\mathbb{R}}^2$ it holds that $\|\mathbf{v}\| \geq \sqrt{d} \cdot \mathcal{N}(\mathbf{v} \cdot \mathcal{O}_K)^{1/d}$ (by applying the inequality of arithmetic and geometric means to the squares of the coordinates of $\langle \mathbf{v}, \mathbf{v} \rangle_{K_{\mathbb{R}}}$). Further, from the definition of the gap of a module M , we know that $\lambda_1^{\mathcal{N}}(M) = \mathcal{N}(M)^{1/2}/\gamma(M)^d$. Combining these relations provides the upper bound on $\lambda_{2d}(M^{\vee})$.

In order to get the upper bound on $\lambda_1(M^{\vee})^{-1}$, we use the inequality $1 \leq \lambda_{2d}(M) \cdot \lambda_1(M^{\vee})$. Hence, it suffices to bound $\lambda_{2d}(M)$ from above. To do so, we use Lemma V.2.3. We know that there exist d \mathbb{Z} -linearly independent vectors in J_1 of norms $\leq \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$, and similarly in J_2 . Hence, from the representation of M in Lemma V.2.3, we obtain $2d$ linearly independent vectors in M of norms $\leq (\gamma(M) \cdot \sqrt{d} + d/\gamma(M)) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$ (where we reduced the last d vectors using the first d ones). Since $\gamma(M) \geq 1$, this implies that $\lambda_{kd}(M) \leq 2d \cdot \gamma(M) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$. \square

D.3 Missing Proofs from Section V.3

D.3.1 Proof of Lemma V.3.2

The fact that $\overline{\mathbf{b}^\vee} \neq \mathbf{0}$ implies that the map $\mathbf{m} \mapsto \langle \mathbf{b}^\vee, \mathbf{m} \rangle_{K_{\mathbb{R}}}$ is a surjective homomorphism from M to $\mathcal{O}_K/\mathfrak{p}$ whose kernel is M' . This gives the following exact sequence of \mathcal{O}_K -modules:

$$0 \rightarrow M' \rightarrow M \rightarrow \mathcal{O}_K/\mathfrak{p} \rightarrow 0.$$

Now, note that $\mathcal{O}_K/\mathfrak{p}$ is isomorphic to the finite field of size $\mathcal{N}(\mathfrak{p})$. The exact sequence and the finiteness of $\mathcal{O}_K/\mathfrak{p}$ imply that $\mathcal{N}(M') = \mathcal{N}(M) \cdot |\mathcal{O}_K/\mathfrak{p}|$. The proof is completed by noting that $|\mathcal{O}_K/\mathfrak{p}| = \mathcal{N}(\mathfrak{p})$. \square

D.3.2 Proof of Lemma V.3.3

The fact that $\mathfrak{p} \cdot M \subset M'$ implies that $\mathfrak{p}I \subset M'$. We now prove the second property. As $\mathbf{b}I$ is primitive, there exists a pseudo-basis (\mathbf{B}, \mathbb{I}) of M such that $\mathbf{b}_1 = \mathbf{b}$ and $I_1 = I$ (see Definition II.3.3). We start by noting that $\langle M^\vee, \mathbf{u} \rangle_{K_{\mathbb{R}}} \cdot I = \mathcal{O}_K$. Indeed, as (\mathbf{B}, \mathbb{I}) is a pseudo-basis of M , we have that $(\mathbf{B}^{-\dagger}, \mathbb{J})$ is a pseudo-basis of M^\vee , with $J_i = (\overline{I_i})^{-1}$ for all i . Therefore:

$$\langle M^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \cdot I = \sum_i \langle \mathbf{b}_i^{-\dagger}, \mathbf{b}_1 \rangle_{K_{\mathbb{R}}} \cdot \mathcal{O}_K = \mathcal{O}_K.$$

The fact that $\langle M^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \cdot I = \mathcal{O}_K$ implies that the scalar product with \mathbf{b} is a surjective homomorphism $M^\vee \rightarrow I^{-1}$. This induces a surjective homomorphism $M^\vee/\mathfrak{p}M^\vee \rightarrow I^{-1}/\mathfrak{p}I^{-1}$. Because of their respective ranks as \mathcal{O}_K -modules, the cardinality of $I^{-1}/\mathfrak{p}I^{-1}$ is $\mathcal{N}(\mathfrak{p})$ and the cardinality of $M^\vee/\mathfrak{p}M^\vee$ is $\mathcal{N}(\mathfrak{p})^k$. Lagrange's theorem (for groups) then implies that every element of $I^{-1}/\mathfrak{p}I^{-1}$ has exactly $\mathcal{N}(\mathfrak{p})^k/\mathcal{N}(\mathfrak{p}) = \mathcal{N}(\mathfrak{p})^{k-1}$ pre-images in $M^\vee/\mathfrak{p}M^\vee$ by this application. In particular, the zero element of $I^{-1}/\mathfrak{p}I^{-1}$ has $\mathcal{N}(\mathfrak{p})^{k-1}$ pre-images, including $\mathbf{0}$. Since $\overline{\mathbf{b}^\vee}$ is uniform in $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$, this implies that the probability that $\langle \mathbf{b}^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}I^{-1}$ is $(\mathcal{N}(\mathfrak{p})^{k-1} - 1)/\mathcal{N}(\mathfrak{p})^k = 1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^k$ over the choice of $\overline{\mathbf{b}^\vee}$.

To complete the proof, note that $\langle \mathbf{b}^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \notin \mathfrak{p}I^{-1}$ is equivalent to $\mathbf{b}I \not\subset M'$, by definition of M' . \square

D.3.3 Proof of Lemma V.3.4

Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of M with integral coefficient ideals I_i . As seen in Section II.3, the pair $(\mathbf{B}^{-\dagger}, \mathbb{J})$ is a pseudo-basis of M^\vee , where $J_i = (\overline{I_i})^{-1}$ for all i . Take $\mathbf{u} \in \mathbb{J}$ such that $\mathbf{B}^{-\dagger} \cdot \mathbf{u}$ is a representative of $\overline{\mathbf{b}^\vee}$ in M^\vee . We have:

$$\begin{aligned} M' &= \left\{ \mathbf{B} \cdot \mathbf{v} : \mathbf{v} \in \mathbb{I} \text{ and } (\mathbf{B}^{-\dagger} \cdot \mathbf{u})^\dagger \cdot \mathbf{B} \cdot \mathbf{v} \in \mathfrak{p} \right\} \\ &= \left\{ \mathbf{B} \cdot \mathbf{v} : \mathbf{v} \in \mathbb{I} \text{ and } \langle \mathbf{u}, \mathbf{v} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p} \right\}. \end{aligned}$$

Let us define

$$N = \{ \mathbf{v} \in \mathbb{I} : \langle \mathbf{u}, \mathbf{v} \rangle_{K_{\mathbb{R}}} = 0 \} \quad \text{and} \quad N' = \{ \mathbf{v} \in \mathbb{I} : \langle \mathbf{u}, \mathbf{v} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p} \}.$$

We use the \mathbb{Z} -basis $\mathbf{B}_{\mathcal{O}_K} = (b_i)_{i \in [d]}$ of \mathcal{O}_K to identify N with a \mathbb{Z} -lattice corresponding to the orthogonal of an integer vector. A basis of this lattice can be computed in polynomial-time, and

the basis vectors provide a set $(\mathbf{n}_i)_{i \in [kd]}$ of (non K -linearly independent) vectors in \mathcal{O}_K^k such that $N = \sum_i \mathbf{n}_i \mathcal{O}_K$. The module N' is the rank- k module generated by the pseudo-basis

$$N' = \sum_{i=1}^{kd} \mathbf{n}_i \mathcal{O}_K + \sum_{i=1}^k \mathbf{e}_i \mathfrak{p},$$

where \mathbf{e}_i is the i -th canonical unit vector. From the pairs $\{(\mathcal{O}_K, \mathbf{n}_i)\}_i$ and $\{(\mathfrak{p}, \mathbf{e}_i)\}_i$, we compute a Hermite Normal Form $(\mathbf{B}', \mathbb{I}')$ of the integral module N' . By definition of N' , the pair $(\mathbf{B} \cdot \mathbf{B}', \mathbb{I}')$ is a pseudo-basis of M' . \square

D.3.4 Proof of Lemma V.3.5

In Step 2, we use one of the provable variants of the BKZ algorithm mentioned above, which allows us to obtain a basis \mathbf{C} of M^\vee such that $\max_i \|\mathbf{c}_i\| \leq (kd)^{kd/\beta+1} \cdot \lambda_{kd}(M^\vee)$ in time polynomial in the bitsize of the input basis of M^\vee and in 2^β . Note that these analyses of the algorithm under scope only prove that the algorithm solves $(kd)^{kd/\beta}$ -SVP (i.e., outputs one short non-zero vector) and do not mention the approximation factor obtained for SIVP (the Shortest Independent Vector Problem). Hence, to obtain an upper bound on $\max_i \|\mathbf{c}_i\|$, we also use the polynomial-time reduction from $(\sqrt{n}\gamma)$ -SIVP to γ -SVP for lattices of rank- n (see [Ste15, Page 1]), together with the fact that one can transform any set of n short linearly independent vectors of norm $\leq B$ in a rank- n lattice L into a basis of L with vectors of norms $\leq \sqrt{n} \cdot B$. Now, we observe that $\varsigma \geq \sqrt{kd} \cdot \max_i \|\mathbf{c}_i\| \geq \sqrt{kd} \cdot \max_i \|\mathbf{c}_i^*\|$, hence we can apply Lemma II.1.12. This means in particular that the vectors \mathbf{y}_i can be sampled in polynomial time, which completes the runtime analysis.

Let us now prove that the matrix \mathbf{Y} satisfies the conditions of the theorem. First of all, note that since the vectors \mathbf{y}_i are in M^\vee , then for all $\mathbf{v} \in M$ we have $\mathbf{Y} \cdot \mathbf{v} \in \mathcal{O}_K^k$, which proves the first point. For the second point, recall that we use a tail-cut distribution $\tilde{D}_{\mathcal{C}^\vee, \varsigma, \mathbf{t}_i}$ with error 2^{d+3} , hence it holds by Lemma II.1.12 that

$$\|\mathbf{y}_i - \mathbf{t}_i\| \leq \sqrt{5kd} \cdot \varsigma = \varepsilon \cdot R,$$

as desired.

Finally, recall from Lemma II.1.12 that the distribution $\tilde{D}_{\mathcal{C}^\vee, \varsigma, \mathbf{t}_i}$ (which might depend on \mathcal{C}^\vee and hence on (\mathbf{B}, \mathbb{I})) is within statistical distance at most 2^{-kd} from the Gaussian distribution $D_{M^\vee, \varsigma, \mathbf{t}_i}$, which is independent of the known basis of M^\vee . Hence, the distribution of \mathbf{Y} is within statistical distance at most $k \cdot 2^{-\Omega(kd)} = 2^{-\Omega(kd)}$ from a distribution independent of the choice of the pseudo-basis (\mathbf{B}, \mathbb{I}) . \square

D.3.5 Proof of Lemma V.3.6

Wlog, we prove the result for $R = 1$. Note that the operator norm of \mathbf{E} satisfies $\|\mathbf{E}\| \leq k\varepsilon < 1$. Therefore, the matrix $\sum_{i \geq 0} (-\mathbf{E})^i$ is well-defined, and satisfies $\mathbf{Y}^{-1} = \sum_{i \geq 0} (-\mathbf{E})^i$. We have $\mathbf{Y}^{-1} = \mathbf{I}_k + \mathbf{E}'$ with $\mathbf{E}' = -\mathbf{E} + \sum_{i \geq 2} (-\mathbf{E})^i$. Using the operator norm again, we obtain that $\|e'_{ij} + e_{ij}\| \leq (k\varepsilon)^2 / (1 - k\varepsilon) \leq k\varepsilon$ for all $i, j \in [k]$, by using assumption that $k\varepsilon \leq 1/2$. This proves the first statement.

By Hadamard's inequality, we have

$$\begin{aligned} \det(\mathbf{Y}) &\leq \left(\sqrt{(1 + \varepsilon)^2 + (k - 1)\varepsilon^2} \right)^d \\ \det(\mathbf{Y}^{-1}) &\leq \left(\sqrt{(1 + \varepsilon')^2 + (k - 1)\varepsilon'^2} \right)^d, \end{aligned}$$

with $\varepsilon' = (k+1)\varepsilon$. Simplifying the expressions using the facts that $\varepsilon' \leq 1$ and $k\varepsilon \leq 1/2$ leads to the second statement. \square

D.4 Missing Proofs from Section V.4

D.4.1 Proof of Theorem V.4.1

Theorem V.4.1 is a direct corollary of the following more complete statement.

Theorem D.4.1. *Let K be a number field of degree d and $\gamma^+ > 0$. There exist three algorithms `uSVP-to-NTRU`, `LiftVecInternal` and `LiftModInternal` and $q_0 = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ such that the following holds, for any $q \geq q_0$, $\gamma_{\text{NTRU}} > 1$, $\gamma_{\text{HSVP}} \geq \sqrt{d}\Delta_K^{1/(2d)}$ and a pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 module $M \subset K^2$ with $\gamma(M) \leq \gamma^+$.*

- Algorithm `uSVP-to-NTRU` takes as input (\mathbf{B}, \mathbb{I}) , q and γ_{HSVP} and outputs a $(\mathbf{B}', \mathcal{O}_K^2)$, a pseudo-basis of a rank 2 free module $M' \subset \mathcal{O}_K^2$, together with some auxiliary information \mathbf{aux} . If (\mathbf{B}, \mathbb{I}) is a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance with

$$\gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 16\sqrt{2} \cdot d^{3/2} \cdot \delta_K,$$

then $(\mathbf{B}', \mathcal{O}_K^2)$ is a $(\gamma_{\text{NTRU}}, q)$ -NTRU instance. If given access to a $\gamma_{\text{HSVP-id-HSVP}}$ oracle, it runs in time polynomial in its input bitsize, in $\zeta_K(2)$ and in $\exp(\frac{d \log(d)}{\log(q/q_0)})$ and makes one call to the oracle.

- Algorithm `LiftVecInternal` takes as input a non-zero vector $\mathbf{s}' \in M'$ and the auxiliary information \mathbf{aux} . It outputs a non-zero vector $\mathbf{s} \in M$ such that

$$\|\mathbf{s}\| \leq 150 \cdot \gamma_{\text{HSVP}}^{3/2} \cdot d^{7/2} \cdot \delta_K^2 \cdot \frac{\|\mathbf{s}'\|}{\mathcal{N}(M')^{\frac{1}{2d}}} \cdot \mathcal{N}(M)^{\frac{1}{2d}}.$$

If given access to a $\gamma_{\text{HSVP-id-HSVP}}$ oracle, it runs in polynomial time and makes one call to the oracle.

- Algorithm `LiftModInternal` takes as input a pseudo-basis of a rank-1 densest submodule N' of M' and the auxiliary information \mathbf{aux} and outputs a pseudo-basis of a rank-1 densest submodule N of M . It runs in polynomial time.

Proof. Let $V_0 = \left(\text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+)\right)^{2d}$ be as in Lemma V.4.4 (defined using γ^+ instead of $\gamma(M)$). Define

$$q_0 = \frac{V_0^{1/d} \cdot 4d}{\gamma_{\text{HSVP}}}.$$

One can check that q_0 is indeed $\text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+)$ as desired. We prove that the theorem holds for this choice of q_0 .

Algorithm `uSVP-to-NTRU`. On input (\mathbf{B}, \mathbb{I}) , q and γ_{HSVP} , `uSVP-to-NTRU` sets $V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$ and $\beta = \left\lceil \frac{2d \log(2d)}{\log(\sqrt{q/q_0}) + \log(2d)} \right\rceil$. It then runs `PreCond` on input (\mathbf{B}, \mathbb{I}) , V and β , to obtain a matrix $\mathbf{Y} \in \text{GL}_2(K)$.

From the definition of q_0 , V and β , one can check that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Moreover, we have $\gamma(M) \leq \gamma^+$ by assumption, hence we can apply Lemma V.4.4. This implies in particular

that the call to the `PreCond` algorithm runs in time polynomial in the input bitsize, in $2^\beta = 2^{\mathcal{O}(d \log(d) / \log(q/q_0))}$ and in $\zeta_K(2)$.

Algorithm `uSVP-to-NTRU` then runs `Conditioned-to-NTRU` on input $(\mathbf{Y}\mathbf{B}, \mathbb{I})$, q and γ_{HSVP} . It obtains a basis \mathbf{B}' of a free module M' and some auxiliary information \mathbf{aux}' . Algorithm `uSVP-to-NTRU` finally outputs $(\mathbf{B}', \mathcal{O}_K^2)$ and $\mathbf{aux} = (\mathbf{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$.

We know that the call to `Conditioned-to-NTRU` can be done in polynomial time, with one call to the γ_{HSVP} -id-HSVP oracle. This completes the proof on the running time of algorithm `uSVP-to-NTRU`.

Let us assume now that (\mathbf{B}, \mathbb{I}) was a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance, for γ_{uSVP} as in the theorem. We know from Lemma V.4.4 that $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a $\gamma_{\text{uSVP}/(2\sqrt{2})\text{-mod-uSVP}_2}$ instance. Moreover, still from Lemma V.4.4, we know that the module spanned by $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a rank-2 module in \mathcal{O}_K^2 , with the coprime property and such that $\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$. Hence we can apply Lemma V.4.7 and conclude that $(\mathbf{B}', \mathcal{O}_K^2)$ is a γ_{NTRU} instance, as desired (note that V and $\gamma_{\text{uSVP}}/(2\sqrt{2})$ have the desired shape for applying Lemma V.4.7). This proves the first item of the theorem.

Algorithm LiftVecInternal. On input $s' \in M'$ and $\mathbf{aux} = (\mathbf{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$, algorithm `LiftVecInternal` runs `LiftVec`($\mathbf{aux}', \gamma_{\text{HSVP}}, \mathbf{B}', s'$) and gets a nonzero vector \mathbf{t} . It then outputs $\mathbf{Y}^{-1} \cdot \mathbf{t}$. By Lemma V.4.9, we know that the call to `LiftVec` can be performed in polynomial time, with one call to the id-HSVP oracle. This proves the running time of `LiftVecInternal`.

By Lemma V.4.9 again, we know that $\|\mathbf{t}\| \leq \|s'\| \cdot 68 \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$. From the shape of \mathbf{Y} , Lemma V.3.6 instantiated with $\varepsilon = 1/5$ and Lemma D.4.2, we obtain¹

$$\|\mathbf{Y}^{-1} \cdot \mathbf{t}\| \leq \frac{2.2 \cdot \mathcal{N}(M)^{\frac{1}{2d}}}{V^{\frac{1}{2d}}} \cdot \|\mathbf{t}\| \leq 150 \cdot \gamma_{\text{HSVP}}^{\frac{3}{2}} \cdot d^{\frac{7}{2}} \cdot \delta_K^2 \cdot \frac{\|s'\|}{\sqrt{q}} \cdot \mathcal{N}(M)^{\frac{1}{2d}}.$$

Using the fact that $\mathcal{N}(M') = q^d$ provides the desired upper bound on the output size. Note also that by construction, $\mathbf{Y}^{-1} \cdot \mathbf{t}$ is indeed a non-zero vector in M .

Algorithm LiftModInternal. Let us call \widetilde{M} the intermediate module $(\mathbf{Y} \cdot \mathbf{B}) \cdot \mathbb{I}$ computed by algorithm `uSVP-to-NTRU`.

On input a pseudo-basis (\mathbf{v}', J') of a densest rank-1 module of M' and $\mathbf{aux} = (\mathbf{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$, algorithm `LiftModInternal` runs `LiftMod`($\mathbf{aux}', \mathbf{B}', (\mathbf{v}', J')$) and gets a vector \mathbf{w} . It then computes J such that $\text{span}(\mathbf{w}) \cap \widetilde{M} = \mathbf{w} \cdot J$, sets $\mathbf{v} = \mathbf{Y}^{-1} \cdot \mathbf{w}$ and outputs the pseudo-basis (\mathbf{v}, J) .

From Lemma V.4.8, we know that algorithm `LiftModInternal` runs in polynomial time. Moreover, since (\mathbf{v}', J') was a densest submodule of M' , we know that $\mathbf{w} \cdot J$ is a densest submodule of the module \widetilde{M} . Recall that we proved that \widetilde{M} is a $\gamma_{\text{uSVP}}/(2\sqrt{2})\text{-mod-uSVP}_2$ instance, hence we have

$$\mathcal{N}(\mathbf{w} \cdot J)^{\frac{1}{d}} = \lambda_1^{\mathcal{N}(\widetilde{M})} \leq \frac{2\sqrt{2}}{\gamma_{\text{uSVP}}} \cdot \mathcal{N}(\widetilde{M})^{\frac{1}{2d}}.$$

From the special shape of \mathbf{Y} , Lemma V.3.6 implies that $\mathbf{Y}^{-1} = (1/R) \cdot (\mathbf{I} + \mathbf{E}')$ where $\mathbf{E}' = (e'_{i,j})$ satisfies $\max_{i,j} \|e'_{i,j}\| \leq 3/5$ and $R = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)}$. It then holds that for any

¹J. Constante vérifiée et mise à jour.

Pour avoir cette constante, on applique d'abord l'inégalité triangulaire avant Lemmas V.3.6 and D.4.2, et ça nous donne $1 + 2 \cdot 3/5 = 2.2$.

embedding $\sigma_i : K_{\mathbb{R}} \rightarrow \mathbb{C}$ (for $1 \leq i \leq d$),

$$\begin{aligned} |\sigma_i(\|(\mathbf{I} + \mathbf{E}') \cdot \mathbf{w}\|_{K_{\mathbb{R}}})| &\leq \|(\mathbf{I}_2 + |\sigma_i(\mathbf{E}')|) \cdot \sigma_i(\mathbf{w})\|_2 \\ &\leq 2 \cdot \|\mathbf{I} + |\sigma_i(\mathbf{E}')|\|_{\infty} \cdot \|\sigma_i(\mathbf{w})\| \\ &\leq 2 \cdot (1 + 3/5) \cdot \|\sigma_i(\mathbf{w})\| \\ &\leq 4 \cdot \|\sigma_i(\mathbf{w})\|. \end{aligned}$$

Multiplying the previous inequality for $i = 1, \dots, d$ and dividing it by R^d gives

$$\mathcal{N}(\mathbf{Y}^{-1} \cdot \mathbf{w}) \leq 4^d \cdot R^{-d} \cdot \mathcal{N}(\mathbf{w}).$$

Hence, we obtain

$$\mathcal{N}(\mathbf{v} \cdot J)^{\frac{1}{d}} \leq 4 \cdot \frac{2\sqrt{2}}{R \cdot \gamma_{\text{uSVP}}} \cdot \mathcal{N}(\widetilde{M})^{\frac{1}{2d}} \leq \frac{16}{\gamma_{\text{uSVP}}} \cdot \mathcal{N}(M)^{\frac{1}{2d}},$$

where we used the definition of R and the fact that $\mathcal{N}(\widetilde{M}) \leq 2^d \cdot V$ (by Lemma V.4.4). Since $\gamma_{\text{uSVP}} > 16$, we conclude that $\mathbf{v} \cdot J$ is a rank-1 submodule of M with $\mathcal{N}(\mathbf{v} \cdot J) < \mathcal{N}(M)^{1/2}$. From the fact that the densest module is unique, we conclude that $\mathbf{v} \cdot J$ is indeed the densest submodule of M . \square

We used the following norm inequalities for matrix in $K_{\mathbb{R}}$.

Lemma D.4.2. *Let $k \geq 1$, for any matrix $\mathbf{A} = (a_{i,j})_{1 \leq i,j \leq k} \in K_{\mathbb{R}}^{k \times k}$ and $\mathbf{x} \in K_{\mathbb{R}}^k$, we have*

$$\|\mathbf{A} \cdot \mathbf{x}\| \leq k \cdot \max_{i,j} \|a_{i,j}\| \cdot \|\mathbf{x}\|.$$

Proof. By abuse of notation, we let $\Phi(\mathbf{A})$ denote the block-diagonal matrix in $\mathbb{C}^{dk \times dk}$ where the block (i, j) is $\text{diag}(\Phi(a_{i,j}))$ for $1 \leq i, j \leq k$. For any $n \geq 1$ and $\mathbf{M} \in \mathbb{C}^{n \times n}$ we let $\|\mathbf{M}\|_F = (\sum_{i,j} M_{i,j}^2)^{1/2}$ denote the Frobenius norm and $\|\|\mathbf{M}\|\|$ the operator norm $\sup_{\|x\|=1} \|\mathbf{M} \cdot x\|$ of \mathbf{M} . It holds that $\|\|\mathbf{M}\|\| \leq \|\mathbf{M}\|_F$. We have

$$\begin{aligned} \|\mathbf{A} \cdot \mathbf{x}\| &\leq \|\|\Phi(\mathbf{A})\|\| \cdot \|\mathbf{x}\| \\ &\leq \|\Phi(\mathbf{A})\|_F \cdot \|\mathbf{x}\| = \left(\sum_{1 \leq i,j \leq k} \|a_{i,j}\|^2 \right)^{\frac{1}{2}} \cdot \|\mathbf{x}\| \\ &\leq \sqrt{k^2} \cdot \left(\max_{i,j} \|a_{i,j}\|^2 \right)^{\frac{1}{2}} \cdot \|\mathbf{x}\| \\ &= k \cdot \max_{i,j} \|a_{i,j}\| \cdot \|\mathbf{x}\|. \end{aligned}$$

\square

D.4.2 Proof of Lemma V.4.4

The algorithm PreCond is as follows.

Proof. Let P be the polynomial from Lemma V.2.2 and define

$$V_0^{\frac{1}{2d}} = 10\sqrt{10} \cdot d \cdot \gamma(M) \cdot \left(P(\Delta_K^{1/d}, 2, d, 5\sqrt{2d}, 4d^{3/2} \cdot \gamma(M)^2 \cdot \delta_K \cdot \Delta_K^{1/(2d)}) + (2d)^{3/2} \right).$$

Algorithm D.4.1 Algorithm PreCond

Input: A pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 module $M \subseteq K^2$, two parameters $V > 0, B \geq 2$ and a block-size $\beta \in [2, 2d]$

Output: A matrix $\mathbf{Y} \in \text{GL}_2(K)$

- 1: Set $\varsigma = V^{1/(2d)} \cdot (5\sqrt{10d})^{-1} \cdot \mathcal{N}(M)^{-1/(2d)}$
- 2: **repeat**
- 3: Sample $\mathbf{Y} := (\mathbf{y}_1, \mathbf{y}_2)^T \leftarrow \text{DualRound}((\mathbf{B}, \mathbb{I}), \varsigma, \beta, 1/5)$
- 4: **until** $\mathbf{y}_1 \cdot \mathcal{O}_K$ is a primitive submodule of M^\vee
- 5: **Return** \mathbf{Y}

We will prove that the lemma holds for this choice of V_0 . Note that $V_0^{1/(2d)}$ is indeed equal to $\text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma(M))$ as desired.

Let us first observe that, by using the lower bound on V , the definition of ς and V_0 and Lemma V.2.5, one can prove that the lower bound $\varsigma \geq (2d)^{2d/\beta+3/2} \cdot \lambda_{2d}(M^\vee)$ required in Lemma V.3.5 is satisfied.

Applying Lemma V.3.5, we know that the calls to Algorithm DualRound will take a time polynomial in the input bitsize and in 2^β . To estimate the number of such calls, let us use Lemma V.2.2. The definition of V_0 ensures that the condition of Lemma V.2.2, namely $\varsigma \geq \lambda_{2d}(M^\vee) \cdot P(\Delta_K^{1/d}, 2, d, \|\mathbf{t}_1\|/\varsigma, \lambda_{kd}(M^\vee)/\lambda_1(M^\vee))$ are met. Hence, we know that $\Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{t}_1}}(\mathbf{y} \cdot \mathcal{O}_K \text{ is primitive in } M^\vee) \geq 1/(4\zeta_K(2))$. Using the fact that $\text{SD}(D_{M^\vee, \varsigma, \mathbf{t}_1}, \tilde{D}_{M^\vee, \varsigma, \mathbf{t}_1}) \leq 4^{-d}$ by Lemma V.3.5 and that for any $d \geq 2$ we have $4^{-d} \leq (9/10) \cdot (4\zeta_K(2))^{-1}$ we conclude that the probability to exit the while loop is at least $1/(4\zeta_K(2)) - 4^{-d} \geq \zeta_K(2)/40$ at every iteration of the algorithm. This proves the expected running time of the algorithm.

The fact that $\mathbf{Y} = R \cdot \mathbf{I}_2 + \mathbf{E}$ with $\|e_{ij}\| \leq R/5$ and that $M' := \mathbf{Y}\mathbf{B} \cdot \mathbb{I}$ is included in \mathcal{O}_K^2 follows from Lemma V.3.5 (instantiated with $\varepsilon = 1/5$). Since $\varepsilon = 1/5 \leq 1/4$, we can also use Lemma V.3.6. This implies in particular that $\det(\mathbf{Y}) \in [1/2^d, 2^d] \cdot R^{2d}$, where $R = 5 \cdot \sqrt{10d} \cdot \varsigma = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)}$ by definition of ς . This proves that \mathbf{Y} is invertible, and so M' is indeed a rank-2 module. This also proves that $\mathcal{N}(M') = \det(\mathbf{Y}) \cdot \mathcal{N}(M) \in [1/2^d, 2^d] \cdot V$.

Let us now show that if (\mathbf{B}, \mathbb{I}) was a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance, then $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a $\gamma'_{\text{uSVP-mod-uSVP}_2}$ instance. Let $\mathbf{s} \in M$ be a short vector such that $\|\mathbf{s}\| \leq 1/\gamma_{\text{uSVP}} \cdot \mathcal{N}(M)^{1/(2d)}$ (such a short vector exists if (\mathbf{B}, \mathbb{I}) is a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance). Define $\mathbf{s}' = \mathbf{Y} \cdot \mathbf{s}$, which is a vector of M' . We have

$$\|\mathbf{s}'\| \leq R \cdot \|\mathbf{s}\| + \|\mathbf{E} \cdot \mathbf{s}\| \leq 2R \cdot \|\mathbf{s}\| \leq 2R \cdot \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M)^{1/(2d)}.$$

Recall that $\mathcal{N}(M') = \det(\mathbf{Y}) \cdot \mathcal{N}(M) \geq 1/2^d \cdot R^{2d} \cdot \mathcal{N}(M)$. This finally implies that $\|\mathbf{s}'\| \leq 2\sqrt{2} \cdot \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M')^{1/(2d)}$, and so $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is indeed a $\gamma'_{\text{uSVP-mod-uSVP}_2}$ instance.

It finally remains to show that the module M' has the coprime property. This is implied by the fact that $\mathbf{y}_1 \cdot \mathcal{O}_K$ is primitive in M^\vee . Indeed, by definition of M' , we have that $\{x \in \mathcal{O}_K \mid \exists y \in \mathcal{O}_K \text{ s.t. } (x, y)^T \in M'\} = \{(\mathbf{y}_1, \mathbf{z})_{K_{\mathbb{R}}} \mid \mathbf{z} \in M\}$. One can see from the definition that this set is an ideal of \mathcal{O}_K . Assume by contradiction that it is not equal to \mathcal{O}_K and let \mathfrak{p} be a prime ideal dividing it. Then it holds that $\mathbf{y}_1 \cdot \mathfrak{p}^{-1} \subset M^\vee$. But this is a rank-1 submodule of M^\vee containing strictly the rank-1 module $\mathbf{y}_1 \cdot \mathcal{O}_K$, contradicting the assumption that $\mathbf{y}_1 \cdot \mathcal{O}_K$ is primitive in M^\vee . Hence, we conclude that M' has the coprime property. \square

D.4.3 Proof of Lemma V.4.5

The algorithm BalanceIdeal is as follows.

Algorithm D.4.2 Algorithm `BalanceIdeal`**Input:** A \mathbb{Z} -basis of a fractional ideal $I \subset K$ and a parameter $\gamma_{\text{HSVP}} \geq 1$ **Output:** An element $x \in K$ \triangleright Using a γ_{HSVP} -id-HSVP oracle to get short linearly independent vectors of I^{-1}

- 1: Call a γ_{HSVP} -id-HSVP solver on I^{-1} to get $y \in I^{-1}$
- 2: Let $\mathbf{B} = (y \cdot b_1^{\mathcal{O}_K}, \dots, y \cdot b_d^{\mathcal{O}_K})$ (where $(b_1^{\mathcal{O}_K}, \dots, b_d^{\mathcal{O}_K})$ is a \mathbb{Z} -basis of \mathcal{O}_K . This is a \mathbb{Z} -basis of $\langle y \rangle$)

 \triangleright Using the short vectors to find a balanced element in I^{-1} by solving CVP

- 3: Let $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(I)^{-1/d}$ and $t = (\sigma, \dots, \sigma)$
- 4: Write $t = \sum_i t_i \cdot y \cdot b_i^{\mathcal{O}_K}$, with $t_i \in \mathbb{R}$
- 5: Define $s = \sum_i \lfloor t_i \rfloor \cdot y \cdot b_i^{\mathcal{O}_K}$
- 6: **Return** $x = s^{-1}$

Proof. One can check that all the steps of the algorithm, except for the one call to the γ_{HSVP} -HSVP oracle, can be performed in polynomial time.

Let us then prove correction, and start with $I \subseteq \langle x \rangle$. We know that $s \in \langle y \rangle$, by definition of s . Since $y \in I^{-1}$, it holds that $\langle s \rangle \subseteq I^{-1}$, which implies $I \subseteq \langle s \rangle^{-1} = \langle x \rangle$ as desired (provided that $s \neq 0$, which we will show below).

Let us now look at how balanced are the coordinates of s (and x). We have

$$\begin{aligned}
\|s - t\|_\infty &\leq \sum_i 1/2 \cdot \|y \cdot b_i^{\mathcal{O}_K}\|_\infty \\
&\leq 1/2 \cdot \sum_i \|y\|_\infty \cdot \|b_i^{\mathcal{O}_K}\|_\infty \\
&\leq d/2 \cdot \|y\| \cdot \delta_K \\
&\leq 1/(2d) \cdot \sigma,
\end{aligned}$$

where we used in the last inequality the fact that y is the output of the γ_{HSVP} -id-HSVP solver on I^{-1} , and hence $\|y\| \leq \gamma_{\text{HSVP}} \cdot \mathcal{N}(I)^{-1/d}$. Since $\sigma_i(s)$ is the i -th coordinate of s and all the coordinates of t are equal to σ , this implies that $|\sigma_i(s)| \in [\sigma \cdot (1 - 1/(2d)), \sigma \cdot (1 + 1/(2d))]$ (and in particular $\sigma_i(s) \neq 0$ for all i 's, so s is invertible). Using the facts that $\sigma_i(x) = \sigma_i(s)^{-1}$ and the convexity of the function $x \mapsto 1/x$ over $[1/2, 2]$ conclude the proof. \square

D.4.4 Proof of Lemma V.4.6

We know from preliminaries (cf Section II.3) that the HNF basis of a module can be computed in polynomial time. From Lemma V.4.5 we know that the algorithm `BalanceIdeal` runs in polynomial time and make one call to the γ_{HSVP} oracle. Note that the input ideal J_2 is indeed fractional (and even integral) since $M \subset \mathcal{O}_K^2$ and that $\gamma_{\text{HSVP}} \geq \sqrt{d} \Delta_K^{1/(2d)}$ hence we can indeed run algorithm `BalanceIdeal`. Finally, the multiplications and rounding in the third step of the algorithm can be performed in polynomial time too. \square

D.4.5 Proof of Lemma V.4.7

Let us fix some $\delta, \gamma_{\text{HSVP}}, \gamma_{\text{NTRU}}$ and q as in the theorem and define V and γ_{uSVP} accordingly.

Let (\mathbf{B}, \mathbb{I}) be the input pseudo-basis, spanning a rank-2 module $M_1 \subset \mathcal{O}_K^2$ with $\mathcal{N}(M_1) \in [1/2^{2d}, 2^{2d}] \cdot V$, with the coprime property, and which we know contains a non-zero vector $\mathbf{s}_1 = (u, v)^T \in M_1$ such that $\|\mathbf{s}_1\| \leq 1/\gamma_{\text{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. We will see step by step how the

module M_1 is modified by the algorithm, and what happens to its short non-zero vectors. This is summarized on Figure V.2.

First step: HNF. After the HNF computation, we have a new pseudo-basis of the form

$$\begin{bmatrix} J_1 & J_2 \\ \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix} \end{bmatrix}$$

for some $a \in K$ and $J_1, J_2 \subset K$ (cf Section II.3). This pseudo-basis generates a rank-2 module M_2 which is the same as the input module M_1 . Hence, M_2 contains a short non-zero vector $\mathbf{s}_2 := \mathbf{s}_1$.

Since our module M_2 is integral, we know that both ideals J_1 and J_2 are integral. Also, since module $M_2 = M_1$ has the coprime property, we know that $J_1 = \mathcal{O}_K$. Finally, because of the shape of the pseudo-basis, it holds that $\mathcal{N}(J_1) \cdot \mathcal{N}(J_2) = \mathcal{N}(M_2) = \mathcal{N}(M_1)$, which yields $\mathcal{N}(J_2) \geq \mathcal{N}(M_1)$.

Second step: from pseudo-basis to basis. Let M_3 be the free module generated by the pseudo-basis

$$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ a & b \end{pmatrix} \end{bmatrix},$$

where $b \leftarrow \text{BalanceIdeal}(J_2, \gamma_{\text{HSVP}})$. Since $J_2 \subseteq \langle b \rangle$ by Lemma V.4.5 and $J_1 = \mathcal{O}_K$, we conclude that $M_2 \subseteq M_3$. Hence, the short vector $\mathbf{s}_3 := \mathbf{s}_2$ is still in M_3 .

Before moving to the next step, let us have a closer look at b . We know from Lemma V.4.5 that $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$ for all $i \leq d$, with $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_2)^{-1/d}$. Using the lower bound on $\mathcal{N}(J_2)$ that we computed above, this shows that $\sigma \leq \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(M_1)^{-1/d}$.

Third step: transforming b into q . Let M_4 be the free module generated by the pseudo-basis

$$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix} \end{bmatrix},$$

where $h = \lfloor a \cdot q/b \rfloor$. This is the basis output by Algorithm Conditioned-to-NTRU. The new module M_4 does not contain M_1 anymore, however we will show that its geometry is close to the one of M_3 , so that it has a short non-zero vector if M_3 does.

Recall that M_3 contains a short vector $\mathbf{s}_3 = (u, v)^T$ such that $\|\mathbf{s}_3\| \leq 1/\gamma_{\text{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. Let $x \in \mathcal{O}_K$ be such that $\mathbf{s}_3 = u \cdot (1, a)^T + x \cdot (0, b)^T$. Define $\mathbf{s}_4 = u \cdot (1, h)^T + x \cdot (0, q)^T \in M_4 \setminus \{\mathbf{0}\}$. Unrolling the definition of h and using the equation $u \cdot a + x \cdot b = v$, one can rewrite $\mathbf{s}_4 = (u, v \cdot q/b - u \cdot \{a \cdot q/b\})^T$. We can upper bound the euclidean norm of \mathbf{s}_4 as follows

$$\begin{aligned} \|\mathbf{s}_4\| &\leq \|u\| + \|v\| \cdot \|q/b\|_\infty + \|u\| \cdot \|\{a \cdot q/b\}\|_\infty \\ &\leq \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot (1 + q \cdot 2\sigma + d\delta_K) \\ &\leq \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot (q \cdot 2\sigma + 2 \cdot d\delta_K) \\ &\leq 2\gamma_{\text{uSVP}}^{-1} \cdot d \cdot \delta_K \cdot (\mathcal{N}(M_1)^{1/(2d)} + q \cdot d \cdot \gamma_{\text{HSVP}} \cdot \mathcal{N}(M_1)^{-1/(2d)}) \\ &\leq 1/\gamma_{\text{NTRU}} \cdot \sqrt{q}, \end{aligned}$$

where in the last step we used the fact that $\mathcal{N}(M_1)^{1/(2d)} \in [1/2, 2] \cdot V^{1/(2d)}$ and the definitions of V and γ_{uSVP} . We conclude that the pseudo-basis output by Algorithm Conditioned-to-NTRU is indeed a γ_{NTRU} -NTRU instance, as desired. \square

D.4.6 Proof of Lemma V.4.8

Algorithm `LiftMod` is as follows.

Algorithm D.4.3 Algorithm `LiftMod`

Input: Two elements $a, b \in K$, an NTRU instance $((\mathbf{c}_1, \mathbf{c}_2), \mathcal{O}_K^2)$ and a pseudo-basis (\mathbf{v}, J) of a rank-1 module in K^2 .

Output: A vector $\mathbf{w} \in K^2$

- 1: Compute $x, y \in K$ such that $\mathbf{v} = x \cdot \mathbf{c}_1 + y \cdot \mathbf{c}_2$
 - 2: Define $\mathbf{w} = x \cdot (1, a)^T + y \cdot (0, b)^T$
 - 3: **Return** \mathbf{w}
-

Proof. The running time follows from inspection of the algorithm.

Let \mathbf{s}_1 be a shortest vector of M_1 . Since $\gamma_{\text{uSVP}} > 1$, we know from Lemma II.3.4 that \mathbf{s}_1 belongs to the densest rank-1 submodule of M_1 , i.e., the densest submodule of M_1 is equal to $\text{span}_K(\mathbf{s}_1) \cap M_1$.

We use the notations M_1, M_2, M_3 and M_4 as in Figure V.2 (and the proof of Lemma V.4.7). Recall from the proof of Lemma V.4.7 that \mathbf{s}_1 is still a vector of the rank-2 module M_3 spanned by $(1, a)^T, (0, b)^T$. Let $u, r \in \mathcal{O}_K$ be such that $\mathbf{s}_1 = u \cdot (1, a)^T + r \cdot (0, b)^T$. Recall again from the proof of Lemma V.4.7 that $\mathbf{s}_4 = u \cdot \mathbf{c}_1 + r \cdot \mathbf{c}_2$ is an unexpectedly short vector of the output NTRU module M_4 . More precisely, we proved that $\|\mathbf{s}_4\| \leq 1/\gamma_{\text{NTRU}} \cdot \mathcal{N}(M_4)^{1/(2d)}$.

Using Lemma II.3.4 again and the fact that $\gamma_{\text{NTRU}} > 1$, we know that \mathbf{s}_4 belongs to the densest submodule of M_4 . Since (\mathbf{v}, J) is a pseudo-basis of this densest submodule, it should be that \mathbf{v} and \mathbf{s}_4 are K -collinear, i.e., there exists $z \in K$ such that $\mathbf{v} = z \cdot \mathbf{s}_4 = zu \cdot \mathbf{c}_1 + zr \cdot \mathbf{c}_2$.

Hence, the elements x, y computed in the algorithms are equal to zu and zr respectively. This proves that $\mathbf{w} = x \cdot (1, a)^T + y \cdot (0, b)^T = z \cdot \mathbf{s}_1$. Hence, $\text{span}_K(\mathbf{w}) = \text{span}_K(\mathbf{s}_1)$ and the densest submodule of M is $\text{span}_K(\mathbf{w}) \cap M_1$. \square

D.4.7 Proof of Lemma V.4.9

Algorithm `LiftVec` is as follows.

Algorithm D.4.4 Algorithm `LiftVec`

Input: Some auxiliary information $\text{aux} = (a, b, J_1, J_2)$, a parameter γ_{HSVP} , an NTRU instance $((\mathbf{c}_1, \mathbf{c}_2), \mathcal{O}_K^2)$ and a vector $\mathbf{s} \in \mathcal{C} \cdot \mathcal{O}_K^2$

Output: A vector $\mathbf{w} \in K^2$

- 1: Compute $x, y \in \mathcal{O}_K$ such that $\mathbf{s} = x \cdot \mathbf{c}_1 + y \cdot \mathbf{c}_2$
 - 2: Run $z \leftarrow \text{BalanceIdeal}(\langle b \rangle \cdot J_1^{-1} \cdot J_2^{-1}, \gamma_{\text{HSVP}})$
 - 3: Compute $\mathbf{t} = z^{-1} \cdot (x \cdot (1, a)^T + y \cdot (0, b)^T)$
 - 4: **Return** \mathbf{t}
-

Proof. The running-time follows from inspection of the algorithm and from Lemma V.4.5.

Let us show that the output \mathbf{t} of the algorithm is indeed in the module M_1 . Let us keep the notations M_1, M_2, M_3 and M_4 from Figure V.2. In particular, $M_1 = M_2$ is the module generated by the pseudo-basis $((1, a)^T, (0, b)^T), (J_1, \langle b^{-1} \rangle \cdot J_2)$.

From Lemma V.4.5, we know that $\langle b \rangle \cdot J_1^{-1} \cdot J_2^{-1} \subseteq \langle z \rangle$, i.e., $z^{-1} \in J_1 \cdot J_2 \cdot \langle b^{-1} \rangle$. Using the fact that J_1 and $J_2 \cdot \langle b^{-1} \rangle$ are both integral (recall that $J_2 \subseteq \langle b \rangle$ and that M_1 is in \mathcal{O}_K^2), this implies

that $z^{-1} \in J_1 \cap J_2 \cdot \langle b^{-1} \rangle$. Since $x, y \in \mathcal{O}_K$, we conclude that $\mathbf{t} = z^{-1} \cdot x \cdot (1, a)^T + z^{-1} \cdot y \cdot (0, b)^T$ is in M_1 as desired.

Let us now upper bound the size of \mathbf{t} . Let us write $\mathbf{s} = (s_1, s_2)^T$ and express the coordinates of \mathbf{t} in terms of s_1 and s_2 . From the equation $\mathbf{s} = x \cdot (1, \lfloor a \cdot q/b \rfloor)^T + y \cdot (0, q)^T$, we obtain $x = s_1$ and $s_2 = x \lfloor a \cdot q/b \rfloor + yq$. This implies that $\mathbf{t} = z^{-1} \cdot (s_1, b/q \cdot (s_2 + s_1 \cdot \{a \cdot q/b\}))^T$. From this, we can upper bound

$$\|\mathbf{t}\| \leq \|z^{-1}\|_\infty \cdot \|\mathbf{s}\| \cdot (1 + \|b/q\|_\infty \cdot (1 + d\delta_K)).$$

Recall from the proof of Lemma V.4.7 that for all $i \leq d$, we have $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$, with $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_2)^{-1/d}$. Hence,

$$\|b/q\|_\infty \leq \frac{2 \cdot \mathcal{N}(J_2)^{1/d}}{\gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot q} \quad \text{and} \quad |\mathcal{N}(b)| \geq \frac{\mathcal{N}(J_2)}{(2\gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K)^d}.$$

From Lemma V.4.5, we similarly know that $|\sigma_i(z)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma_z^{-1}$ for all $i \leq d$, where $\sigma_z = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d}$. Hence we obtain

$$\|z^{-1}\|_\infty \leq 2 \cdot \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d} \leq 4 \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2,$$

where we used the fact that $J_1 = \mathcal{O}_K$ thanks to the coprime property of M_1 .

Finally, recall that $\mathcal{N}(J_2) = \mathcal{N}(M_1) \leq 2^{2d} \cdot V$, with $V^{1/d} = q \cdot \gamma_{\text{HSVP}} \cdot d$. Combining everything provides the desired upper bound on $\|\mathbf{t}\|$. \square

D.5 Removing $\zeta_K(2)$ from Theorem V.4.1

D.5.1 Tail-cut of $\zeta_K(2)$ and B -coprime property

In order to remove the dependence in $\zeta_K(2)$ in the running-time of Theorem V.4.1, we will relax the coprime property.

Definition D.5.1 (B -coprime property). *We say that a rank-2 module $M \subseteq \mathcal{O}_K^2$ has the B -coprime property if it holds that*

$$\mathcal{N}(\{x \in \mathcal{O}_K \mid \exists y \in \mathcal{O}_K, (x, y)^T \in M\}) \leq B.$$

In other words, the module M has the coprime property if the ideal spanned by the first coordinate of all the vectors of M has norm at most B .

To decide which B to use, we have to give bounds on the tail of the defining serie of $\zeta_K(2)$.

Lemma D.5.2. *There exists an absolute constant $\kappa > 1$ such that*

$$\sum_{\mathcal{N}(\mathbf{a}) > B} \frac{1}{\mathcal{N}(\mathbf{a})^2} \leq \frac{(\Delta_K \cdot d^d)^\kappa}{B}$$

for every $B \geq 2^d$.

Proof. We have, that for any $\alpha > 1$ and $n \geq 1$,

$$\frac{1}{(n+1)^\alpha} \leq \int_n^{n+1} \frac{dx}{x^\alpha} \leq \frac{1}{n^\alpha},$$

and then, as $\alpha > 1$ we have for any $a \geq 1$

$$\sum_{n>a} \frac{1}{n^\alpha} = \sum_{n \geq a} \frac{1}{(n+1)^\alpha} \leq \int_a^\infty \frac{dx}{x^\alpha} = \frac{1}{(\alpha-1) \cdot a^{\alpha-1}}.$$

Let a_n be the number of ideals of \mathcal{O}_K with norm exactly n . Note that $N_K(x) = \sum_{k \leq x} a_k$. Note also that $(2n+1)/(n^2(n+1)^2) \leq 3/n^3$ for any $n \geq 1$. Let $2^d < a < b$ two integers; it holds that

$$\begin{aligned} \sum_{a < k \leq b} \frac{a_k}{k^2} &= \sum_{a < k \leq b} \frac{N_K(k) - N_K(k-1)}{k^2} \\ &= \sum_{a < k \leq b} \frac{N_K(k)}{k^2} - \sum_{a-1 < k \leq b-1} \frac{N_K(k)}{(k+1)^2} \\ &= \sum_{a < k \leq b-1} N_K(k) \cdot \frac{2k+1}{k^2 \cdot (k+1)^2} + \frac{N_K(b)}{b^2} - \frac{N_K(a)}{(a+1)^2} \\ &\leq \sum_{a < k \leq b-1} \frac{3 \cdot N_K(k)}{k^3} + \frac{N_K(b)}{b^2}. \end{aligned}$$

The asymptotic growth of $N_K(b)$ as $\rho_K \cdot b$ when b goes to infinity implies that (with $\eta = 1/(16 \ln(d))$)

$$\begin{aligned} \sum_{k>a} \frac{a_k}{k^2} &\leq \sum_{k>a} \frac{3 \cdot N_K(k)}{k^3} \\ &\leq 3\rho_K \cdot \sum_{k>a} \frac{1}{k^2} + 3 \cdot M'(K) \cdot \sum_{k>a} \frac{1}{k^{2+\eta}} \quad (\text{Theorem III.1.2}) \\ &\leq \frac{3\rho_K}{a} + \frac{3 \cdot M'(K)}{(1+\eta) \cdot a^{1+\eta}} \\ &\leq \frac{3\rho_K + 3M'(K)}{a}. \end{aligned}$$

By Theorem II.2.10, we have that $\rho_K \leq (e \cdot \log(\Delta_K)/(2(d-1)))^{d-1}$, so there exists $\kappa \geq 1$ such that $3(M'(K) + \rho_K) \leq (\Delta_K \cdot d^d)^\kappa$. \square

We make use of the following lemmas.

Lemma D.5.3. *Let $M \subset K_{\mathbb{R}}^k$ be a rank- k module and let $\mathbf{y}_1 \in M^\vee$. We define*

$$\mathbf{y}_1 \cdot M := \{ \langle \mathbf{y}_1, \mathbf{m} \rangle_{K_{\mathbb{R}}}, \mathbf{m} \in M \} \subseteq \mathcal{O}_K.$$

It holds that $\mathbf{y}_1 \cdot M$ is an integral ideal of K . Furthermore, for any integral ideal \mathfrak{a} ,

$$\mathbf{y}_1 \cdot M = \mathfrak{a} \Leftrightarrow \mathbf{y}_1 \cdot \mathcal{O}_K \text{ primitive in } \mathfrak{a} \cdot M^\vee.$$

Proof. The fact that $\mathbf{y}_1 \cdot M$ is an integral ideal comes from its definition. Assume that $\mathbf{y}_1 \cdot \mathcal{O}_K$ is primitive in $\mathfrak{a} \cdot M^\vee$. The fact that $\mathbf{y}_1 \in \mathfrak{a} M^\vee$ implies that $\mathbf{y}_1 \cdot M \subseteq \mathfrak{a}$. Now assume that $\mathbf{y}_1 \cdot M = \mathfrak{b} \cdot \mathfrak{a}$ for some integral ideal \mathfrak{b} . Then if $\mathfrak{b} \neq \mathcal{O}_K$, it holds that $\mathbf{y}_1 \cdot \mathfrak{b}^{-1} \subset \mathfrak{a} M^\vee$, contradicting the primitivity of $\mathbf{y}_1 \cdot \mathcal{O}_K$.

Conversely, assume that $\mathbf{y}_1 \cdot M = \mathfrak{a}$. This implies that $\mathbf{y}_1 \cdot \mathcal{O}_K \subset \mathfrak{a} \cdot M^\vee$. Assume that $\mathbf{y}_1 \cdot \mathcal{O}_K$ is not primitive in $\mathfrak{a} \cdot M^\vee$, i.e., there exists an integral ideal $\mathfrak{b} \neq \mathcal{O}_K$ such that $\mathbf{y}_1 \cdot \mathfrak{b}^{-1} \subset \mathfrak{a} \cdot M^\vee$, i.e., $\mathbf{y}_1 \cdot \mathcal{O}_K \subset \mathfrak{a} \mathfrak{b} \cdot M^\vee$. In particular, it holds that $\mathfrak{a} = \mathbf{y}_1 \cdot M \subset \mathfrak{a} \cdot \mathfrak{b}$ which contradicts $\mathfrak{b} \neq \mathcal{O}_K$, hence the result. \square

We can now give a bound on the probability to be B -coprime for a random module.

Lemma D.5.4. *Let κ as in Lemma D.5.2. There exists an absolute polynomial \tilde{P} such that the following holds. For any degree- d number field K , real $B \geq (\Delta_K \cdot d^d)^\kappa$, $\delta \geq 0$ and rank-2 module $M \subset K_{\mathbb{R}}^2$, if $\mathbf{c} \in \text{span}_{K_{\mathbb{R}}}(M^\vee)$ and $\varsigma > 0$ are such that $\|\mathbf{c}\| \leq \delta \cdot \varsigma$ and*

$$\varsigma \geq \lambda_{kd}(M^\vee) \cdot \tilde{P}\left(B^{1/d}, \delta, \frac{\lambda_{2d}(M^\vee)}{\lambda_1(M^\vee)}\right),$$

then it holds that

$$\Pr_{[\mathbf{y}_1, \mathbf{y}_2]^T \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}^2}([\mathbf{y}_1, \mathbf{y}_2]^T \cdot M \text{ has the } B\text{-Coprime property}) \geq \frac{1}{10}.$$

Proof. Let κ be the same as defined in Lemma D.5.2, B as in the statement (note that $B \geq 2^d$), \mathfrak{a} an integral ideal of norm less than B and

$$\varsigma \geq \sqrt{d} \Delta_K^{1/(2d)} \cdot B^{1/d} \cdot \lambda_{2d}(M^\vee) \cdot P\left(\Delta_K^{1/d}, 2, d, \delta, \frac{B^{1/d} \cdot \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \lambda_{2d}(M^\vee)}{\lambda_1(M^\vee)}\right),$$

where P is the polynomial defined in Lemma V.2.2. It holds that $\lambda_1(\mathfrak{a} \cdot M^\vee) \geq \lambda_1(M^\vee)$ and that

$$\lambda_{2d}(\mathfrak{a} \cdot M^\vee) \leq \lambda_1^{(\infty)}(\mathfrak{a}) \cdot \lambda_{2d}(M^\vee) \leq \sqrt{d} \cdot \mathcal{N}(\mathfrak{a})^{\frac{1}{d}} \cdot \Delta_K^{\frac{1}{2d}} \cdot \lambda_{2d}(M^\vee).$$

The standard deviation ς satisfies the hypothesis of Lemma V.2.2 for the module $\mathfrak{a} \cdot M^\vee$. Note that the definition of P also implies (See the proof of Lemma V.2.2) that $\varsigma \geq \eta_\varepsilon(\mathfrak{a} M^\vee) \geq \eta_\varepsilon(M^\vee)$ for $\varepsilon = 2^{-4d-5}$. This, along with Lemma II.1.6, give that

$$\Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \in \mathfrak{a} M^\vee) \geq \frac{1 - \varepsilon}{1 + \varepsilon} \cdot \frac{\text{Vol}(M^\vee)}{\text{Vol}(\mathfrak{a} M^\vee)} \geq \frac{9}{10 \cdot \mathcal{N}(\mathfrak{a})^2}.$$

For any event $E(\cdot)$ depending on some $\mathbf{y} \in M^\vee$, by the definition of the Gaussian distribution, it holds that

$$\Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(E(\mathbf{y}) \mid \mathbf{y} \in \mathfrak{a} M^\vee) = \Pr_{\mathbf{y} \leftarrow D_{\mathfrak{a} \cdot M^\vee, \varsigma, \mathbf{c}}}(E(\mathbf{y})).$$

Let $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2]^T \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}^2$ and $M' = \mathbf{Y} \cdot M$. The ideal spanned by first coordinates of the vectors of $\mathbf{Y} \cdot M$ is the ideal $\mathbf{y}_1 \cdot M$. By Lemma D.5.3, we have $\mathbf{y}_1 \cdot M = \mathfrak{a}$ if and only if \mathbf{y}_1 is primitive in $\mathfrak{a} \cdot M^\vee$, by Lemma V.2.2 it then holds that

$$\begin{aligned} & \Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \cdot M = \mathfrak{a}) \\ &= \Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \text{ primitive in } \mathfrak{a} \cdot M^\vee \mid \mathbf{y} \in \mathfrak{a} \cdot M^\vee) \cdot \Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \in \mathfrak{a} \cdot M^\vee) \\ &= \Pr_{\mathbf{y} \leftarrow D_{\mathfrak{a} \cdot M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \text{ primitive in } \mathfrak{a} \cdot M^\vee) \cdot \Pr_{\mathbf{y} \leftarrow D_{M^\vee, \varsigma, \mathbf{c}}}(\mathbf{y} \in \mathfrak{a} \cdot M^\vee) \\ &\geq \frac{1}{4 \cdot \zeta_K(2)} \cdot \frac{9}{10 \cdot \mathcal{N}(\mathfrak{a})^2} \\ &\geq \frac{1}{5 \cdot \zeta_K(2) \cdot \mathcal{N}(\mathfrak{a})^2}. \end{aligned}$$

The fact that the events $[\mathbf{y} \cdot M = \mathbf{a}]$ are pairwise exclusive for every \mathbf{a} implies that

$$\begin{aligned}
& \Pr_{\mathbf{Y}=[\mathbf{y}_1, \mathbf{y}_2]^T \leftarrow D_{M^\vee, \varsigma, \mathbf{e}}^2} (\mathbf{Y} \cdot M \text{ has the } B\text{-coprime property}) \\
&= \sum_{\substack{\mathcal{N}(\mathbf{a}) \leq B \\ \mathbf{a} \text{ integral}}} \Pr_{\mathbf{y}_1 \leftarrow D_{M^\vee, \varsigma, \mathbf{e}}} (\mathbf{y} \cdot M = \mathbf{a}) \\
&\geq \frac{1}{5\zeta_K(2)} \cdot \sum_{\substack{\mathcal{N}(\mathbf{a}) \leq B \\ \mathbf{a} \text{ integral}}} \frac{1}{\mathcal{N}(\mathbf{a})^2} \\
&= \frac{1}{5} \cdot \left(1 - \frac{1}{\zeta_K(2)} \cdot \sum_{\substack{\mathcal{N}(\mathbf{a}) > B \\ \mathbf{a} \text{ integral}}} \frac{1}{\mathcal{N}(\mathbf{a})^2} \right) \\
&\geq \frac{1}{5} \cdot \left(1 - \frac{(\Delta_K \cdot d^d)^\kappa}{B \cdot \zeta_K(2)} \right), \quad \text{by Lemma D.5.2.}
\end{aligned}$$

Note that $\zeta_K(2) \leq \zeta(2)^d$. Taking $B \geq (\Delta_K \cdot d^d)^\kappa \geq 2 \cdot (\Delta_K \cdot d^d)^\kappa / \zeta_K(2)$ gives the claimed result. \square

D.5.2 Proof of Theorem V.4.2

Theorem V.4.2 is a direct corollary of the following more complete statement (the blue background highlight the difference with the results of the previous sections).

Theorem D.5.5 (Updated Theorem D.4.1). *Let $\kappa > 1$ as defined in Lemma D.5.2, K a number field of degree d , let $\gamma^+ > 0$. There exist three algorithms `uSVP-to-NTRU`, `LiftVecInternal` and `LiftModInternal` (the same as in Theorem D.4.1) and $q_0 = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ such that the following holds. For any $q \geq q_0$, $\gamma_{\text{NTRU}} > 1$, $\gamma_{\text{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$ and (\mathbf{B}, \mathbb{I}) pseudo-basis of a rank-2 module $M \subset K^2$ with $\gamma(M) \leq \gamma^+$.*

- Algorithm `uSVP-to-NTRU` takes as input (\mathbf{B}, \mathbb{I}) , q and γ_{HSVP} and outputs $(\mathbf{B}', \mathcal{O}_K^2)$ a pseudo-basis of a rank 2 free module $M' \subset \mathcal{O}_K^2$, together with some auxiliary information \mathbf{aux} . If (\mathbf{B}, \mathbb{I}) is a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance with

$$\gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 16\sqrt{2} \cdot d^{\frac{3+\kappa}{2}} \cdot \Delta_K^{\frac{\kappa}{2d}} \cdot \delta_K,$$

then $(\mathbf{B}', \mathcal{O}_K^2)$ is a $(\gamma_{\text{NTRU}}, q)$ -NTRU instance. If given access to a $\gamma_{\text{HSVP-id-HSVP}}$ oracle, it makes one call to the oracle and runs in time polynomial in its input bitsize and in $\exp(\frac{d \log(d)}{\log(q/q_0)})$.

- Algorithm `LiftVecInternal` takes as input a non-zero vector $\mathbf{s}' \in M'$ and the auxiliary information \mathbf{aux} . It outputs a non-zero vector $\mathbf{s} \in M$ such that

$$\|\mathbf{s}\| \leq 150 \cdot \gamma_{\text{HSVP}}^{3/2} \cdot \Delta_K^{\frac{\kappa}{2d}} \cdot d^{\frac{7+3\kappa}{2}} \cdot \delta_K^2 \cdot \frac{\|\mathbf{s}'\|}{\mathcal{N}(M')^{\frac{1}{2d}}} \cdot \mathcal{N}(M)^{\frac{1}{2d}}.$$

If given access to a $\gamma_{\text{HSVP-id-HSVP}}$ oracle, it runs in polynomial time and makes one call to the oracle.

- *Algorithm LiftModInternal* takes as input a pseudo-basis of a rank-1 densest submodule N' of M' and the auxiliary information \mathbf{aux} and outputs a pseudo-basis of a rank-1 densest submodule N of M . It runs in polynomial time.

It should be noted that κ in Theorem D.5.5 is directly related to the decrease rate of the tail $\sum_{\mathcal{N}(\mathbf{a}) > B} \mathcal{N}(\mathbf{a})^{-2}$ when B goes to infinity. We state here the theorem in full generality (the constant κ applies to all number fields), but for some family of fields it might be possible prove that $\sum_{\mathcal{N}(\mathbf{a}) > B} \mathcal{N}(\mathbf{a})^{-2}$ converges to 0 quicker than $(\Delta_K \cdot d^d)^\kappa / B$, leading to a smaller κ in the theorem.

Proof of Theorem D.5.5. This proof is a modified version of the proof of Theorem D.4.1. We fix $B = (\Delta_K \cdot d^d)^\kappa$ for the rest of the proof.

Let $V_0 = (\text{poly}(B^{1/d}, \delta_K, \gamma^+))^{2d}$ be as in Lemma D.5.6 (defined using γ^+ instead of $\gamma(M)$). Define

$$q_0 = \frac{V_0^{1/d} \cdot 4d}{\gamma_{\text{HSVP}}}.$$

One can check that q_0 is indeed $\text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+)$ as desired. We prove that the theorem holds for this choice of q_0 .

Algorithm uSVP-to-NTRU. On input (\mathbf{B}, \mathbb{I}) , q and γ_{HSVP} , uSVP-to-NTRU sets $V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$ and $\beta = \lceil \frac{2d \log(2d)}{\log(\sqrt{q/q_0}) + \log(2d)} \rceil$. It then runs PreCond on input (\mathbf{B}, \mathbb{I}) , V and β , to obtain a matrix $\mathbf{Y} \in \text{GL}_2(K)$.

From the definition of q_0 , V and β , one can check that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Moreover, we have $\gamma(M) \leq \gamma^+$ by assumption, hence we can apply Lemma D.5.6. This implies in particular that the call to the PreCond algorithm runs in time polynomial in the input bitsize and in $2^\beta = 2^{\mathcal{O}(d \log(d) / \log(q/q_0))}$.

Algorithm uSVP-to-NTRU then runs Conditioned-to-NTRU on input $(\mathbf{Y}\mathbf{B}, \mathbb{I})$, q and γ_{HSVP} . It obtains a basis \mathbf{B}' of a free module M' and some auxiliary information \mathbf{aux}' . Algorithm uSVP-to-NTRU finally outputs $(\mathbf{B}', \mathcal{O}_K^2)$ and $\mathbf{aux} = (\mathbf{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$.

We know that the call to Conditioned-to-NTRU can be done in polynomial time, with one call to the γ_{HSVP} -id-HSVP oracle. This completes the proof on the running time of algorithm uSVP-to-NTRU.

Let us assume now that (\mathbf{B}, \mathbb{I}) was a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance, for γ_{uSVP} as in the theorem. We know from Lemma D.5.6 that $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a $\gamma_{\text{uSVP}}/(2\sqrt{2})$ -mod-uSVP₂ instance. Moreover, still from Lemma D.5.6, we know that the module spanned by $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a rank-2 module in \mathcal{O}_K^2 , with the B -coprime property and such that $\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$. Hence we can apply Lemma D.5.7 and conclude that $(\mathbf{B}', \mathcal{O}_K^2)$ is a γ_{NTRU} instance, as desired (note that V and $\gamma_{\text{uSVP}}/(2\sqrt{2})$ have the desired shape for applying Lemma D.5.7). This proves the first item of the theorem.

Algorithm LiftVecInternal. On input $\mathbf{s}' \in M'$ and $\mathbf{aux} = (\mathbf{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$, algorithm LiftVecInternal runs LiftVec(\mathbf{aux}' , γ_{HSVP} , \mathbf{B}' , \mathbf{s}') and gets a nonzero vector \mathbf{t} . It then outputs $\mathbf{Y}^{-1} \cdot \mathbf{t}$. By Lemma D.5.8, we know that the call to LiftVec can be performed in polynomial time, with one call to the id-HSVP oracle. This proves the running time of LiftVecInternal.

By Lemma D.5.8 again, we know that $\|\mathbf{t}\| \leq \|\mathbf{s}'\| \cdot 68 \cdot B^{3/(2d)} \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$. From the shape of \mathbf{Y} , Lemma V.3.6 instantiated with $\varepsilon = 1/5$, the value of B and Lemma D.4.2, we obtain²

$$\begin{aligned} \|\mathbf{Y}^{-1} \cdot \mathbf{t}\| &\leq \frac{2.2 \cdot \mathcal{N}(M)^{\frac{1}{2d}}}{V^{\frac{1}{2d}}} \cdot \|\mathbf{t}\| \\ &\leq 150 \cdot \gamma_{\text{HSVP}}^{\frac{3}{2}} \cdot \Delta_K^{\frac{\kappa}{2d}} \cdot d^{\frac{7+3\kappa}{2}} \cdot \delta_K^2 \cdot \frac{\|\mathbf{s}'\|}{\sqrt{q}} \cdot \mathcal{N}(M)^{\frac{1}{2d}}. \end{aligned}$$

Using the fact that $\mathcal{N}(M') = q^d$ provides the desired upper bound on the output size. Note also that by construction, $\mathbf{Y}^{-1} \cdot \mathbf{t}$ is indeed a non-zero vector in M .

Algorithm LiftModInternal. The proof is identical to the one of Theorem D.4.1. \square

D.5.3 Updating Lemma V.4.4

We can now state Algorithm D.5.1, which is the equivalent of Lemma V.4.4 for the algorithm `PreCond'`.

Lemma D.5.6 (Updated Lemma V.4.4). *Let $B \geq (\Delta_K \cdot d^d)^\kappa$. Let (\mathbf{B}, \mathbb{I}) be a pseudo-basis of a rank-2 module $M \subset K^2$ with gap $\gamma(M) \geq 1$. There exists some $V_0 > 0$ with $V_0^{1/(2d)} = \text{poly}(B^{1/d}, \delta_K, \gamma(M))$ and an algorithm `PreCond'` such that the following holds.*

Let $\beta \in \{2, \dots, 2d\}$ and $V > 0$ be such that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Then, on input (\mathbf{B}, \mathbb{I}) , V and β , algorithm `PreCond'` outputs a matrix $\mathbf{Y} \in \text{GL}_2(K)$ such that

- *if (\mathbf{B}, \mathbb{I}) is a $\gamma_{\text{uSVP-mod-uSVP}_2}$ instance, then $(\mathbf{Y}\mathbf{B}, \mathbb{I})$ is a $\gamma'_{\text{uSVP-mod-uSVP}_2}$ instance for $\gamma'_{\text{uSVP}} = \gamma_{\text{uSVP}}/(2\sqrt{2})$;*
- *the rank-2 module $M' := \mathbf{Y}\mathbf{B} \cdot \mathbb{I}$ is contained in \mathcal{O}_K^2 ;*
- *$\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$;*
- *M' has the B -coprime property;*
- *$\mathbf{Y} = R \cdot \mathbf{I}_2 + \mathbf{E}$ for some $R = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)} > 0$ and $\|e_{ij}\| \leq R/5$ for all $1 \leq i, j \leq 2$.*

Algorithm `PreCond'` runs in expected time polynomial in its input bitsize and in 2^β .

Proof. The proof is the same as the one of Lemma V.4.4, with the difference that V_0 has to be chosen so that ς matches the hypothesis of Lemma D.5.4 and that the expected running number of repeat event is ≤ 10 . \square

²J: Pour avoir cette constante, on applique d'abord l'inégalité triangulaire avant Lemmas V.3.6 and D.4.2, et ça nous donne $1 + 2 \cdot 3/5 = 2.2$.

Algorithm D.5.1 Algorithm PreCond'

Input: A pseudo-basis (\mathbf{B}, \mathbb{I}) of a rank-2 module $M \subseteq K^2$, two parameters $V > 0, B \geq 2$ and a block-size $\beta \in [2, 2d]$

Output: A matrix $\mathbf{Y} \in \text{GL}_2(K)$

1: Set $\varsigma = V^{1/(2d)} \cdot (5\sqrt{2d})^{-1} \cdot \mathcal{N}(M)^{-1/(2d)}$

2: **repeat**

3: Sample $\mathbf{Y} := (\mathbf{y}_1, \mathbf{y}_2)^T \leftarrow \text{DualRound}((\mathbf{B}, \mathbb{I}), \varsigma, \beta, 1/5)$

4: **until** The norm of the first row of $\mathbf{Y} \cdot (\mathbf{B}, \mathbb{I})$ is at most B .

5: **Return** \mathbf{Y}

D.5.4 Updating Lemma V.4.7

Lemma D.5.7. Let $B \geq 1$, $\gamma_{\text{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, $\gamma_{\text{NTRU}} > 1$ and $q \in \mathbb{Z}_{>0}$ be some parameters. Define

$$V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$$

$$\text{and } \gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 8 \cdot B^{1/(2d)} \cdot d^{3/2} \cdot \delta_K.$$

Let (\mathbf{B}, \mathbb{I}) be any γ_{uSVP} -mod-uSVP₂ instance in \mathcal{O}_K^2 , with the B -coprime property and with norm in $[1/2^{2d} \cdot V, 2^{2d} \cdot V]$. Then on input $(\mathbf{B}, \mathbb{I}), \gamma_{\text{HSVP}}, q$, the algorithm Conditioned-to-NTRU outputs $(\mathbf{B}_4, \mathbf{aux})$ such that \mathbf{B}_4 is a $(\gamma_{\text{NTRU}}, q)$ -NTRU instance.

Proof. Modified version of the proof of Lemma V.4.7. Let us fix some $\delta, \gamma_{\text{HSVP}}, \gamma_{\text{NTRU}}$ and q as in the theorem and define V and γ_{uSVP} accordingly.

Let (\mathbf{B}, \mathbb{I}) be the input pseudo-basis, spanning a rank-2 module $M_1 \subset \mathcal{O}_K^2$ with $\mathcal{N}(M_1) \in [1/2^{2d}, 2^{2d}] \cdot V$, with the B -coprime property, and which we know contains a non-zero vector $\mathbf{s}_1 = (u, v)^T \in M_1$ such that $\|\mathbf{s}_1\| \leq 1/\gamma_{\text{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. We will see step by step how the module M_1 is modified by the algorithm, and what happens to its short non-zero vectors.

First step: HNF. After the HNF computation, we have a new pseudo-basis of the form

$$\begin{bmatrix} J_1 & J_2 \\ \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix} \end{bmatrix}$$

for some $a \in K$ and $J_1, J_2 \subset K$. This pseudo-basis generates a rank-2 module M_2 which is the same as the input module M_1 . Hence, M_2 contains a short non-zero vector $\mathbf{s}_2 := \mathbf{s}_1$.

Since our module M_2 is integral, we know that both ideals J_1 and J_2 are integral. Also, since module $M_2 = M_1$ has the B -coprime property, we know that $\mathcal{N}(J_1) \leq B$. Finally, because of the shape of the pseudo-basis, it holds that $\mathcal{N}(J_1) \cdot \mathcal{N}(J_2) = \mathcal{N}(M_2) = \mathcal{N}(M_1)$, which yields $\mathcal{N}(J_2) \geq \mathcal{N}(M_1)/B$.

Second step: from pseudo-basis to basis. Let M_3 be the free module generated by the pseudo-basis

$$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ a & b \end{pmatrix} \end{bmatrix},$$

where $b \leftarrow \text{BalanceIdeal}(J_2, \gamma_{\text{HSVP}})$. Since $J_2 \subseteq \langle b \rangle$ by Lemma V.4.5 and J_1 is an integral ideal, we conclude that $M_2 \subseteq M_3$. Hence, the short vector $\mathbf{s}_3 := \mathbf{s}_2$ is still in M_3 .

Before moving to the next step, let us have a closer look at b . We know from Lemma V.4.5 that $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$ for all $i \leq d$, with $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_2)^{-1/d}$. Using the lower bound on $\mathcal{N}(J_2)$ that we computed above, this shows that $\sigma \leq \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot B^{1/(2d)} \cdot \mathcal{N}(M_1)^{-1/d}$.

Third step: transforming b into q . Let M_4 be the free module generated by the pseudo-basis

$$\begin{bmatrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix} \end{bmatrix},$$

where $h = \lfloor a \cdot q/b \rfloor$. This is the basis output by algorithm `Conditioned-to-NTRU`. The new module M_4 does not contain M_1 anymore, however we will show that its geometry is close to the one of M_3 , so that it has a short non-zero vector if M_3 does.

Recall that M_3 contains a short vector $\mathbf{s}_3 = (u, v)^T$ such that $\|\mathbf{s}_3\| \leq 1/\gamma_{\text{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. Let $x \in \mathcal{O}_K$ be such that $\mathbf{s}_3 = u \cdot (1, a)^T + x \cdot (0, b)^T$. Define $\mathbf{s}_4 = u \cdot (1, h)^T + x \cdot (0, q)^T \in M_4 \setminus \{\mathbf{0}\}$. Unrolling the definition of h and using the equation $u \cdot a + x \cdot b = v$, one can rewrite $\mathbf{s}_4 = (u, v \cdot q/b - u \cdot \{a \cdot q/b\})^T$. We can upper bound the Euclidean norm of \mathbf{s}_4 as follows

$$\begin{aligned} \|\mathbf{s}_4\| &\leq \|u\| + \|v\| \cdot \|q/b\|_\infty + \|u\| \cdot \|\{a \cdot q/b\}\|_\infty \\ &\leq \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot (1 + q \cdot 2\sigma + d\delta_K) \\ &\leq \gamma_{\text{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot (q \cdot 2\sigma + 2 \cdot d\delta_K) \\ &\leq 2\gamma_{\text{uSVP}}^{-1} \cdot d \cdot \delta_K \cdot (\mathcal{N}(M_1)^{1/(2d)} + q \cdot d \cdot \gamma_{\text{HSVP}} \cdot B^{1/(2d)} \cdot \mathcal{N}(M_1)^{-1/(2d)}) \\ &\leq 1/\gamma_{\text{NTRU}} \cdot \sqrt{q}, \end{aligned}$$

where in the last step we used the fact that $\mathcal{N}(M_1)^{1/(2d)} \in [1/2, 2] \cdot V^{1/(2d)}$ and the definitions of V and γ_{uSVP} . We conclude that the pseudo-basis output by `Conditioned-to-NTRU` is indeed a γ_{NTRU} -NTRU instance, as desired. \square

Note that Lemma V.4.8 still works when replacing Lemma V.4.7 by Lemma D.5.7: except from the gap variation, the densest submodule is not changed by taking modules with the B -coprime property. We now state the modified version of Lemma V.4.9.

D.5.5 Updating Lemma V.4.9

Lemma D.5.8 (Updated Lemma V.4.9). *There exists an algorithm `LiftVec`³ such that the following holds. Let q, γ_{HSVP} and (\mathbf{B}, \mathbb{I}) be as in Lemma D.5.7. Let M_1 denote the rank-2 module generated by (\mathbf{B}, \mathbb{I}) , $[\mathbf{C}, \mathbf{aux}] \leftarrow \text{Conditioned-to-NTRU}((\mathbf{B}, \mathbb{I}), q, \gamma_{\text{HSVP}})$ and let M_4 denote the rank-2 free module generated by \mathbf{C} .*

Let $\mathbf{s} \in M_4$. Then, on input $\mathbf{aux}, \gamma_{\text{HSVP}}, (\mathbf{C}, \mathcal{O}_K^2)$ and \mathbf{s} , algorithm `LiftVec` outputs a vector $\mathbf{t} \in M$ such that $\|\mathbf{t}\| \leq \|\mathbf{s}\| \cdot 68 \cdot B^{3/(2d)} \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$.

If given access to a γ_{HSVP} -id-HSVP oracle, algorithm `LiftVec` runs in polynomial time and makes 1 call to the oracle.

³The algorithm is the same.

Proof. Modified version of the proof of Lemma V.4.9. The running-time of the algorithm follows from inspection and from Lemma V.4.5.

Let us show that the output \mathbf{t} of the algorithm is indeed in the module M_1 . Let us keep the notations M_1, M_2, M_3 and M_4 from Figure V.2. In particular, $M_1 = M_2$ is the module generated by the pseudo-basis $((1, a)^T, (0, b)^T), (J_1, \langle b^{-1} \rangle \cdot J_2)$.

From Lemma V.4.5, we know that $\langle b \rangle \cdot J_1^{-1} \cdot J_2^{-1} \subseteq \langle z \rangle$, i.e., $z^{-1} \in J_1 \cdot J_2 \cdot \langle b^{-1} \rangle$. Using the fact that J_1 and $J_2 \cdot \langle b^{-1} \rangle$ are both integral (recall that $J_2 \subseteq \langle b \rangle$ and that M_1 is in \mathcal{O}_K^2), this implies that $z^{-1} \in J_1 \cap J_2 \cdot \langle b^{-1} \rangle$. Since $x, y \in \mathcal{O}_K$, we conclude that $\mathbf{t} = z^{-1} \cdot x \cdot (1, a)^T + z^{-1} \cdot y \cdot (0, b)^T$ is in M_1 as desired.

Let us now upper bound the size of \mathbf{t} . Let us write $\mathbf{s} = (s_1, s_2)^T$ and express the coordinates of \mathbf{t} in terms of s_1 and s_2 . From the equation $\mathbf{s} = x \cdot (1, \lfloor a \cdot q/b \rfloor)^T + y \cdot (0, q)^T$, we obtain $x = s_1$ and $s_2 = x \lfloor a \cdot q/b \rfloor + yq$. This implies that $\mathbf{t} = z^{-1} \cdot (s_1, b/q \cdot (s_2 + s_1 \cdot \{a \cdot q/b\}))^T$. From this, we can upper bound

$$\|\mathbf{t}\| \leq \|z^{-1}\|_\infty \cdot \|\mathbf{s}\| \cdot (1 + \|b/q\|_\infty \cdot (1 + d\delta_K)).$$

Recall from the proof of Lemma D.5.7 that for all $i \leq d$, we have $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$, with $\sigma = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot B^{1/(2d)} \cdot \mathcal{N}(J_2)^{-1/d}$. Hence,

$$\left\| \frac{b}{q} \right\|_\infty \leq \frac{2 \cdot \mathcal{N}(J_2)^{1/d}}{\gamma_{\text{HSVP}} \cdot B^{1/(2d)} \cdot d^2 \cdot \delta_K \cdot q} \quad \text{and} \quad |\mathcal{N}(b)| \geq \frac{\mathcal{N}(J_2)}{\sqrt{B} \cdot (2\gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K)^d}.$$

From Lemma V.4.5, we similarly know that $|\sigma_i(z)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma_z^{-1}$ for all $i \leq d$, where $\sigma_z = \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d}$. Hence we obtain

$$\|z^{-1}\|_\infty \leq 2 \cdot \gamma_{\text{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d} \leq 4 \cdot \gamma_{\text{HSVP}}^2 \cdot d^4 \cdot \delta_K^2 \cdot B^{3/(2d)},$$

where we used the fact that $\mathcal{N}(J_1) \leq B$ thanks to the B -coprime property of the module M_1 .

Finally, recall that $\mathcal{N}(J_2) \leq \mathcal{N}(M_1) \leq 2^{2d} \cdot V$, with $V^{1/d} = q \cdot \gamma_{\text{HSVP}} \cdot d$. Combining everything provides the desired upper bound on $\|\mathbf{t}\|$. \square

D.6 Missing Proofs from Section V.5

D.6.1 Proof of Theorem V.5.2

Note that the assumptions on B and γ in the theorem statement enable the use of all theorems and lemmas from Sections V.5.1, V.5.2 and V.5.3. The runtime statement follows from the runtime statements in Theorems V.5.6 and V.5.9. By using Theorems V.5.6 and V.5.9, we also obtain that the pseudo-basis output by `Randomize` spans a rank-2 and norm-1 module.

Let M' be the module spanned by the output $(\mathbf{B}', \mathbb{I}')$ of `Randomize`, when given as input a module with gap γ . By Theorems V.5.6 and V.5.9, the distribution of M' (over the internal randomness of `Randomize`) is within statistical distance $2^{-\Omega(d)}$ from $\text{QRSF-2-Mod}(D_{B,\gamma}^{\text{rand}})$, where $D_{B,\gamma}^{\text{rand}}$ is as defined in Definition V.5.10. Now, we apply `QRSF-2-Mod` to all the distributions of Definition V.5.10. By the probability preservation properties of the statistical distance and Rényi divergence, and by Lemmas D.6.1, D.6.2, D.6.3, D.6.4 and D.6.6, any event that occurs with probability $\varepsilon \geq 2^{-\Omega(d)}$ for $\text{QRSF-2-Mod}(D_{B,\gamma}^{\text{target}})$ also holds with probability $\Omega(\varepsilon^4)$ for $\text{QRSF-2-Mod}(D_{B,\gamma}^{\text{rand}})$. By observing that $\text{QRSF-2-Mod}(D_{B,\gamma}^{\text{rand}})$ is exactly $D_{B,\gamma}^{\text{module}}$, we obtain that any event that holds for $D_{B,\gamma}^{\text{module}}$ with probability $\varepsilon \geq 2^{-o(d)}$ also holds for M' with probability $\Omega(\varepsilon^4)$ over the internal randomness of `Randomize`.

We now analyze **Recover**. Let M be the module spanned by (\mathbf{B}, \mathbb{I}) . Let U be its densest rank-1 submodule. Let $((\mathbf{B}', \mathbb{I}'), \mathbf{aux})$ be an output of **Randomize** when given (\mathbf{B}, \mathbb{I}) as input, and U' be a densest rank-1 submodule of M' . By Theorems V.5.6 and V.5.9, we have that with probability $1 - 2^{-\Omega(d)}$, the module M' has gap larger than 1 and its densest rank-1 submodule is

$$U' = (\mathcal{N}(\mathfrak{p}) \cdot \det(\mathbf{D}))^{-\frac{1}{2d}} \cdot \mathbf{D} \cdot U \cdot \mathfrak{q}\mathfrak{p}.$$

This completes the proof. \square

D.6.2 Proof of Theorem V.5.3

Assume that $\mathbf{u} \cdot J_1 \not\subseteq M'$, which holds with probability $1 - (1/B)^{\Omega(1)}$ by Lemma V.5.4. We fix x as in Lemma V.5.5. Let $M'' = \mathbf{u} \cdot \mathfrak{p}J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2$. By Lemmas V.5.4 and V.5.5, we have that $M'' \subseteq M'$. By construction, the norm of M'' is $\mathcal{N}(\mathfrak{p}) \cdot \mathcal{N}(M)$, which is equal to $\mathcal{N}(M')$ by Lemma V.3.2, leading to the equality $M' = M''$. This completes the proof of the first statement.

Assume that we have $M' = \mathbf{u} \cdot \mathfrak{p}J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2$ and $\gamma(M) \geq B^{1/(2d)}$, and that $\mathbf{u} \cdot J_1$ is the densest rank-1 submodule of M . As $\mathbf{u} \cdot \mathfrak{p}J_1$ is a rank-1 submodule of M' , we have:

$$\gamma(M') \geq \left(\frac{\sqrt{\mathcal{N}(M')}}{\mathcal{N}(\mathbf{u} \cdot \mathfrak{p}J_1)} \right)^{\frac{1}{d}} = \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \left(\frac{\sqrt{\mathcal{N}(M)}}{\mathcal{N}(\mathbf{u} \cdot J_1)} \right)^{\frac{1}{d}} = \frac{\gamma(M)}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}}.$$

As $\mathcal{N}(\mathfrak{p}) \leq B$, we obtain that $\gamma(M') \geq \gamma(M)/B^{1/(2d)} > 1$. By Lemma II.3.4, we know that M' has a unique densest rank-1 submodule. Now, using the equalities above and the inequalities $\gamma(M) \geq B^{1/(2d)}$ and $\mathcal{N}(\mathfrak{p}) \leq B$, we have

$$\mathcal{N}(\mathbf{u} \cdot \mathfrak{p}J_1) = \mathcal{N}(\mathfrak{p})^{\frac{1}{2}} \cdot \frac{\mathcal{N}(M')^{\frac{1}{2}}}{\gamma(M)^d} \leq \mathcal{N}(M')^{\frac{1}{2}}.$$

Lemma II.3.4 then implies that $\mathbf{u} \cdot \mathfrak{p}J_1$ is contained in the densest rank-1 submodule of M' . By primitivity (see Definition II.3.3), we conclude that it is the densest rank-1 submodule of M' . \square

D.6.3 Proof of Lemma V.5.4

As $\mathbf{u}\mathfrak{p}J_1$ is a primitive rank-1 submodule of M , we can use Lemma V.3.3. It implies that the result holds, except with probability $1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^2$ over the choice of $\bar{\mathbf{b}}^\vee$.

The overall probability (including over the choice of \mathfrak{p}) that $\mathbf{u} \cdot J_1 \subset M'$ holds satisfies:

$$\begin{aligned} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \Pr(\mathfrak{p}) \cdot \Pr(\langle \bar{\mathbf{b}}^\vee, \mathbf{u} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}J_1^{-1} \mid \mathfrak{p}) &= \frac{1}{\pi_K(B)} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \left(\frac{1}{\mathcal{N}(\mathfrak{p})} - \frac{1}{\mathcal{N}(\mathfrak{p})^2} \right) \\ &\leq \frac{1}{\pi_K(B)} \sum_{p \leq B} \sum_{\mathfrak{p} \mid p} \frac{1}{\mathcal{N}(\mathfrak{p})} \\ &\leq \frac{d}{\pi_K(B)} \sum_{p \leq B} \frac{1}{p}, \end{aligned}$$

where the sums indexed by \mathfrak{p} are over the prime ideals of \mathcal{O}_K and the sums indexed by p are over the prime integers. The last inequality comes from the facts that there are at most d ideals \mathfrak{p} over p , and each of them has norm $\geq p$. As $\sum_{p \leq B} 1/p = \log \log B + O(1)$ (see, e.g., [Apo98, Theorem 4.2]) and $\pi_K(B) = \Theta(B/\log B)$, we obtain that the probability above is $\leq (1/B)^{\Omega(1)}$. \square

D.6.4 Proof of Lemma V.5.5

Let $j \in J_1$ with $ju \notin M'$. Since $\langle \mathbf{b}^\vee, ju \rangle_{K_{\mathbb{R}}}$ belongs to $\mathcal{O}_K \setminus \mathfrak{p}$ (by definition of j), we can take a representative $a \in \mathcal{O}_K$ of its inverse in $\mathcal{O}_K/\mathfrak{p}$. We define $y = -\langle \mathbf{b}^\vee, \mathbf{v} \rangle_{K_{\mathbb{R}}} \cdot a \in J_2^{-1}$. By construction, we have $\langle \mathbf{b}^\vee, \mathbf{v} + jyu \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}J_2^{-1}$. This implies that $(\mathbf{v} + jyu) \cdot J_2 \subset M'$. Setting $x = jy$ provides the result. \square

D.6.5 Proof of Theorem V.5.6

The running time bound follows from Theorem V.5.3 and Lemma II.2.12. Now, we write

$$M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) = \mathbf{u} \cdot J_1 + \mathbf{v} \cdot J_2.$$

Let $\mathfrak{p}, \overline{\mathbf{b}^\vee}$ and \mathfrak{q} refer to the random variables sampled during the execution of **Real-CR** and let \mathbf{b}^\vee be a representative of $\overline{\mathbf{b}^\vee}$ in M^\vee . By Theorem V.5.3, we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_{\mathbb{R}}} \notin \mathfrak{p}J_1^{-1}$ with probability $1 - (1/B)^{\Omega(1)}$. In the following, we assume that this holds. We also replace the distribution of \mathfrak{q} by the uniform distribution over norm-1 ideals. By Lemma II.2.12, these two distributions are within 2^{-d} statistical distance from one another. These two assumptions account for the statistical distance upper bound in the theorem statement.

Let $x \in J_1J_2^{-1}$ as in Theorem V.5.3. We have $\langle \mathbf{b}^\vee, \mathbf{v} + x\mathbf{u} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p}J_2^{-1}$. For any choice of x such that the latter holds, the module M' corresponding to the output of **Real-CR** is, by Theorem V.5.3:

$$M' = \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot (\mathbf{u} \cdot J_1\mathfrak{p}\mathfrak{q} + \mathbf{v}' \cdot J_2\mathfrak{q}),$$

where $\mathbf{v}' = \mathbf{v} + x\mathbf{u}$. Note that the QR-factorization of the matrix $[\mathbf{u}|\mathbf{v}']$ is:

$$[\mathbf{u}|\mathbf{v}'] = \mathbf{Q} \cdot \begin{pmatrix} \frac{1}{\gamma} & \gamma \cdot (r+x) \\ 0 & \gamma \end{pmatrix}.$$

We define the norm-1 ideal $J = J_2\mathfrak{q}$. We have:

$$\begin{aligned} M' &= \frac{1}{\gamma \cdot \mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1J_2^{-1}J\mathfrak{p} + \gamma^2 \cdot \begin{bmatrix} r+x \\ 1 \end{bmatrix} \cdot J \right) \\ &= \frac{1}{\sqrt{\gamma'}} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1J_2^{-1}J \frac{\mathfrak{p}}{\mathcal{N}^{1/d}(\mathfrak{p})} + \gamma'^2 \cdot \begin{bmatrix} r+x \\ 1 \end{bmatrix} \cdot J \right), \end{aligned}$$

where $\gamma' = \gamma/\mathcal{N}(\mathfrak{p})^{1/(2d)}$. As the ideal \mathfrak{q} is distributed uniformly over the set of norm-1 ideals, so is J . This implies that the distribution of M' is the same as the distribution of

$$\text{QRSF-2-Mod}(\text{Ideal-CR}_B(\mathbf{Q}, \gamma, J_1, J_2, r)).$$

Still assuming that we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_{\mathbb{R}}} \notin \mathfrak{p}J_1^{-1}$, Theorem V.5.3 gives us that the densest rank-1 submodule of M' is:

$$\frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot \frac{1}{\gamma} \mathbf{u} \cdot J_1\mathfrak{p}\mathfrak{q} = \frac{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}}{\gamma} \cdot \mathbf{Q} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1\mathfrak{q} \frac{\mathfrak{p}}{\mathcal{N}^{1/d}(\mathfrak{p})}.$$

This completes the proof of the theorem. \square

D.6.6 Proof of Lemma V.5.7

For the first statement, we prove that for $\mathbf{D}_0 \in \mathbb{R}^{2 \times 2}$ sampled from $\mathcal{D}(0, 1)^{2 \times 2}$, we have $|\det \mathbf{D}_0| \geq 1/d$ with probability $1 - O((\log d)/d)$. As D_{distort} consists in $\leq d$ independent copies of the latter distribution, the probability of accepting a sample from $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^{2 \times 2}$ when rejecting to D_{distort} is at least $1/d^{O(1)}$.

Observe that $\mathbf{D}_0 = \|\mathbf{d}_1\| \cdot \|\mathbf{d}_2^*\|$, where $\mathbf{d}_1 \sim \mathcal{D}(0, 1)^2$ is the first column of \mathbf{D}_0 and \mathbf{d}_2^* is the projection of the second column orthogonally to the first. As \mathbf{D}_0 is invariant under rotations, conditioned on \mathbf{d}_1 , the vector \mathbf{d}_2^* is distributed as a sample from $\mathcal{D}(0, 1)$ multiplied with a unit vector orthogonal to \mathbf{d}_1 . For these reasons, it suffices to show that with probability $O((\log d)/d)$, the product of two iid samples x, y from $\mathcal{D}(0, 1)$ has magnitude $\geq 1/d$. We have

$$\begin{aligned} \Pr_{x, y \leftarrow \mathcal{D}(0, 1)} [|xy| < 1/d] &\leq O(1/d) + 4 \cdot \Pr_{x, y \leftarrow \mathcal{D}(0, 1)} [xy < 1/d \wedge x, y \in [1/d, 1]] \\ &\leq O(1/d) + c \cdot \Pr_{x, y \leftarrow \mathcal{U}([1/d, 1])} [xy < 1/d], \end{aligned}$$

for some constant c . The latter is $O((\log d)/d)$, allowing to complete the proof of the first statement.

The second statement comes from the invariances of the determinant and vector Gaussian distribution under multiplication by an orthogonal matrix. \square

D.6.7 Proof of Lemma V.5.8

We first show that without the conditioning, the matrix \mathbf{D} from the lemma statement is distributed from $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^{2 \times 2}$. Let us write $\mathbf{D} = [\mathbf{d}_1 | \mathbf{d}_2]$. Then \mathbf{d}_1 is the product of a uniform unit vector and an element sampled from $\chi_{K_{\mathbb{R}}}$. It is hence distributed as a Gaussian vector. Now, as the Gaussian vector distribution is invariant by multiplication by an orthogonal matrix, the distribution of $\mathbf{d}_2 = \mathbf{Q} \cdot (b, c)^T$ is $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^2$, independently of \mathbf{Q} and a .

To conclude, note that the conditioning is with respect to the event “ $\forall i : |\det(\sigma_i(\mathbf{D}))| \geq 1/d^i$ ”, for both D and D_{distort} . \square

D.6.8 Proof of Theorem V.5.9

The runtime claim follows from Lemma V.5.8. Now, let $\mathbf{D} \leftarrow D_{\text{distort}}$ be the matrix sampled in Step 1 of Real-GR. The matrix $\mathbf{D} \cdot \mathbf{Q}$ is also distributed from D_{distort} , by Lemma V.5.7. By Lemma V.5.8 we can write $\mathbf{D}\mathbf{Q} = \mathbf{Q}' \cdot \begin{pmatrix} a & b \\ 0 & c \end{pmatrix}$ with $\mathbf{Q}' \leftarrow \mathcal{U}(\mathcal{O}_2(K_{\mathbb{R}}))$, $a \leftarrow \chi_{K_{\mathbb{R}}}$ and $b, c \leftarrow \mathcal{D}(0, 1)$, conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$. We can then write:

$$\mathbf{D} \cdot \mathbf{M} = \mathbf{Q}' \cdot \begin{pmatrix} a & b \\ 0 & c \end{pmatrix} \cdot \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \cdot \begin{bmatrix} 1/\gamma \cdot J_1 \\ \gamma \cdot J_2 \end{bmatrix} = \mathbf{Q}' \cdot \begin{pmatrix} a & b + ar + b \\ 0 & c \end{pmatrix} \cdot \begin{bmatrix} 1/\gamma \cdot J_1 \\ \gamma \cdot J_2 \end{bmatrix}.$$

Using the equality $\det \mathbf{D} = \mathcal{N}(ab)$, we obtain:

$$\mathbf{M}' = |\det \mathbf{D}|^{-\frac{1}{2d}} \cdot \mathbf{D} \cdot \mathbf{M} = \mathbf{Q}' \cdot \begin{pmatrix} 1 & r' \\ 0 & 1 \end{pmatrix} \cdot \begin{bmatrix} 1/\gamma' \cdot J_1' \\ \gamma' \cdot J_2' \end{bmatrix},$$

where $r' = (b + ar)/c$, $\gamma' = \mathcal{N}(c/a)^{1/(2d)} \cdot \gamma$, $J_1' = (a/\mathcal{N}^{1/d}(a))J_1$ and $J_2' = (c/\mathcal{N}^{1/d}(c))J_2$. This proves the equality of distributions.

We now study $\gamma(M')$. For this, by the above, we can consider **Ideal-GR**. Thanks to the conditioning on the distribution of (a, c) , we have:

$$\gamma' = \mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}} \cdot \gamma = \frac{\mathcal{N}(ac)^{\frac{1}{2d}}}{\mathcal{N}(a)^{\frac{1}{d}}} \gamma \geq \frac{1}{\sqrt{d} \cdot \mathcal{N}(a)^{\frac{1}{d}}} \gamma.$$

Now, note that without the conditioning, the coefficient a would be normally distributed, and the Gaussian tailbound would imply that $\mathcal{N}(a)^{1/d} \leq \sqrt{d}$ with probability $1 - 2^{-\Omega(d)}$. As the rejection occurs with probability at most $1 - 1/d^{O(1)}$ over the choice of (a, c) Gaussian, we still have that $\mathcal{N}(a)^{1/d} \leq \sqrt{d}$ with probability $1 - 2^{-\Omega(d)}$ for (a, c) distributed as in **Ideal-GR**. Overall, we obtain that $\gamma' \geq \gamma/d$ with probability $1 - 2^{-\Omega(d)}$. Using to the QR-standard form of M' with J'_1 and J'_2 of norm 1, we obtain that $\gamma(M') \geq \gamma' > 1$. By Lemma **II.3.4**, the module M' has a unique rank-1 densest submodule. The QR-standard form leads us to consider the following rank-1 submodule of M' :

$$U' = \frac{1}{\gamma'} \cdot \mathbf{Q}' \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J'_1 = |\det \mathbf{D}|^{-1/(2d)} \cdot \mathbf{D} \cdot U.$$

It satisfies $\mathcal{N}(U') = 1/\gamma'^d \leq \gamma(M')^d$. By Lemma **II.3.4**, it is contained in the unique densest rank-1 submodule of M . By primitivity, we have equality. \square

D.6.9 Relations between the distributions of Definition **V.5.10**

Let us first recall the definitions of the considered distributions.

$$\begin{aligned} D_{B,\gamma}^{\text{rand}} &: \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \frac{a}{\mathcal{N}^{1/d}(a)} J_1 J_2^{-1} J \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, \frac{c}{\mathcal{N}^{1/d}(c)} \cdot J, \frac{b + a(r+x)}{c} \right), \\ D_{B,\gamma}^{(1)} &: \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \mathcal{N}^{\frac{1}{d}} \left(\frac{c}{a}\right) \cdot \frac{au}{c} \cdot J_1 J_2^{-1} J \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, J, u \frac{b + a(r+x)}{c} \right), \\ D_{B,\gamma}^{(2)} &: \left(\mathbf{Q}, \gamma \cdot \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, I(J_1, J_2), J, u \frac{b + a(r+x)}{c \text{Exp}(\zeta)} \right), \\ D_{B,\gamma}^{(3)} &: \left(\mathbf{Q}, \gamma', I(J_1, J_2), J, \frac{B^{\frac{1}{d}}}{\mathcal{N}^{1/d}(\mathbf{p})} \cdot u \frac{b + a(r+x)}{c \text{Exp}(\zeta)} \right), \\ D_{B,\gamma}^{(4)} &: \left(\mathbf{Q}, \gamma', I(J_1, J_2), J, r''(J_1, J_2) \right), \\ D_{B,\gamma}^{\text{target}} &: \left(\mathbf{Q}, \gamma', I_1, I_2, r' \right). \end{aligned}$$

Where B, γ, J_1, J_2 and the random variables $\mathbf{Q}, a, b, c, x, \mathbf{p}, I_1, I_2, I(\cdot, \cdot), J, \zeta, u, r', r''$ are defined in Definition **V.5.10**.

Lemma D.6.1. *For any $B \geq 2, \gamma > 0, r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have*

$$D_{B,\gamma}^{\text{rand}}(J_1, J_2, r) = D_{B,\gamma}^{(1)}(J_1, J_2, r).$$

Proof. Let

$$A = \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \frac{a}{\mathcal{N}^{1/d}(a)} J_1 J_2^{-1} J \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, \frac{c}{\mathcal{N}^{1/d}(c)} \cdot J, \frac{b + a(r+x)}{c} \right)$$

be a sample from $D_{B,\gamma}^{\text{rand}}(J_1, J_2, r)$. As the distribution $\mathcal{U}(\mathcal{I}_1)$ is invariant by multiplication by a norm-1 ideal, the random variable $J' = c/\mathcal{N}^{1/d}(c) \cdot J$ is uniformly distributed in \mathcal{I}_1 (over the randomness of J , which is statistically independent of all other random variables). We have

$$A = \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{a}{c} \cdot J_1 J_2^{-1} J' \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, J', \frac{b + a(r+x)}{c} \right).$$

Now let u be uniform in $\{x \in K_{\mathbb{R}}, \forall i \in [d] : |\sigma_i(x)| = 1\}$, and $c' = cu$. As the distribution $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)$ is invariant by multiplication by an element in this set, and the conditioning on (a, c) translates identically to (a, c') , the random variable (a, c') follows the same distribution as the random variable (a, c) (which is statistically independent of all other random variables). We have

$$A = \left(\mathbf{Q}, \gamma \frac{\mathcal{N}\left(\frac{c'}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathbf{p})^{\frac{1}{2d}}}, \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{au}{c'} \cdot J_1 J_2^{-1} J' \frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})}, J', u \frac{b + a(r+x)}{c'} \right).$$

We recognize the distribution $D_{B,\gamma}^{(1)}(J_1, J_2, r)$. □

□

Lemma D.6.2. *For any $B \geq 2, \gamma > 0, r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\text{RD}_2\left(D_{B,\gamma}^{(2)}(J_1, J_2, r) \parallel D_{B,\gamma}^{(1)}(J_1, J_2, r)\right) = O(1).$$

Proof. The result follows from the fact that $\mathcal{N}(c \cdot \text{Exp}(\zeta)) = \mathcal{N}(c)$, the data processing inequality and the bound:

$$\text{RD}_2(c \cdot \text{Exp}(\zeta) \parallel c) = O(1).$$

The rest of the proof is devoted to establishing the latter.

Let $\zeta \in E$ fixed with $\|\zeta\|_{\infty} \leq 1/d$. When $d \geq 2$, we have that $2 - \exp(\zeta_i) > 0$ for all i . Therefore, by Lemma D.1.1 and the fact that $\mathcal{N}(\text{Exp}(\zeta)) = 1$, we have:

$$\text{RD}_2(\mathcal{D}_{K_{\mathbb{R}}}(0, \text{Exp}(\zeta)) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0, 1)) = \mathcal{N}(2 - \text{Exp}(\zeta))^{-\frac{1}{2}}.$$

As $|\zeta_i| \leq 1/d$ holds for all i , each embedding coefficient of $|(2 - \text{Exp}(\zeta))|$ is $\leq 1 - 1/d$. We hence obtain that

$$\mathcal{N}(2 - \text{Exp}(\zeta))^{-\frac{1}{2}} \leq (1 - 1/d)^{-\frac{d}{2}} = O(1).$$

To complete the proof, let us consider ζ as a random variable again. We use Lemma D.1.2 with $K_{\mathbb{R}}$ in place of \mathbb{R} (which is fine, by the multiplicativity property of the Rényi divergence), to obtain:

$$\text{RD}_2(c \cdot \text{Exp}(\zeta) \parallel c) \leq \mathbb{E}_{\zeta} \left(\text{RD}_2(\mathcal{D}_{K_{\mathbb{R}}}(0, \text{Exp}(\zeta)) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0, 1))^{\frac{1}{2}} \right)^2.$$

By the analysis above, the latter upper bound is $O(1)$. □

Lemma D.6.3 (Assuming ERH). *For any $B \geq (\log \Delta_K)^{\Omega(1)}, \gamma > 0, r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\text{RD}_2\left(D_{B,\gamma}^{(3)}(J_1, J_2, r) \parallel D_{B,\gamma}^{(2)}(J_1, J_2, r)\right) = O(1).$$

Proof. Note that $D^{(3)}$ is obtained from $D^{(2)}$ by replacing all occurrences of c by $c \cdot \mathcal{N}^{1/d}(\mathfrak{p})/B^{1/d}$. The result then follows from the data processing inequality and the bound:

$$\text{RD}_2\left(c \cdot \frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}} \parallel c\right) = O(1)$$

The rest of the proof is devoted to proving the latter.

Let us fix a \mathfrak{p} of norm $\leq B$, this implies that $2 - \mathcal{N}^{1/d}(\mathfrak{p})/B^{1/d} > 0$ so by Lemma D.1.1 we have

$$\text{RD}_2\left(c \cdot \frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}} \parallel c\right) \leq \mathcal{N}\left(\frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}} \cdot \left(2 - \frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}}\right)\right)^{-\frac{1}{2}} \leq \left(\frac{B}{\mathcal{N}(\mathfrak{p})}\right)^{\frac{1}{2}}.$$

Now, we consider \mathfrak{p} as a random variable again. Thanks to the above, we have:

$$\mathbb{E}_{\mathfrak{p}}\left(\text{RD}_2\left(c \cdot \frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}} \parallel c\right)^{\frac{1}{2}}\right) \leq \frac{B^{\frac{1}{4}}}{\pi_K(B)} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}}.$$

Abel's summation formula gives (see, e.g., [Apo98, Theorem 4.2]):

$$\begin{aligned} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}} &= \frac{\pi_K(B)}{B^{\frac{1}{4}}} + \frac{1}{4} \int_2^B \pi_K(t) t^{-\frac{5}{4}} dt \\ &\leq 1.1 \frac{B^{\frac{3}{4}}}{\log B} + 1.1 \int_{B_0}^B \frac{t^{-1/4}}{\log(t)} dt + \int_2^{B_0} \pi_K(t) t^{-\frac{5}{4}} dt, \end{aligned}$$

where $B_0 = (\log \Delta_K)^{\Omega(1)}$ is such that for $B \geq B_0$ we have $\pi_K(B) \leq 1.1B/\log B$ (see Section V.2). The last term in the upper bound is $\leq 2^{-1/4} \pi_K(B_0) \leq B_0$. Assuming that $B \geq B_0^2$, the latter is $\leq B^{1/2}$. Overall, we obtain that

$$\sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}} \leq 5 \frac{B^{\frac{3}{4}}}{\log(B)}.$$

Using the lower bound $\pi_K(B) \geq 0.9B/\log B$ from Section V.2, we then obtain that

$$\mathbb{E}_{\mathfrak{p}}\left(\text{RD}_2\left(c \cdot \frac{\mathcal{N}^{1/d}(\mathfrak{p})}{B^{1/d}} \parallel c\right)^{\frac{1}{2}}\right) \leq O(1).$$

Finally, Lemma D.1.2 allows us to conclude. □ □

Lemma D.6.4. For $B \geq 2$, $\gamma \geq d^{1/4} \Delta_K^{1/(2d)}$, $r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:

$$\text{SD}\left(D_{B,\gamma}^{(3)}(J_1, J_2, r), D_{B,\gamma}^4(J_1, J_2)\right) \leq 2^{-\Omega(d)}.$$

To prove Lemma D.6.4, we will use the following result on the closeness to uniformity of a Gaussian distribution over $K_{\mathbb{R}}$, when it is folded modulo an ideal lattice.

Lemma D.6.5 (Adapted from [PRS17, Lemma 6.9]). Let I an ideal, $s \in K_{\mathbb{R}}^+$ and $\mathbf{s} = (\sigma_i(s))_{i \in [d]}$. If $\mathcal{N}(s) \geq \Delta_K \cdot \mathcal{N}(I)$, then we have:

$$\text{SD}(\mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{s}) \bmod I, \mathcal{U}(K_{\mathbb{R}} \bmod I)) \leq 2^{-\Omega(d)}.$$

Proof of Lemma D.6.4. We consider the following sample from $D_{B,\gamma}^{(3)}(J_1, J_2, r)$:

$$\left(\mathbf{Q}, \gamma', I(J_1, J_2), J, B^{\frac{1}{d}} \cdot u \frac{b + a(r+x)}{c \text{Exp}(\zeta) \mathcal{N}^{1/d}(\mathbf{p})} \right).$$

Note that $b \sim \mathcal{D}_{K_{\mathbb{R}}}(0, 1)$ is independent of all other variables and occurs only once in the sample above. Let $b' = B^{1/d} \cdot sb / (c \text{Exp}(\zeta) \mathcal{N}^{1/d}(\mathbf{p}))$. Over the randomness of b (and assuming all other random variables are fixed), it is distributed as $\mathcal{D}_{K_{\mathbb{R}}}(0, B^{1/d} / (|c| \text{Exp}(\zeta) \mathcal{N}^{1/d}(\mathbf{p})))$. We now consider the folding of b' modulo the ideal $I' := \gamma'^{-2} I(J_1, J_2) \cdot J^{-1}$. Lemma D.6.5 implies that if

$$\frac{B}{\mathcal{N}(c) \mathcal{N}(\mathbf{p})} \geq \Delta_K \cdot \mathcal{N}(I'),$$

then $\text{SD}(c' \bmod I', \mathcal{U}(K_{\mathbb{R}} \bmod I')) \leq 2^{-\Omega(d)}$, leading to the result. It hence suffices to prove the premise.

As $I(J_1, J_2), J \in \mathcal{I}_1$ and $\mathcal{N}(\mathbf{p}) \leq B$, using the definition of γ' , it suffices that we have $\gamma^{2d} \geq \Delta_K \mathcal{N}(a)$. By the Gaussian tail bound, we have $\mathcal{N}(a) \leq d^{d/2}$ with probability $1 - 2^{-\Omega(d)}$, which suffices for our purposes. \square

Lemma D.6.6. *For $B \geq (d^d \Delta_k)^{\Omega(1)}$, $\gamma > 0$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\text{SD}\left(D_{B,\gamma}^{(4)}(J_1, J_2), D_{B,\gamma}^{\text{target}}\right) \leq 2^{-\Omega(d)}.$$

Proof. By Lemma II.2.12, we have that the distribution of $\frac{\mathbf{p}}{\mathcal{N}^{1/d}(\mathbf{p})} \cdot u \exp(-\zeta)$ is within statistical distance $2^{-\Omega(d)}$ from $\mathcal{U}(\mathcal{I}_1)$. The latter distribution being invariant by multiplication by norm-1 ideals, we obtain that the distribution of $I(J_1, J_2)$ is at statistical distance $2^{-\Omega(d)}$ from $\mathcal{U}(\mathcal{I}_1)$, over the random choices of \mathbf{p} , s and ζ . As they are independent of \mathbf{Q} , γ' and J , and as the distribution of the last tuple entry is a function of the others, we obtain the result. \square

D.7 Missing Proofs from Section V.6

D.7.1 Proof of Lemma V.6.2

Using the notations from Definition V.5.1, the gap of the module M is equal to $\gamma(M) = \gamma' \mathcal{N}(c/a)^{1/(2d)} / B^{1/(2d)}$. Now, by the conditioning on the pair (a, c) , we have $\mathcal{N}(c) \geq 1/(d^d \mathcal{N}(a))$. Also, by the Gaussian tail bound, we have $\|a\| \leq \sqrt{d}$ with probability $1 - 2^{-\Omega(d)}$. The inequality $\mathcal{N}(a) \leq \|a\|/\sqrt{d}$ then leads to the result. \square

D.7.2 Proof of Lemma V.6.3

Let us write $(\mathbf{B}, \mathbb{I}) = \mathbf{Q} \cdot ((1, 0)^T \cdot 1/\gamma J_1 + (r, 1)^T \cdot \gamma J_2)$ and $\mathbf{Y} = R \cdot (\mathbf{I} + (2d)^{-3/2} \cdot \mathbf{E})$ with R as define in DualRound (Algorithm V.3.1) and $\|e_{ij}\| \leq 1$ for all $i, j \in [2]$ (see Lemma V.3.5). We consider the QR-factorization of $\mathbf{Q}^{-1} \cdot \mathbf{Y} \cdot \mathbf{Q}$:

$$\mathbf{Q}^{-1} \cdot \mathbf{Y} \cdot \mathbf{Q} = R \cdot \mathbf{Q}' \cdot \begin{bmatrix} x & y \\ 0 & z \end{bmatrix},$$

for some $x, y, z \in K_{\mathbb{R}}$. In particular, we have that $\mathcal{N}(x)$ is the algebraic norm of the first column of $\mathbf{I}_2 + \mathbf{E}'$, where $\mathbf{E}' = \mathbf{Q}^{-1} \cdot \mathbf{E} \cdot \mathbf{Q}$ satisfies $\|e'_{ij}\|_{\infty} \leq \sqrt{2d}$ for $i, j \in [2]$. This implies that $\mathcal{N}(x) \leq 1 + 1/(2d)$. In the same vein as in the proof of Theorem V.5.9, this implies that

$$\mathcal{N}(\mathbf{Y} \cdot U) \leq R^d \cdot \left(1 + \frac{1}{2d}\right)^d \cdot \mathcal{N}(U) \leq R^d \cdot \sqrt{e} \cdot \mathcal{N}(U),$$

where $U = \mathbf{Q} \cdot (1, 0)^T \cdot 1/\gamma J_1$. The result then follows from Lemma V.3.6 and the fact that $\mathcal{N}(\mathbf{Y} \cdot M) = \det(\mathbf{Y}) \cdot \mathcal{N}(M)$. \square

D.7.3 Proof of Lemma V.6.5

Wlog, we may assume that the gap of the γ' -wc-mod-uSVP $_2^{\mathcal{N}, \text{mod}}$ instance (\mathbf{B}, \mathbb{I}) satisfies $\gamma' \leq 2^d \Delta_K^{O(1/d)}$, as otherwise the problem can be solved in polynomial time using LLL [LLL82]. We cover the interval $[2 \log(\Delta_K)^{O(1/d)} \cdot \gamma, 2^d \Delta_K^{O(1/d)}]$ by at most $O(d^2 + \log \Delta_K)$ intervals of the form $\gamma \cdot [(1 + 1/(3d))^i, (1 + 1/(3d))^{i+1}]$, and guess uniformly the i for which contains the gap of the module M spanned by (\mathbf{B}, \mathbb{I}) . The guess is correct with probability $\Omega(1/(d^2 + \log \Delta_K))$ and, in the following, we only analyze what happens when this occurs.

The next step is to find a prime ideal \mathfrak{p} such that $\mathcal{N}(\mathfrak{p})^{1/(2d)} \in \gamma' \cdot [(1 + 1/(3d))^{i-1}, (1 + 1/(3d))^i]$. As $\gamma' \geq 2 \log(\Delta_K)^{O(1/d)}$, we can use Lemma II.2.11 to sample \mathfrak{p} uniformly among the prime ideals with norms $\leq (1 + 1/(3d))^{di}$. By the estimates on π_K stated in Section V.2, the value $\mathcal{N}(\mathfrak{p})^{1/(2d)}$ belongs to the appropriate interval with probability $\Omega(1)$. We assume this is the case. Note that we then have that $\gamma'/\mathcal{N}(\mathfrak{p}) \in \gamma \cdot [1, 1 + 1/d]$.

We then sample $\overline{\mathbf{b}}^\vee$ uniformly in $(M^\vee / \mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$, and sparsify the module M by $(\overline{\mathbf{b}}^\vee, \mathfrak{p})$, using Lemma V.3.4. By Lemmas V.3.2 and V.3.3, the gap of the sparsified module M' is $\gamma'/\mathcal{N}(\mathfrak{p})$, with probability $\Omega(1)$, and the pseudo-basis of M' is a valid γ^\approx -wc-mod-uSVP $_2^{\mathcal{N}, \text{mod}}$ instance. Finally, note that when the latter event occurs, if U is the densest rank-1 submodule of M , then $\mathfrak{p}U$ is the densest rank-1 submodule of M' (as in the proof of Theorem V.5.6). This completes the description and the analysis of the reduction. \square

D.7.4 Proof of Lemma V.6.6

By Lemma V.6.2, samples from $D_{\gamma'}^{\text{mod-uSVP}_2}$ and $D_{\gamma'}^{\text{mod-uSVP}_2}$ are indeed γ -mod-uSVP $^{\mathcal{N}}$ instances.

Now, note that $D_{\gamma'}^{\text{mod-uSVP}_2}$ is obtained from $D_{\gamma}^{\text{mod-uSVP}_2}$ by replacing all the occurrences of c by $c \cdot (\gamma'/\gamma)^2$ in Definition V.5.1. The result then follows from the data processing inequality and the bound:

$$\text{RD}_2(c \parallel c \cdot (\gamma'/\gamma)^2) = O(1)$$

The rest of the proof is devoted to proving the latter. We have $\gamma'/\gamma \geq 1$, implying that $2(\gamma'/\gamma)^2 - 1 \geq 1$. By Lemma D.1.1 this implies that

$$\text{RD}_2(c \parallel c \cdot (\gamma'/\gamma)^2) \leq \mathcal{N} \left(\frac{(\gamma'/\gamma)^4}{2(\gamma'/\gamma)^2 - 1} \right)^{1/2} \leq (1 + 1/d)^{2d} = O(1).$$

\square

D.7.5 Proof of Lemma V.6.7

The reduction first runs algorithm `Randomize $_B$` from Theorem V.5.2. It then calls the algorithm `DualRound $_{\zeta, \beta, \varepsilon}$` and `HNF`. The parameters B, ζ, β and ε are set exactly as in the sampling algorithm for $D^{\text{mod-uSVP}_2}$. It then calls the $(D_{\gamma'}^{\text{mod-uSVP}_2}, \gamma'')$ -mod-uSVP $_2^{\mathcal{N}, \text{mod}}$ oracle and pulls the returned rank-1 submodule back to a rank-1 submodule of the input module, using the \mathbf{Y} matrix from `DualRound` and the `aux` output from `Randomize`.

The runtime bound comes from Theorem V.5.2 and Lemma V.3.5. Correctness follows from Theorem V.5.2, Lemmas V.6.3 and Lemma V.6.6. \square