Homomorphic Encryption and Lattices, Spring 2011	Instructor: Shai Halevi
Problem Set #2	
March 17, 2011	Due March 24

1 Lattices and their Determinant

A. Prove that if $\Lambda \subset \mathbb{Z}^n$ is a full-rank integer lattice with prime determinant, then it has no nontrivial refinements. Namely, if $\Lambda \subseteq \Lambda'$ for some integer lattice Λ' then $\Lambda' = \Lambda$ or $\Lambda' = \mathbb{Z}^n$.

B. Prove the converse: if $\Lambda \subset \mathbb{Z}^n$ is a full rank lattice and $\det(\Lambda)$ is a composite, then Λ has a nontrivial refinement. Namely, there exists a lattice Λ' such that $\Lambda \subsetneq \Lambda' \subsetneq \mathbb{Z}^n$.

2 Gram-Schmidt, LLL, and Dual Lattices

Recall that the Gram-Schmidt orthogonalization of a basis $B = (b_1, \ldots, b_n)$ is $\tilde{B} = (\tilde{b_1}, \ldots, \tilde{b_n})$ such that the $\tilde{b_i}$'s are orthogonal to each other and $b_i = \tilde{b_i} + \sum_{j<i} \mu_{i,j}\ \tilde{b_j}$, where $\mu_{i,j} = \langle b_i, \tilde{b_j}\rangle / \|\tilde{b_j}\|^2$.

Recall also that a basis $B = (b_1, \ldots, b_n)$ is LLL reduced if its Gram-Schmidt orthogonalization satisfies

$$\forall\ 1 \le j < i \le n, \qquad |\mu_{i,j}| \ \le\ 1/2 \tag{1}$$

$$\forall\ 1 \le i < n, \qquad \|\tilde{b}_{i-1}\|^2 \cdot \frac{3}{4} \ \le\ \|\tilde{b_i} + \mu_{i,i-1}\tilde{b}_{i-1}\|^2 \tag{2}$$

Note that all the "smallness" properties of LLL-reduced bases actually rely on a weaker first condition, namely that

$$\forall\ 1 \le j < n, \qquad |\mu_{j+1,j}| \ \le\ 1/2 \tag{3}$$

(The stronger condition from Equation (1) is only needed to prove that the numbers do not grow too large during the LLL procedure.) Below we call a basis "*effectively LLL-reduced*" if it satisfies Equations (3) and (2).

Let $B = (b_1, \ldots, b_n)$ be a basis of a full rank lattice Λ, let D' be the dual basis (i.e., $D' = (B^{-1})^t$), and let $D = (d_1, \ldots, d_n)$ be the matrix D' with the order of the columns reversed. Namely

$$\langle b_i, d_j \rangle = \begin{cases} 1 & \text{if } i = n+1-j \\ 0 & \text{otherwise} \end{cases}$$

A. Prove that the following relation holds for all i:

$$\tilde{b_i} \ =\ \tilde{d}_{n+1-i} / \|\tilde{d}_{n+1-i}\|^2 \tag{4}$$

B. Using Equation (4), prove that the following relation holds for all i:

$$\langle b_i, \tilde{b}_{i-1}\rangle / \|\tilde{b}_{i-1}\|^2 = -\langle d_{n+2-i}, \tilde{d}_{n+1-i}\rangle / \|\tilde{d}_{n+1-i}\|^2 \tag{5}$$

C. Using Equations (4) and (5), prove that if B is effectively LLL-reduced then so is D.

3 Lattice-Based Cryptanalysis

The purpose of this question is to cryptanalyze the following simple candidate for a "weak pseudo-random function" (wPRF).

There is a public prime modulus p. (We will assume for convenience that p is very close to a power of two, say $2^n > p > 2^n - 2^{n/2}$ with n the security parameter, hence the bits of a random element modulo p are almost uniform and independent.) The secret key for the weak-PRF is a randomly chosen integer $\tau \in \mathbb{Z}_p$, and on input $x \in \mathbb{Z}_p$ the function outputs $f_\tau(x) = \mathsf{MSB}_k(\tau x \bmod p)$. Namely, reduce $\tau \cdot x$ modulo p (into the interval $[0, p-1]$) and output the k most-significant bits of the result, where k is a parameter. (Think about $k = O(\sqrt{n})$.)

Consider now an attacker that can obtain polynomially many pairs (x_i, y_i) where the x_i's are chosen uniformly in $\mathbb{Z}_p$ and independently, and the y_i's are computed as $y_i = \mathsf{MSB}_k(\tau x_i \bmod p)$. The attacker's goal is to recover the secret τ. Assume that the attacker has d pairs (x_i, y_i) (for some parameter d), and denote $\vec{u} = 2^{n-k} \cdot \langle y_1, \ldots, y_d, 0\rangle$. Consider the $(d+1)$-dimensional lattice with basis

$$B = \begin{pmatrix} p & 0 & \cdots & 0 & x_1 \\ 0 & p & \cdots & 0 & x_2 \\ & \vdots & \ddots & \vdots & \\ 0 & 0 & \cdots & p & x_d \\ 0 & 0 & \cdots & 0 & 1/p \end{pmatrix}$$

A. Prove that for the secret τ and appropriately chosen integers $\kappa_1, \ldots, \kappa_d$, the lattice vector $\vec{v} = B \cdot \langle \kappa_1, \ldots, \kappa_d, \tau\rangle^t$ satisfies $\|\vec{v} - \vec{u}\| \le \sqrt{d+1} \cdot p/2^k$.

B. Prove that for any parameters d and μ, and for randomly chosen $x_1, \ldots, x_d$ (and their corresponding y_i's), it holds with probability at least $1 - p/2^{d(\mu-1)}$ (over the x_i's) that *every vector* $\vec{v} \in \Lambda(B)$ which is as close to $\vec{u}$ as $\|\vec{v} - \vec{u}\| \le p/2^\mu$, has to be of the form $\vec{v} = B \cdot \langle \kappa_1, \ldots, \kappa_d, \tau'\rangle^t$ for some $\tau' = \tau \pmod p$ and some κ_i's.

C. Using A and B, describe a polynomial-time algorithm that recovers the secret τ, assuming that the parameter k is larger than (say) $3\lceil\sqrt{n}\rceil$. Use the fact that LLL can be used to get an approximation algorithm for the closest-vector-problem (CVP) with approximation factor $2^{(d-1)/2}$ in dimension d.