In El Gamal encryption, we can calculate an encryption of a product of two messages from their ciphertexts only. Recall that encryptions have the form $\operatorname{Enc}(P K, M)=\left(g^{R}, P K^{R} \cdot M\right)$ for a random $R$, so given encryptions $(h, k)$ and $\left(h^{\prime}, k^{\prime}\right)$ of messages $M$ and $M^{\prime}$, their pointwise product ( $h h^{\prime}, k k^{\prime}$ ) gives an encryption of $M M^{\prime}$. This property is called a homomorphism of encryptions.

Homomorphisms are useful for implementing certain multiparty computations. Suppose a large number of voters want to elect Alice or Bob as their leader. Usually an election committee collects the private votes, tallies them, and publish a count. This amounts to a multiparty computation of the function

$$
f\left(x_{1}, \ldots, x_{n}\right)=x_{1}+\cdots+x_{n}
$$

where the inputs represent the voters' choices, say 1 for Alice, -1 for Bob, 0 if neutral.
Physical elections are messy and sometimes violent so there has been some interest in making the process more digital. In principle this can be done securely with the methods from the last few lectures, but you would imagine that any reasonable implementation would be quite inefficient especially when there are many voters.

Here is an alternative. The election committee generates a key-pair for El Gamal encryption and posts the public key. Each voter then posts an encryption $C_{i}=\operatorname{Enc}\left(P K, g^{x_{i}}\right)$ of their vote (together with a zero-knowledge proof that $x_{i}$ takes one of the values $-1,0$, or 1 ). The committee then calculates an encryption of the product $g^{x_{1}} \cdots g^{x_{n}}=g^{x_{1}+\cdots+x_{n}}$ and decrypts it. It finds and posts the vote tally by a brute-force discrete $\log$ calculation, which is easily feasible if $n$ is in the millions or so.

Besides its practicality, this mechanism has several desirable features. Message indistinguishability of encryptions guarantees that votes are anonymous relative to the other voters. The election committee can convince the voters that the count was carried out properly by providing a zeroknowledge proof of knowledge.

In 1978 Rivest, Adleman, and Dertouzos asked if there are encryptions that support homomorphism of multiplication and addition. Their intended application was secure outsourcing of computation. Suppose Bob wants Alice to carry out some long computation for him but he wants to keep his input secret. Bob sends Alice encryptions of her inputs under his public key. Alice then homomorphically evaluates the (arithmetic) circuit representing the computation, eventually obtaining an encryption of the output, which she then sends to Bob. Bob decrypts to obtain his output and Alice hasn't learned anything.

## 1 Fully homomorphic encryption

An encryption scheme is fully homomorphic if there is a homomorphic evaluation algorithm that, given a circuit $f$ and encryptions of its inputs, outputs an encryption of its outputs. There are two variants of the definition. In the weaker one, the output of the homomorphic evaluator is required to decrypt correctly.

Definition 1. Algorithm Hom is a weak homomorphic evaluator for a given encryption scheme if for every circuit $f$ and inputs $x_{1}, \ldots, x_{n}, \operatorname{Dec}(S K, C)=f\left(x_{1}, \ldots, x_{n}\right)$, where

$$
C=\operatorname{Hom}\left(P K, f, \operatorname{Enc}\left(P K, x_{1}\right), \ldots, \operatorname{Enc}\left(P K, x_{n}\right)\right)
$$

This is sufficient for secure outsourcing of computation. The evaluator Alice gets no information beyond the encryptions of Bob's input, which she can simulate on her own. The outsourcing protocol extends naturally to this two-party computation proposal:

## Two-party computation from homomorphic encryption

1. Bob generates $S K, P K$ and sends encryptions $B_{1}=\operatorname{Enc}\left(P K, y_{1}\right), \ldots, B_{n}=\operatorname{Enc}\left(P K, y_{n}\right)$ of his inputs together with $P K$ to Alice.
2. Alice creates encryptions $A_{1}=\operatorname{Enc}\left(P K, x_{1}\right), \ldots, A_{n}=\operatorname{Enc}\left(P K, x_{n}\right)$ of her own inputs, computes $C=\operatorname{Hom}\left(P K, f, A_{1}, \ldots, A_{n}, B_{1}, \ldots, B_{n}\right)$ and sends $C$ to Bob.
3. Bob outputs $\operatorname{Dec}(S K, C)$.

This implementation is still secure for Bob, but it is not clear why it should be secure for Alice. It could be that the ciphertext $C$ contains Alice's input in some part that is ignored by the decryption algorithm. To rule out this possibility, the stronger definition imposes the additional security requirement that the output of the homomorphic evaluator looks like an encryption of $f\left(x_{1}, \ldots, x_{n}\right)$ that is independent of its inputs.
Definition 2. Hom is a strong homomorphic evaluator if the random variables

$$
\left(P K, \operatorname{Enc}\left(P K, x_{1}\right), \ldots, \operatorname{Enc}\left(P K, x_{n}\right), C\right) \text { and }\left(P K, \operatorname{Enc}\left(P K, x_{1}\right), \ldots, \operatorname{Enc}\left(P K, x_{n}\right), \operatorname{Enc}(P K, y)\right)
$$

are $\varepsilon$-statistically close, where $y$ and $C$ are the outputs of $f$ and Hom, respectively.
Definition 2 implies Definition 1 because the decryption algorithm itself can be used as a distinguisher. This distinguisher is not efficient since it needs to sample a secret key consistent with the given public key (which is one reason why the indistinguishability requirement is statistical).
The above two-party computation protocol is honest-but-curious secure assuming Hom is strongly secure because Bob can simulate his view from his output. The protocol can be extended to multiple parties. Here is a 3 -party computation that is secure with respect to any pair of honest-but-curious parties.

## Three-party computation from homomorphic encryption

1. Alice and Bob run a two-party computation to generate a public key $P K$ and private random shares $S K_{A}$ and $S K_{B}$ that add up to the secret key $S K$.
2. Alice and Bob publish $P K$ and send encryptions of their input to Charlie.
3. Charlie encrypts his own inputs, performs the homomorphic evaluation, and publishes the encrypted output $C$.
4. Alice and Bob run the two-party computation $\operatorname{Dec}\left(S K_{A}+S K_{B}, C\right)$. Bob privately outputs the outcome.

Besides the improvement in security, another advantage over the BGW protocol from last lecture is that the number of messages exchanged is independent of the size of the circuit being evaluated. In summary, fully homomorphic encryption provides simple and efficient solutions to difficult cryptographic problems.

In 2008 Gentry proposed the first candidate homomorphic encryption scheme. This was a significant proof of concept that has since been simplified and improved in several ways. Despite remarkable progress current proposals are still impractical and rely on non-standard security assumptions. To explain how these work we first need to revisit LWE-based encryption.

## 2 Learning with errors revisited

Let's start with the LWE-based encryption from Homework 2 with some of the notation changed. We will be underlining all bounded objects (scalars, vectors, matrices) because boundedness will be important for the discussion. The secret key is a bounded vector $\underline{s}$ and the public key $P K$ is the shortLWE instance $(A, b=\underline{e}+\underline{s} A)$. Then $\underline{s} A-b$ is equal to $\underline{e}$. In matrix form this equation can be written as

$$
\underline{s k} \cdot P K=\left[\begin{array}{ll}
\underline{s} & -1
\end{array}\right] \cdot\left[\begin{array}{c}
A \\
\underline{e}+\underline{s} A
\end{array}\right]=\underline{e},
$$

where $\underline{s k}$ is the vector $(\underline{s},-1)$. To encrypt a message $m$ Bob sends $\left(A \underline{x}+\underline{e}^{\prime}, b \underline{x}+\underline{e}^{\prime \prime}+m\right)$ to Alice. In matrix form, the encryption is given by

$$
\operatorname{Enc}(P K, m)=P K \cdot \underline{x}+\left[\begin{array}{l}
\underline{e^{\prime}} \\
\underline{e}^{\prime \prime}
\end{array}\right]+\left[\begin{array}{l}
0 \\
m
\end{array}\right] .
$$

To decrypt $C=\operatorname{Enc}(P K, m)$, Alice calculates $s k \cdot C$ to obtain

$$
s k \cdot C=s k \cdot P K \cdot \underline{x}+\underline{s k} \cdot\left[\begin{array}{c}
\underline{e}^{\prime} \\
\underline{e}^{\prime \prime}
\end{array}\right]+\underline{s k} \cdot\left[\begin{array}{c}
0 \\
m
\end{array}\right]=\underline{e} \cdot \underline{x}+\underline{s k} \cdot\left[\begin{array}{c}
\underline{e}^{\prime} \\
\underline{e}^{\prime \prime}
\end{array}\right]-m .
$$

If the message is represented by the significant bits of $m$ Alice can recover it because all the "noise" adds up to a bounded scalar.

It turns out that the encryption remains secure even if Bob forgets to add the noise ( $\underline{e}^{\prime}, \underline{e}^{\prime \prime}$ ) (assuming $A$ has sufficiently many columns). Under this simplification, the decryption algorithm computes $s k \cdot C=\underline{e} \cdot \underline{x}-m$, so the secret key no longer needs to be bounded for decryption to succeed. This is the original LWE-based encryption scheme for which Oded Regev was awarded the 2018 Gödel prize a few weeks ago.

Regev encryption. The secret and public keys are the vector and matrix

$$
s k=\left[\begin{array}{ll}
s & -1
\end{array}\right] \in \mathbb{Z}_{q}^{n}, \quad P K=\left[\begin{array}{c}
A \\
\underline{e}+s A
\end{array}\right]
$$

where $(A, \underline{e}+s A)$ is a (regular) LWE instance with $b$-bounded noise. To encrypt, Bob chooses a random vector $\underline{x}$ with $-1,0,1$ entries and outputs

$$
\operatorname{Enc}(P K, m)=P K \cdot \underline{x}+\left[\begin{array}{l}
0 \\
m
\end{array}\right]
$$

where $m$ is encoded in the most significant bits. To decrypt $C$, Alice calculates $s k \cdot C$, flips its sign, and rounds to the closest message. Decryption is correct because

$$
s k \cdot C=s k \cdot P K \cdot \underline{x}+\left[\begin{array}{ll}
s & -1
\end{array}\right] \cdot\left[\begin{array}{c}
0 \\
m
\end{array}\right]=\underline{e} \cdot \underline{x}-m
$$

and the scalar $\underline{e} \cdot \underline{x}$ is $b n$-bounded in magnitude.
We don't have all the tools to rigorously prove the security of Regev encryption but it will be instructive to sketch the proof anyway. Message simulatability requires that ( $P K, \operatorname{Enc}(P K, m)$ ) can be simulated without knowing $m$. This random variable has the form ( $P K, P K \cdot \underline{x}+c(m)$ ) where $c(m)$ is the column vector $(0, m)$. By the LWE assumption this is indistinguishable from $(R, R \cdot \underline{x}+c(m))$ for a random matrix $R$. The missing ingredient is the so-called leftover hash lemma which implies that $(R, R \cdot \underline{x})$ is $\varepsilon$-statistically close to a pair of uniform and independent
entries $(R, r)$ (provided $R$ has at least $n \log q+2 \log (1 / \varepsilon)$ columns). Then $(R, R \cdot \underline{x}+c(m))$ is also $\varepsilon$-statistically close to $(R, r)$, so encryptions are indistinguishable from uniform random strings.

For the proof of security it is irrelevant how the message $m$ is represented in the column vector $c(m)$. If $c(m)$ was instead chosen to be the column vector $(m, 0)$, encryption and decryption would give

$$
\operatorname{Enc}(P K, m)=P K \cdot \underline{x}+\left[\begin{array}{c}
m \\
0
\end{array}\right] \quad s k \cdot C=s k \cdot P K \cdot \underline{x}+\left[\begin{array}{ll}
s & -1
\end{array}\right] \cdot\left[\begin{array}{c}
m \\
0
\end{array}\right]=\underline{e} \cdot \underline{x}+m \cdot s k_{1},
$$

where $s k_{1}$ is the first entry of the secret key.
Assuming the message is a single bit $\underline{m} \in\{0,1\}$ it can still be read off from $s k \cdot C$ with good probability: $s k \cdot C$ is $b n$-close to zero when $\underline{m}=0$ and $b n$-close to $s k_{1}$ when $\underline{m}=1$. Unless $s k_{1}$ itself happens to be of magnitude at most $2 b n, \underline{m}$ can be reliably decrypted. Since $s k_{1}$ is random, the probability of a decryption failure is at most $4 b n / q$ over the choice of the secret key.

The failure probability can be reduced exponentially by repeating the encryption independently $n$ times, using a different position to encode the message in every instantiation. To encrypt a $m$, Bob sends the $n$-tuple

$$
\operatorname{Enc}(P K, \underline{m})=\left(P K \cdot \underline{x_{1}}+c_{1}(\underline{m}), P K \cdot \underline{x_{2}}+c_{2}(\underline{m}), \ldots, P K \cdot \underline{x_{n}}+c_{n}(\underline{m})\right),
$$

where $c_{i}(\underline{m})$ is the column vector with $m$ in position $i$ and zeroes elsewhere, and $\underline{x_{1}}, \ldots, \underline{x_{n}}$ are independent random $\{-1,0,1\}$-valued column vectors. This encryption can compactly be represented in the matrix form

$$
\begin{equation*}
\operatorname{Enc}(P K, \underline{m})=P K \cdot \underline{X}+\underline{m} \cdot \underline{I} \tag{1}
\end{equation*}
$$

where $\underline{X}$ is a random $n$-row $\{-1,0,1\}$-valued matrix and $\underline{I}$ is the $n \times n$ identity matrix.
To decrypt a bit $m$, Alice computes the vector-matrix product

$$
\begin{equation*}
s k \cdot C=s k \cdot P K \cdot \underline{X}+s k \cdot \underline{m} \cdot \underline{I}=\underline{e} \cdot \underline{X}+\underline{m} \cdot s k \tag{2}
\end{equation*}
$$

then rounds all the entries to 0 or $q / 2$ depending on which number is closer modulo $q$. If $\underline{m}=0$ and $b n<q / 4$ then all entries will round to zero and decryption is correct. If $\underline{m}=1$ then the probability that all entries round to zero is at most $2^{-n+1}$ because the first $n-1$ entries of the secret key are random. ${ }^{1}$

## 3 Shallow homomorphic encryption

Let's investigate the homomorphic properties of encryption 1 . Since every circuit can be built from NOT and AND gates it is enough to study these two operations. In $\mathbb{Z}_{q}$ arithmetic, NOT $\underline{m}$ is the function $1-\underline{m}$ while $\underline{m}$ AND $\underline{m}^{\prime}$ is multiplication $m \cdot m^{\prime}$, so we are looking for a way to calculate encryptions of differences and products of messages from their respective encryptions.

In this respect equation (2) has a very useful algebraic property. Messages are "approximate" (modular) eigenvalues of their ciphertext matrices, with the secret key being the corresponding "approximate" eigenvector. If the relations $s k \cdot C=\underline{m} \cdot s k$ and $s k \cdot C^{\prime}=\underline{m^{\prime}} \cdot s k$ were exact then scalar operations on messages could be emulated by the same matrix operations on their corresponding ciphertexts

$$
s k \cdot\left(C-C^{\prime}\right)=\left(\underline{m}-\underline{m}^{\prime}\right) \cdot s k \quad \text { and } \quad s k \cdot C C^{\prime}=\underline{m m^{\prime}} \cdot s k .
$$

[^0]resulting in a fully homomorphic encryption scheme.
Let us now consider the effect of the noise. Assume that $C$ and $C^{\prime}$ now satisfy the equations
$$
s k \cdot C=\underline{m} \cdot s k+\underline{e} \quad \text { and } \quad s k \cdot C^{\prime}=\underline{m}^{\prime} \cdot s k+\underline{e}^{\prime}
$$
for some $b$-bounded $e$ and $e^{\prime}$. Then
$$
s k \cdot\left(C-C^{\prime}\right)=\left(\underline{m}-\underline{m}^{\prime}\right) \cdot s k+\left(\underline{e}-\underline{e}^{\prime}\right)
$$
so $C-C^{\prime}$ represents the message $\underline{m}-\underline{m}^{\prime}$, but the best we can say about $\underline{e}-\underline{e}^{\prime}$ is that it is $2 b$-bounded so the noise has grown in magnitude. For ciphertext multiplication, we have
\[

$$
\begin{align*}
s k \cdot C C^{\prime} & =(\underline{m} \cdot s k+\underline{e}) \cdot C^{\prime} \\
& =\underline{m} \cdot s k \cdot C^{\prime}+\underline{e} \cdot C^{\prime} \\
& =\underline{m} \cdot\left(\underline{m}^{\prime} \cdot s k+\underline{e}^{\prime}\right)+\underline{e} \cdot C^{\prime}  \tag{3}\\
& =\underline{m m^{\prime}} \cdot s k+\underline{m} \cdot \underline{e}^{\prime}+\underline{e} \cdot C^{\prime} .
\end{align*}
$$
\]

Owing to the likely presence of large entries in the matrix $C^{\prime}$ this expression is not of the desired form $\underline{m m^{\prime}} \cdot s k$ plus some bounded vector.

This seemingly intractable difficulty can be resolved by an incredibly naive idea: Before multiplying, replace each entry in the matrix $C^{\prime}$ by its bit decomposition. For example if $q=8$ and you see a 3 in $C$ replace it by a 0 , a 1 , and a 1 . Miraculously, all that remains is to tweak the encryption to be syntactically compatible with this change. To describe this modification we need a bit of bit decomposition math.

Bit decomposition The bit decomposition function $\underline{D}(z)$ is the function that maps a number $z \in \mathbb{Z}_{q}$ to its $\log q$-long bit represnentation. For example, $\underline{D}(3)=(0,1,1)$ when $q=8$. Although the function $\underline{D}$ outputs a sequence of bits, we will think of it as a function from $\mathbb{Z}_{q}$ to $\mathbb{Z}_{q}^{t}$ where $t=\log q$. Then $\underline{D}$ has an inverse $D^{\dagger}: \mathbb{Z}_{q}^{t} \rightarrow \mathbb{Z}_{q}$ given by

$$
\begin{equation*}
D^{\dagger}\left(z_{t-1}, \ldots, z_{0}\right)=2^{t-1} z_{t-1}+\cdots+2 z_{1}+z_{0} . \tag{4}
\end{equation*}
$$

The inverse $D^{\dagger}$ is a linear function (between $\mathbb{Z}_{q}$-modules), even though $\underline{D}$ itself is not.
The functions $\underline{D}$ and $D^{\dagger}$ can be applied pointwise to each element of a given vector or matrix. Applying $\underline{D}$ to a matrix increases the number of columns by a factor of $t=\log q$. For example, when $q=8$,

$$
\text { if } C=\left[\begin{array}{ll}
3 & 2 \\
0 & 7
\end{array}\right] \quad \text { then } \quad \underline{D}(C)=\left[\begin{array}{ll}
\underline{D}(3) & \underline{D}(2) \\
\underline{D}(0) & \underline{D}(7)
\end{array}\right]=\left[\begin{array}{llllll}
0 & 1 & 1 & 0 & 1 & 0 \\
0 & 0 & 0 & 1 & 1 & 1
\end{array}\right]
$$

To recover $C$ back from $\underline{D}(C)$, the entries of each row should be divided into triples and each triple should be combined using formula (4). This amounts to left multiplication by the matrix

$$
D^{\dagger}=\left[\begin{array}{llllll}
4 & 2 & 1 & & & \\
& & & 4 & 2 & 1
\end{array}\right] .
$$

You can verify that $D^{\dagger} \cdot \underline{D}(C)$ recovers $C$.
In general, if $C$ has $n$ rows, the operation $\underline{D}(C)$ can be inverted by left-mutliplying its output with the $t n$-by- $n$ matrix $D^{\dagger}$ that has the vector $\left(2^{t-1}, \ldots, 2,1\right)$ arranged diagonally $n$ times (i.e., the tensor product of this vector and the $n$-by- $n$ identity matrix).

Gentry-Sahai-Waters (GSW) shallow homomorphic encryption. The keys are the same as in Regev encryption. The encryption procedure is

$$
\begin{equation*}
E n c(P K, \underline{m})=P K \cdot \underline{X}+\underline{m} \cdot D^{\dagger} . \tag{5}
\end{equation*}
$$

The only difference from (1) is that the identity matrix is replaced by the inverse bit decomposition matrix $D^{\dagger}$. This does not affect security. Regarding decryptability, because $D^{\dagger}$ contains the identity matrix, the encryption (1) is embedded in those columns of the GSW encryption that are indexed by multiples of $t$. Therefore the same decryption procedure works after discarding the other columns. If all the columns are preserved, however, left multiplication by the secret key gives

$$
\begin{equation*}
s k \cdot C=s k \cdot\left(P K \cdot \underline{X}+\underline{m} \cdot D^{\dagger}\right)=\underline{e} \cdot \underline{X}+\underline{m} \cdot s k \cdot D^{\dagger} . \tag{6}
\end{equation*}
$$

As before, homomorphic subtraction of ciphertexts is entrywise matrix subtraction. Given ciphertexts $C$ and $C^{\prime}$ satisfying

$$
s k \cdot C=\underline{m} \cdot s k \cdot D^{\dagger}+\underline{e} \quad \text { and } \quad s k \cdot C^{\prime}=\underline{m}^{\prime} \cdot s k \cdot D^{\dagger}+\underline{e}^{\prime}
$$

for noise vectors $\underline{e}, \underline{e}^{\prime}$, by linearity $C-C^{\prime}$ and $m-m^{\prime}$ satisfy a relation of the same form. If $\underline{e}$ is $b$-bounded and $\underline{e}^{\prime}$ is $b^{\prime}$-bounded, the noise vector $\underline{e}+\underline{e}^{\prime}$ is $\left(b+b^{\prime}\right)$-bounded.

Homomorphic multiplication of ciphertexts is now given by the formula $C \cdot \underline{D}\left(C^{\prime}\right)$. To analyze it we repeat calculation (3) for GSW ciphertexts:

$$
\begin{aligned}
s k \cdot C \cdot \underline{D}\left(C^{\prime}\right) & =\left(\underline{m} \cdot s k \cdot D^{\dagger}+\underline{e}\right) \cdot \underline{D}\left(C^{\prime}\right) \\
& =\underline{m} \cdot s k \cdot D^{\dagger} \underline{D}\left(C^{\prime}\right)+\underline{e} \cdot \underline{D}\left(C^{\prime}\right) \\
& =\underline{m} \cdot s k \cdot C^{\prime}+\underline{e} \cdot \underline{D}\left(C^{\prime}\right) \\
& =\underline{m} \cdot\left(\underline{m^{\prime}} \cdot s k \cdot D^{\dagger}+\underline{e}^{\prime}\right)+\underline{e} \cdot \underline{D}\left(C^{\prime}\right) \\
& =\underline{m m^{\prime}} \cdot s k \cdot D^{\dagger}+\underline{m} \cdot \underline{e}^{\prime}+\underline{e} \cdot \underline{D}\left(C^{\prime}\right) .
\end{aligned}
$$

This is precisely of the desired form: If $\underline{e}$ and $\underline{e}^{\prime}$ are $b$ and $b^{\prime}$-bounded, respectively, then $\underline{e}^{\prime}+\underline{e} \cdot \underline{D}\left(C^{\prime}\right)$ is $\left(n b+b^{\prime}\right)$-bounded.

To summarize, homomorphic subtraction and multiplication preserve the form of ciphertexts, but the noise magnitude grows with each operation. If a circuit $f$ is evaluated homomorphically gate-by-gate, the ratio between the noise at the output and the noise at the input grows by a factor that is exponential in its depth. ${ }^{2}$ If the depth of $f$ is smaller than the logarithm of the modulus $q$ then the noise never grows past the decryption threshold and the output of the homomorphic evaluator decrypts correctly. ${ }^{3}$

## 4 Denoising

The final component for weak fully homomorphic encryption is a mechanism for denoising ciphertexts. This is a procedure Den that, given a ciphertext $C$ (and the public key), outputs another ciphertext $C^{\prime}$ that encrypts the same message but with less noise: If $s k \cdot C=\underline{m} \cdot s k \cdot D^{\dagger}+\underline{e}$ for some $b^{+}$-bounded $\underline{e}$ then $s k \cdot C^{\prime}=\underline{m} \cdot s k \cdot D^{\dagger}+\underline{e}^{\prime}$ for some $b^{-}$-bounded $\underline{e}^{\prime}$ with $b^{-}<b^{+}$.

So far all the LWE transformations increase the amount of noise, so it is not clear that denoising is at all plausible. Where should we even begin to look? The answer is again incredibly naive:

[^1]If you prefer equations to mantras, that translates to

$$
\begin{equation*}
\operatorname{Den}(P K, C)=\operatorname{Hom}\left(P K, D e c_{C}, \operatorname{Enc}(P K, s k)\right) \tag{7}
\end{equation*}
$$

where $D e c_{C}$ is the decryption circuit with the input ciphertext $C$ hardwired into it.
Decryption functionality says that $D e c_{C}(s k)$ outputs the message $m$. Homomorphic evaluation functionality says that a homomorphic evaluation of $D e c_{C}$ given an encryption of $s k$ outputs a new encryption $C^{\prime}$ of $m$.

What about noise? As long as $b^{+} \leq q / 4 n$ decryption is guaranteed to succeed. On the other hand, the noise of the homomorphic evaluator's output only depends on the noise of its input and the depth of the decryption circuit. Since the input is an actual encryption its noise is completely determined by the LWE parameters. In the case of GSW encryption, (5) and (6) reveal that the input noise is at most $n$ times the LWE noise bound. Using a suitably shallow implementation of GSW decryption, its contribution can be bounded by some polynomial in $n$. In conclusion, the amount $b^{-}$of noise in $C^{\prime}$ can be bounded by some fixed polynomial in $n$ that is independent of $q$.

To implement weak fully homomorphic evaluation, start with shallow evaulation until the upper noise threshold $b^{+}$is reached, denoise down to $b^{-}$, and keep going until the output gate is reached.

There is one more issue to be discussed. To implement (7), the denoising circuit needs to know the encryption of the secret key. The only place it could possibly obtain this information from is the public key itself. To support full homomorphic evaluation, the public key of GSW encryption ought to be augmented with an encryption of the (bits of the) secret key. Encryption schemes that are immune against such disclosures are said to satisfy circular security.

Definition 3. A public-key encryption scheme is circular-secure if for every message $M$, the random variable $(P K, \operatorname{Enc}(P K, s k), \operatorname{Enc}(P K, M))$ is indistinguishable from the output of a simulator with the usual parameters.

Circular security is not a direct consequence of (multiple) message simulatability. In the definition from Lecture 4 the challenge message does not depend on the secret key. It is not known if El Gamal encryption is circular-secure under the DDH assumption, or if the Regev and GSW encryptions are circular-secure under the LWE assumption, but it appears plausible that they might be. ${ }^{4}$

## 5 Homomorphic encryption and obfuscation

On several occasions we used virtual black-box (VBB) obfuscation to argue the plausibility of cryptographic constructions like public-key encryption and signatures based on the Fiat-Shamir heuristic. We are now in a position to sketch the proof that general-purpose VBB obfuscation does not exist.

VBB security demands that for every circuit $C$, the output of any learner that is given an obfuscation $\operatorname{Obf}(C)$ is indistinguishable from the output of a simulator that had oracle access to $C$. For starters we will argue that VBB obfuscation for a pair of circuits is impossible: Namely, there exist circuits $P$ and $T$ and a learner that is given their obfuscations whose output cannot be simulated by oracle access to $P$ and $T$.

[^2]The circuit names $P$ and $T$ stand for point and test, respectively. Their descriptions involve three secret keys in, out, and feature that are chosen uniformly at random from $\{0,1\}^{k}$. The point circuit maps $k$-bit inputs to $k$-bit outputs but evaluates to zero everywhere except at in:

$$
P(x)= \begin{cases}\text { out, }, & \text { if } x=\text { in } \\ 0, & \text { if not }\end{cases}
$$

The test circuit takes the description of a circuit $C$ as its input and outputs feature only if $C$ maps in to out:

$$
T(C)= \begin{cases}\text { feature }, & \text { if } C(\text { in })=\text { out } \\ 0, & \text { if not }\end{cases}
$$

Given the circuits $P$ and $T$, feature can be recovered by running $T$ on $P$. The same is true if $\operatorname{Obf}(T)$ is evaluated on $\operatorname{Obf}(P)$, so there is a linear-size learner that outputs feature. However, any VBB simulator is very unlikely to output feature: To do this, the simulator must either happen to query $P$ on in to get out, or if not query $T$ on a circuit that maps $i n$ to out, or if not guess feature at random. Each of these events has probability at most $s \cdot 2^{-k}$ for a simulator of size $t$.

How about obfuscating a single circuit? To rule this out, it is natural to try and combine $P$ and $T$ into a single hard circuit $H$. This is usually done by furnishing $H$ with an extra bit of input that tells it whether to operate in $P$-mode or in $T$-mode:

$$
H(\text { mode }, z)= \begin{cases}P(z), & \text { if } \text { mode }=\mathrm{P} \\ T(z), & \text { if mode }=\mathrm{T}\end{cases}
$$

To mimic the proof we just gave, we first need to show how to extract feature from an obfuscation $O=\operatorname{Obf}(H)$. To do this $T$ should be evaluated on $P$, which amounts to first printing the circuit $O_{\mathrm{P}}(x)=O(\mathrm{P}, x)$ and then computing the output $O\left(\mathrm{~T}, O_{\mathrm{P}}\right)$. This involves running the circuit $O$ on itself as an input. Unless the size of $O$ is smaller than the length of its input (which means $O$ should be very primitive), $O$ cannot be evaluated on itself.

This is where fully homomorphic encryption comes into play. It allows for the evaluation of the test circuit to be outsourced to the learner. Consider instead the circuit $H$ that outputs

$$
H(\text { mode }, z)= \begin{cases}P K, \operatorname{Enc}(P K, i n), & \text { if } \text { mode }=\mathrm{I} \\ \text { out }, & \text { if } \text { mode }=\mathrm{P} \text { and } z=\text { in } \\ \text { feature }, & \text { if mode }=\mathrm{T} \text { and } \operatorname{Dec}(s k, z)=\text { out }\end{cases}
$$

and outputs zero otherwise.
Given the obfuscated circuit $O=\operatorname{Obf}(H)$ the learner can now extract feature like this: First, it runs $O(\mathrm{I})$ to obtain $P K$ and the encryption of $i n$. It then evaluates the circuit $O_{\mathrm{P}}$ homomorphically on the encryption of $i n$ to obtain an encryption $C$ of out. Finally it runs $O(\mathrm{~T}, C)$ to output feature.

From the simulator's perspective, oracle access to $H$ is $(s, \varepsilon)$-indistinguishable from oracle access to a circuit that is the same except that it outputs a simulated encryption in mode I. Then queries in mode I give no information about $P$ or $T$ so $H$ outputs feature with probability at most $s \cdot 2^{-k}+\varepsilon$.

The conclusion is that if fully homomorphic encryption is possible, then VBB obfuscation isn't. It turns out that fully homomorphic encryption can be constructed from VBB obfuscation (without extra assumptions like circular security of LWE-based encryptions), so VBB obfuscation bites its own tail: If it exists, then it doesn't exist.


[^0]:    ${ }^{1}$ The decryption error can be eliminated at a small price in security.

[^1]:    ${ }^{2}$ The depth of the circuit is the maximum path length in the underlying graph.
    ${ }^{3}$ More precisely, the depth ought to be at $\operatorname{most}^{\log }{ }_{n}(q / 4 b n)$ where $b$ is the encryption noise.

[^2]:    ${ }^{4}$ In contrast, they are secure when the encrypted message depends on the public key.

