Problem 1.

Let $a = \sigma, b = H(m)$. It suffices to find 2 primes p', q' with poly (λ) length, $n' = p'q'$, and $e' \in \mathbb{Z}_{\phi(p'q')}$, such that $a^{e'} \equiv b \pmod{n'}$.

Sample $p' = 2 \cdot 3^k + 1, q' = 2 \cdot 5^\ell + 1$ to be two primes, where $k, \ell = \text{poly}(\lambda)$. Since $\phi(p') = 2 \cdot 3^k$, the cylic group $\mathbb{Z}_{p'}^*$ has $2 \cdot 3^{k-1}$ generators. Due to the randomness of a, b, with $\Omega(1)$ probability, a and b are both generators of $\mathbb{Z}_{p'}$. To verify this, one can compute $\left[a^{(p'-1)/2} \mod p'\right]$ and $\left[a^{(p'-1)/3} \mod p'\right]$. If both of them are not $1, a$ is a generator. Similar things holds for q'. In conclusion, with probability $\Omega(1)$, one finds primes $p' = 2 \cdot 3^k + 1$, $q' = 2 \cdot 5^\ell + 1$ such that a, b are both generators of $\mathbb{Z}_{p'}$ and $\mathbb{Z}_{q'}$.

Now one compute $e_1 = [\log_a b \pmod{p'-1}]$ and $e_2 = [\log_a b \pmod{q'-1}]$. Since $p'-1$ and $q'-1$ have small prime factors, the Dlog can be efficiently computed. Finally, solve the following equation

$$\begin{cases} e' \equiv e_1 \pmod{p'-1} \\ e' \equiv e_2 \pmod{q'-1}. \end{cases} \tag{1}$$

One the one hand, $\gcd(p'-1, q'-1) = 2$. On the other hand, since a, b are both generators of $\mathbb{Z}_{p'}$ and $\mathbb{Z}_{q'}$, we have e_1 and e_2 are both odds. Thus (1) has a solution e'. This e' satisfies $a^{e'} \equiv b \pmod{p'q'}$.

Problem 2.

In order for our security reduction to work, the signer must give out at most one preimage (signature) of a given point (message). We simply assume the signer is stateful and will produce the same signature if any message is queried multiple times. This assumption can be remove using a PRF (or the random oracle itself) to implement "repeatable randomness" in a standard way. The PRF is used to generate the random coins of the preimage sampler. If the sampler's randomness complexity is large, this solution may be impractical and the PFDH (Probabilistic Full-Domain Hash) scheme below may be a better option. For simplicity, we describe the stateful version of the scheme.

Stateful Signer. The scheme is built upon a collection of collision-resistant PSFs given by $(\mathsf{TrapGen}, \mathsf{SampleDom}, \mathsf{SamplePre})$ and operates relative to a function $H = H_n : \{0,1\}^* \to R_n$ that is modeled as a random oracle (recall that D_n and R_n are the efficiently-recognizable domain and range, respectively, of the collection for security parameter n).

- $\mathsf{KeyGen}(1^n)$: let $(a,t) \leftarrow \mathsf{TrapGen}(1^n)$, where a describes a function f_a and t is its trapdoor. The verification key is a and the signing key is t.
- $\mathsf{Sign}(t,m)$: if (m,σ_m) is in local storage, output σ_m. Else, let $\sigma_m \leftarrow \mathsf{SamplePre}(t, H(m))$, store (m,σ_m), and output σ_m.
- $\mathsf{Verify}(a,m,\sigma)$: if $\sigma \in D_n$ and $f_a(\sigma) = H(m)$, accept. Else, reject.

Now we prove that the scheme described above is strongly existentially unforgeable under a chosen-message attack.

It is clear that the scheme is complete, by the properties of the trapdoor collection. Assume, for contradiction, that there is an adversary $\mathcal{A}$ that breaks the existential unforgeability of the signature scheme with probability $\varepsilon = \varepsilon(n)$. We construct a poly-time adversary $\mathcal{S}$ that breaks the trapdoor collision-resistant hash function with probability negligibly close to ϵ. Given an index a describing a function $f_a, \mathcal{S}$ runs $\mathcal{A}$ on public key a, and simulates the random oracle H and signing oracle as follows. Without loss of generality, assume that $\mathcal{A}$ queries H on every message m before making a signing query on m.

- For every query to H on a distinct $m \in \{0,1\}^*, \mathcal{S}$ lets $\sigma_m \leftarrow \mathsf{SampleDom}(1^n)$, stores (m,σ_m), and returns $f_a(\sigma_m)$ to $\mathcal{A}$. (If H was previously queried on $m, \mathcal{S}$ looks up (m,σ_m) and returns $f_a(\sigma_m)$.)
- Whenever $\mathcal{A}$ makes a signing query on $m, \mathcal{S}$ looks up (m,σ_m) in its local storage and returns σ_m as the signature.

Now without loss of generality, assume that before outputting its attempted forgery $(m^*,\sigma^*), \mathcal{A}$ queries H on m^*. When $\mathcal{A}$ produces $(m^*,\sigma^*), \mathcal{S}$ looks up (m^*,σ_{m^*}) in its local storage and outputs (σ^*,σ_{m^*}) as a collision in f_a.

We now analyze the reduction. First, we claim that the view of $\mathcal{A}$ in the real chosen-message attack is identical to its view as provided by $\mathcal{S}$. (This assumes that the properties in (*) are perfect; if they are only statistical, the views are statistically close.) For each distinct query m to H, the value returned by $\mathcal{S}$ is $f_a(\sigma_m)$ where $\sigma_m \leftarrow \mathsf{SampleDom}(1^n)$; by the "uniform output" property of the collection, this is identical to the uniformly random value of $H(m) \in R_n$ in the real system. Now fix the value $H(m)$. Then for every signature query on the message $m, \mathcal{S}$ returns a single value σ_m which is distributed as SampleDom (1^n), given $f_a(\sigma_m) = H(m)$. In the real system, signature queries on m (even repeated ones) are answered by a single value having the same distribution, by the preimage sampleability of SamplePre.

Therefore $\mathcal{A}$ outputs a valid forgery (m^*, σ^*) with probability (negligibly close to) ε. Because σ^* is a valid signature on m^*, we have $\sigma^* \in D_n$ and $f_a(\sigma^*) = H(m^*) = f_a(\sigma_{m^*})$. It simply remains to check that $\sigma^* \neq \sigma_{m^*}$, i.e., that they form a collision in f_a. There are two cases to consider:

1. If $\mathcal{A}$ made a signature query on m^*, it received back the signature σ_{m^*}. Because (m^*, σ^*) is considered a forgery, we have $\sigma^* \neq \sigma_{m^*}$.

2. If $\mathcal{A}$ did not make a signature query on m^*, then for the query to H on $m^*, \mathcal{S}$ stored a tuple (m^*, σ_{m^*}) for $\sigma_{m^*} \leftarrow \mathsf{SampleDom}(1^n)$, and returned $f_a(\sigma_{m^*})$ to $\mathcal{A}$. By the preimage min-entropy property of the hash family, the min-entropy of σ_{m^*} given $f_a(\sigma_{m^*})$ (and the rest of the view of $\mathcal{A}$, which is independent of σ_{m^*}) is $\omega(\log n)$. Thus, the signature $\sigma^* \neq \sigma_{m^*}$, except with negligible probability $2^{-\omega(\log n)}$.

We conclude that $\mathcal{S}$ outputs a valid collision in f_a with probability negligibly close to ε.

Probabilistic FDH. The PFDH scheme replaces the statefulness of FDH with a random "salt" for each signature, and is parameterized by the length k of the salt (for simplicity, we can set $k = n$, though any $k = \omega(\log n)$ will suffice for asymptotic security).

- $\mathsf{KeyGen}(1^n)$: let $(a, t) \leftarrow \mathsf{TrapGen}(1^n)$, where a describes a function f_a and t is its trapdoor. The verification key is a and the signing key is t.

- $\mathsf{Sign}(t, m)$: choose $r \leftarrow \{0,1\}^k$ at random, let $\sigma \leftarrow \mathsf{SamplePre}(t, H(m \| r))$, and output (r, σ).

- $\mathsf{Verify}(a, m, (r, \sigma))$: if $\sigma \in D_n$ and $r \in \{0,1\}^k$ and $f_a(\sigma) = H(m \| r)$, then accept. Else, reject.

We now prove that the scheme described above is strongly existentially unforgeable under a chosen-message attack.

The proof is almost identical the prior one, so we simply describe the main idea. Security can be based on either collision-resistance as above (which can be based on the hardness of SIS), or on claw-free pairs (which can be based on the hardness of ISIS). The essential idea is that repeated signature queries on the same message m will all have distinct salts r (except with negligible probability $Q_{\mathsf{sign}}^2 / 2^k$), so the signer will provide a preimage for independent hash values $H(m \| r)$.

Problem 3.

Part A.

- $\mathsf{SampleDom}(1^n)$: Using [GPV08] to sample $\mathbf{x} \leftarrow D_{\mathbb{Z}^n,s,\vec{0}}$. If $\mathbf{x} \notin \mathcal{D}$ then sample again, by the inequality in the hint we know the success probability is overwhelming.
- $\mathsf{SamplePre}(\mathbf{B}, \mathbf{y})$: First find an e_0 such that $\mathbf{A}\mathbf{e}_0 = \mathbf{y} \mod q$. Then sample $\mathbf{x}' \leftarrow D_{\Lambda(\mathbf{B}),s,-\mathbf{e}_0}$ and return $\mathbf{e}_0 + \mathbf{x}'$.

Next we prove the distributions in (*) are indistinguishable.

Consider function $f_{\mathbf{A}} : \{0,1\}^{m \log q} \to \{0,1\}^{n \log q}$. Denote this funciton family triggered by all possible $\mathbf{A}$ as $\mathcal{H}$. Obviously we have

$$\Pr[h(\mathbf{x}) = h(\mathbf{y})] = \Pr_{\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}^{n \times m}}[\mathbf{A}(\mathbf{x} - \mathbf{y}) = 0] = \frac{1}{q^n}, \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{Z}^m, \mathbf{x} \neq \mathbf{y}.$$

According to the hint, $\forall \mathbf{x} \in \{0,1\}^{n \log q}$, $\Pr[X = \mathbf{x}] \leq \frac{1}{2^m}$, where $X \sim \mathsf{SampleDom}$. Let $\varepsilon = \frac{1}{q^{2n}} = \text{negl}(n)$, then one can verify that $n \log q = m - 2 \log \frac{1}{\varepsilon}$. Then by Leftover Hash Lemma,

$$d_{TV}((H, H(X)), (H, U)) \leq \frac{\varepsilon}{2} = \text{negl}(n).$$

That is to say, $(\mathbf{A}, \mathbf{A}\mathbf{x})$ and $(\mathbf{A}, \mathbf{u})$ are indistinguishable, where $\mathbf{A} \overset{\$}{\leftarrow} \mathbb{Z}^{n \times m}$, $x \overset{\$}{\leftarrow} \mathsf{SampleDom}(1^n)$, and $\mathbf{u} \overset{\$}{\leftarrow} \mathbb{Z}_q^n$.

It remains to show that the distributions of $x|y$ on each side of (*) are identical. This follows from the simple equality below:

$$\frac{\exp(-\frac{\pi}{s^2}\|\mathbf{x}\|^2)}{\sum_{\mathbf{x}' \in D, \mathbf{y} = \mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'} \exp(-\frac{\pi}{s^2}\|\mathbf{x}'\|^2)} = \frac{\exp(-\frac{\pi}{s^2}\|(\mathbf{x} - \mathbf{e}_0) - (-\mathbf{e}_0)\|^2)}{\sum_{\mathbf{x} + \mathbf{e}_0 \in D, \mathbf{A}\mathbf{x}' = 0} \exp(-\frac{\pi}{s^2}\|\mathbf{x}' + \mathbf{e}_0\|^2)}.$$

Part B. One-wayness is implied by collision resistance.

Assume there is a P.P.T algorithm $\mathcal{A}$ that finds collision of $f_{\mathbf{A}}$. Construct an algorithm $\mathcal{A}'$ to solve $\text{SIS}_{q,s,2s\sqrt{m}}$ as follows:

$\mathcal{A}'(1^n, \mathbf{A})$

- Run $\mathcal{A}(1^n, \mathbf{A})$, get output $(\mathbf{x}, \mathbf{x}')$.
- Return $\mathbf{e} = \mathbf{x} - \mathbf{x}'$.

Once $\mathcal{A}$ wins, since $\mathbf{x} \neq \mathbf{x}'$, we have $0 < \|\mathbf{e}\| \leq 2s\sqrt{m}$ and $\mathbf{A}\mathbf{e} = \mathbf{0}$. By assumption, $\text{SIS}_{q,s,2s\sqrt{m}}$ is hard. Thus $\mathcal{A}$ wins with negligible probability.