1 QIP

In this section, we prove that QIP is at least as hard as QR (quadratic residuosity); one way to prove this is to reduce QR to QIP; technically speaking, in doing that we might face the following subproblem: prove that the two distributions below are identical or evaluate/discuss their statistical/computational difference.

Let $s \leftarrow S$ denote the random process of uniformly and independently choosing an element s from the set S. Let $z \leftarrow \mathcal{A}(x, y, \dots)$ denote the output z of an algorithm $\mathcal{A}$ with input $(x, y, \dots)$. Given a security parameter m, let $\mathrm{BW}(1^m)$ denote the set of m-bit primes p such that $p = 3 \bmod 4$. For a positive integer $n > 1$, let $\mathbb{Z}_n^*$ denote the set of positive integers which are less than and coprime to n. Let $QR(n)$ denote the set of quadratic residues modulo n and let $\mathbb{Z}_n^{+1}$ (resp. $\mathbb{Z}_n^{-1}$) be the elements of $\mathbb{Z}_n^*$ which have Jacobi symbol equal to $+1$ (resp. -1). Let $\mathrm{CS}(n, \alpha)$ denote the set of integers s in $\mathbb{Z}_n^*$ which satisfy the condition α. Also, we say that a function $f : \mathbb{Z}^+ \to [0, 1]$ is *negligible in n* if $f(n) < 1/p(n)$ for any polynomial p and sufficiently large n.

Quadratic Residuosity Problem (QRP). The QRP *problem* consists of efficiently distinguishing the following two distributions:

$$
\begin{aligned}
E_0(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; \\
&\qquad s \leftarrow \mathbb{Z}_n^{-1} \cup QR(n) \;:\; (n, s)\} \\
&= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; \\
&\qquad s \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^{-1} \cup QR(n)) \;:\; (n, s)\} \\
E_1(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; \\
&\qquad s \leftarrow \mathbb{Z}_n^* \;:\; (n, s)\} \\
&= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; \\
&\qquad s \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^*) \;:\; (n, s)\}\,.
\end{aligned}
$$

We say that algorithm $\mathcal{A}$ has *advantage* ϵ in solving QRP if we have that:

$$
\begin{aligned}
\Big| \Pr[(n, s) \leftarrow E_0(1^m) : \mathcal{A}(n, s) = 1] \\
- \Pr[(n, s) \leftarrow E_1(1^m) : \mathcal{A}(n, s) = 1] \Big| = \epsilon. \quad (1)
\end{aligned}
$$

We say that QRP is *intractable* if all polynomial time (in m) algorithms have a negligible (in m) advantage in solving QRP.

Quadratic Indistinguishability Problem (QIP). The QIP *problem* consists of efficiently distinguishing the following two distributions:

$$\begin{aligned} D_0(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 - 4h \in \mathbb{Z}_n^{-1} \cup QR(n)) \,:\, (n, h, s)\} \\ D_1(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathbb{Z}_n^* \,:\, (n, h, s)\} \\ &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^*) \,:\, (n, h, s)\}\,. \end{aligned}$$

We say that algorithm $\mathcal{A}$ has *advantage* ϵ in solving QIP if we have that:

$$\begin{aligned} \Big| \Pr[(n, h, s) \leftarrow D_0(1^m) : \mathcal{A}(n, h, s) = 1] \\ - \Pr[(n, h, s) \leftarrow D_1(1^m) : \mathcal{A}(n, h, s) = 1] \Big| = \epsilon. \end{aligned} \quad (2)$$

We say that QIP is *intractable* if all polynomial time (in m) algorithms have a negligible (in m) advantage in solving QIP.

Before proving the equivalence of QRP and QIP, we prove the equivalence of QIP and QIP_0.

QIP_0 *Problem.* The QIP_0 *problem* consists of efficiently distinguishing the following two distributions:

$$\begin{aligned} D_{0,0}(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 - 4h \in \mathbb{Z}_n^{-1} \cup QR(n)) \,:\, (n, h, s)\} \\ D_{0,1}(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 - 4h \in \mathbb{Z}_n^*) \,:\, (n, h, s)\}\,. \end{aligned}$$

Recall that given a finite set C and any two subsets A and B, we have

$$|A \cup B| \leq |C| \quad \text{and} \quad |A \cap B| \geq |A| + |B| - |C|$$

so that

$$|A \Delta B| = |A \cup B| - |A \cap B| \leq |C| - (|A| + |B| - |C|) = 2|C| - |A| - |B|\,. \quad (3)$$

Since addition by a non-zero element is a 1-1 map on $\mathbb{Z}_n$, we have

$$|4h + \mathbb{Z}_n^*| = |\mathbb{Z}_n^*| = (p-1)(q-1)$$

so that, by Equation 3,

$$|4h + \mathbb{Z}_n^* \Delta \mathbb{Z}_n^*| \leq 2pq - (p-1)(q-1) - (p-1)(q-1) = 2(p+q-1)\,.$$

Thus the advantage of an adversary distinguishing the two distributions

$$\begin{aligned} D'_1(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^*) \,:\, (n, h, s)\} \\ D'_{0,1}(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s \in 4h + \mathbb{Z}_n^*) \,:\, (n, h, s)\} \end{aligned} \tag{4}$$

is at most

$$\frac{2(p+q-1)}{(p-1)(q-1)}.$$

In particular, the advantage of an adversary distinguishing the two distributions

$$\begin{aligned} D''_1(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 \in \mathbb{Z}_n^*) \,:\, (n, h, s)\} \\ D''_{0,1}(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 \in 4h + \mathbb{Z}_n^*) \,:\, (n, h, s)\} \end{aligned} \tag{5}$$

is at most

$$4 \cdot \frac{2(p+q-1)}{(p-1)(q-1)/4} = \frac{8(p+q-1)}{(p-1)(q-1)}.$$

Now,

$$s^2 - 4h \in \mathbb{Z}_n^* \iff s^2 \in 4h + \mathbb{Z}_n^* \quad \text{and} \quad s \in \mathbb{Z}_n^* \iff s^2 \in \mathbb{Z}_n^*$$

so that the advantage of an adversary distinguishing the two distributions

$$\begin{aligned} D_1(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^*) \,:\, (n, h, s)\} \\ D_{0,1}(1^m) &= \{p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h \leftarrow \mathbb{Z}_n^{+1}; \\ &\qquad s \leftarrow \mathrm{CS}(n, s^2 - 4h \in \mathbb{Z}_n^*) \,:\, (n, h, s)\} \end{aligned} \tag{6}$$

is at most

$$\frac{8(p+q-1)}{(p-1)(q-1)}.$$

Hence $D_{0,1}(1^m)$ and $D_1(1^m)$ are indistinguishable and the problems QIP and QIP_0 are equivalent (since $D_{0,0}(1^m)$ and $D_0(1^m)$ are the same distribution).

So, to prove the equivalence of QRP and QIP, we only need to prove the equivalence of QRP and QIP_0. Given a distinguisher $\mathcal{A}$ for QIP_0, which takes as input (n, h, s) where n is a Blum-Williams integer, $h \leftarrow \mathbb{Z}_n^{+1}$ and $s \in \mathbb{Z}_n^*$, and outputs a bit b indicating (n, h, s) is from the distribution $D_{0,b}(1^m)$, we simulate a distinguisher $\mathcal{B}$ for QRP, which takes as input (n, s) where n is a Blum-Williams integer and $s \in \mathbb{Z}_n^*$, and outputs a bit b indicating (n, s) is from the distribution $E_b(1^m)$ as follows

1. $\mathcal{B}$ randomly chooses $\sigma \leftarrow \mathbb{Z}_n^*$ and sets $h = \frac{\sigma^2 - s}{4}$.

2. If $h \notin \mathbb{Z}_n^{+1}$, $\mathcal{B}$ discards this σ and goes back to the previous step.
3. $\mathcal{B}$ invokes $\mathcal{A}$ and inputs (n, h, σ) to it.
4. $\mathcal{B}$ outputs $b \leftarrow \mathcal{A}(n, h, \sigma)$.

Since $s = \sigma^2 - 4h$, $(n, h, \sigma) \leftarrow D_{0,b}(1^m)$ implies that

if $b = 0$, $\sigma \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^{-1} \cup QR(n))$;
else if $b = 1$, $\sigma \leftarrow \mathrm{CS}(n, s \in \mathbb{Z}_n^*)$.

Thus $\mathcal{B}$ is correct whenever $\mathcal{A}$ is. Hence QIP_0 is equivalent to QRP.

QIP_1 *Problem.* We define the QIP_1 *problem* as the problem of efficiently distinguishing the following two distributions:

$$\begin{aligned} D_{1,0}(1^m) = \{&p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; d \leftarrow \{0,1\}; h_0, h_1 \leftarrow \mathbb{Z}_n^{+1}; \\ &s \leftarrow \mathrm{CS}(n, s^2 - 4h_d \in \mathbb{Z}_n^{-1} \cup QR(n)) \,:\, (n, h_0, h_1, s)\} \\ D_{1,1}(1^m) = \{&p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h_0, h_1 \leftarrow \mathbb{Z}_n^{+1}; s \leftarrow \mathbb{Z}_n^* \,:\, (n, h_0, h_1, s)\} \end{aligned}$$

We say that algorithm $\mathcal{A}$ has *advantage* ϵ in solving QIP_1 if we have that:

$$\begin{aligned} \big| \Pr[&(n, h_0, h_1, s) \leftarrow D_{1,0}(1^m) : \mathcal{A}(n, h_0, h_1, s) = 1] \\ &- \Pr[(n, h_0, h_1, s) \leftarrow D_{1,1}(1^m) : \mathcal{A}(n, h_0, h_1, s) = 1] \big| = \epsilon. \end{aligned} \tag{7}$$

We say that QIP_1 is *intractable* if all polynomial time (in m) algorithms have a negligible (in m) advantage in solving QIP_1.

By a simple simulation argument, we can prove the following theorem:

Theorem 1. The QIP_1 problem is intractable if and only if the QIP problem is so.

To prove the equivalence of QIP and QIP_1, given D_0 and D_1, we choose randomly $h_1 \leftarrow \mathbb{Z}_n^{+1}$ and create $D_{1,0}$ and $D_{1,1}$. If we can distinguish between these two with probability ϵ, then with prob $\epsilon/2$ we can distinguish between the given two distribution.

QIP_2 *Problem.* We define the QIP_2 *problem* as the problem of efficiently distinguishing the following two distributions:

$$\begin{aligned} D_{2,0}(1^m) = \{&p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h_0, h_1 \leftarrow \mathbb{Z}_n^{+1}; \\ &s \leftarrow \mathrm{CS}(n, s^2 - 4h_0 \in \mathbb{Z}_n^{-1} \cup QR(n)) \,:\, (n, h_0, h_1, s)\} \\ D_{2,1}(1^m) = \{&p, q \leftarrow \mathrm{BW}(1^m); n \leftarrow p \cdot q; h_0, h_1 \leftarrow \mathbb{Z}_n^{+1}; \\ &s \leftarrow \mathrm{CS}(n, s^2 - 4h_1 \in \mathbb{Z}_n^{-1} \cup QR(n)) \,:\, (n, h_0, h_1, s)\} \end{aligned}$$

We say that algorithm $\mathcal{A}$ has *advantage* ϵ in solving QIP_2 if we have that:

$$\begin{aligned} \big| \Pr[&(n, h_0, h_1, s) \leftarrow D_{2,0}(1^m) : \mathcal{A}(n, h_0, h_1, s) = 1] \\ &- \Pr[(n, h_0, h_1, s) \leftarrow D_{2,1}(1^m) : \mathcal{A}(n, h_0, h_1, s) = 1] \big| = \epsilon. \end{aligned} \tag{8}$$

We say that QIP_2 is *intractable* if all polynomial time (in m) algorithms have a negligible (in m) advantage in solving QIP_2.

By a simple hybrid argument, we can prove the following theorem:

Theorem 2. The QIP_2 problem is intractable if and only if the QIP problem is so.

To prove the equivalence of QIP and QIP_2, given D_0 and D_1, we choose randomly $h' \leftarrow \mathbb{Z}_n^{+1}$ and create D_1'. Then we treat D_1 and D_1' as $D_{2,0}$ and $D_{2,1}$. If we can distinguish between these two then we contradict the indistinguishability between D_0 and D_1 and between D_0 and D_1'.