7. Appendix - Proofs

7.1. Proof of Lemma 1

By Condition 1, we know that for any $I \subseteq [n]$, $|I| = n - s$, we have $\mathbf{1} \in \text{span}\{b_i \mid i \in I\}$. In other words, there exists atleast one $x \in \mathbb{R}^{(n-s)}$ such that:

$$xB(I,:) = \mathbf{1} \tag{11}$$

Therefore, by construction, we have: $AB = \mathbf{1}_{\binom{n}{s}\times n}$, and the scheme (A, B) is robust to **any** s stragglers.

7.2. Proof of Theorem 1

Consider any scheme (A, B) robust to **any** s stragglers, with $B \in \mathbb{R}^{n\times k}$. Now, construct a bipartite graph between n workers, $\{W_1, \dots, W_n\}$, and k partitions, $\{P_1, \dots, P_k\}$, where we add an edge (i, j) if worker i and partition j is worker i has access to partition j. In other words, for any $i \in [n], j \in [k]$:

$$e_{ij} = \begin{cases} 1 & \text{if } B(i,j) \neq 0 \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

Now, it is easy to see that the degree of the i^{th} worker W_i is $\|b_i\|_0$.

Also, for any partition P_j, its degree must be at least $(s + 1)$. If its degree is s or less, then consider the scenario where all its neighbors are stragglers. In this case, there is no non-straggler worker with access to P_j, which contradicts robustness to **any** s stragglers.

Based on the above discussion, and using the fact that the sum of degrees of the workers in the bipartite graph must be the same as the sum of degrees of partitions, we get:

$$\sum_{i=1}^{n} \|b_i\|_0 \geq k(s+1) \tag{13}$$

Since we assume all workers get access to the same number of partitions, this gives:

$$\|b_i\|_0 \geq \frac{k(s+1)}{n}, \text{ for any } i \in [n] \tag{14}$$

7.3. Proof of Theorem 2

Consider groups of partitions $\{G_1, \dots, G_{n/(s+1)}\}$ as follows:

$$\begin{aligned} G_1 &= \{P_1, \dots, P_{s+1}\} \\ G_2 &= \{P_{s+2}, \dots, P_{2s+2}\} \\ &\vdots \qquad (15) \\ G_{n/(s+1)} &= \{P_{n-s}, \dots, P_n\} \qquad (16) \end{aligned}$$

Fix some set $I \subseteq [n], |I| = n - s$. Based on our construction, it is easy to observe that for any group G_j, there exists some index in I, say $i_{G_j} \in I$, such that the corresponding row in B, $b_{i_{G_j}}$ has all 1s at partitions in G_j and 0s elsewhere. This is because there are $(s + 1)$ rows of B that correspond in this way to G_j (one in each block $\overline{B}_{\text{block}}$), and so atleast one would survive in the set I of cardinality $(n - s)$. Now, it is trivial to see that:

$$\mathbf{1} \in \text{span}\{b_{i_{G_j}} \mid j = 1, \dots, n/(s+1)\} \tag{17}$$

Also, since

$$\text{span}\{b_{i_{G_j}} \mid j = 1, \dots, n/(s+1)\} \subseteq \text{span}\{b_i \mid i \in I\}, \tag{18}$$

we have $\mathbf{1} \in \text{span}\{b_i \mid i \in I\}$.

Finally, since the above holds for any set I, we get that B satisfies Condition 1. The remainder of the theorem follows from Lemma 1.

7.4. Proof of Theorem 3

Consider the subspace given by the null space of the random matrix matrix H (constructed in Algorithm 2):

$$S = \{x \in \mathbb{R}^n \,|\, Hx = 0\} \tag{19}$$

Note that H has $(n-1)s$ different random values (s for each column), since its last column is simply the negative sum of its previous $(n-1)$ columns. Now, we have the following Lemma listing some properties of H and S.

Lemma 2. *Consider $H \in \mathbb{R}^{ss \times n}$ as constructed in Algorithm 2, and the subspace S as defined in Eq. **??**. Then, the following hold:*

- *Any s columns of H are linearly independent with probability 1*
- *$dim(S) = n - s$ with probability 1*
- *$\mathbf{1} \in S$, where $\mathbf{1}$ is the all-ones vector*

For $i \in [n]$, let S_i denote the set $S_i = \{i \mod n, (i+1) \mod n, \ldots, (i+s) \mod n\}$. Then, S_i corresponds to the support of the i^{th} row of B in our construction, as also given by the support structure in Eq. (10).

Recall that we denote the i^{th} row of B by b_i. By our construction, we have:

$$\begin{aligned} b_i(i) &= 1 \\ b_i(S_i \setminus \{i\}) &= -H^{-1}_{S_i \setminus \{i\}} H_i \end{aligned} \tag{20}$$

Now, we have the following lemma;

Lemma 3. *Consider the i^{th} row of B constructed using Algorithm 2 (also shown in Eq. **??**). Then,*

- *$b_i \in S$*
- *Every element of $b_i(S_i \setminus \{i\})$ is non-zero with probability 1*
- *For any subset $I \subseteq [n]$, $|I| = n - s$, the set of vectors $\{b_i \,|\, i \in I\}$ is linearly independent with probability 1*

Now, using Lemma **??**, we can conclude that for any subset $I \subseteq [n]$, $|I| = n - s$, $dim\left(\text{span}\{b_i \,|\, i \in I\}\right) = n - s$ and $\text{span}\{b_i \,|\, i \in I\} \subseteq S$. Consequently, from Lemma **??**, since $dim(S) = n - s$ and $\mathbf{1} \in S$, this implies that:

$$\text{span}\{b_i \,|\, i \in I\} = S \text{ with probability 1} \tag{21}$$

and, $\mathbf{1} \in \text{span}\{b_i \,|\, i \in I\}$. Taking union bound over every I shows that B satisfies Condition 1. The remainder of the theorem follows from Lemma 1.

7.4.1. Proof of Lemma ??

Consider any subset $I \subseteq n$, $|I| = s$ such that $n \notin I$. Then, all the elements of H_I are independent, and $det(H_I)$ is a polynomial in the elements of H_I. Consequently, since every element is drawn from a continuous probability distribution (in particular, gaussian), the set $\{H_I \,|\, det(H_I) = 0\}$ is a zero measure set. So, $P\left(det(H_I) \neq 0\right) = 1$, and thus the columns of H_I are linearly independent with probability 1.

If $n \in I$, then we have:

$$det(H_I) = det(\widetilde{H}) \tag{22}$$

where we let $\widetilde{H} = \left[H_{I \setminus \{n\}}, -\sum_{i \in [n] \setminus I} H_i\right]$. The elements of $\widetilde{H}$ are independent, so using the same argument as above, we again have $P(det(H_I) = det(\widetilde{H}) \neq 0) = 1$. Finally, taking a union bound over all sets I of cardinality s shows that any s columns of H are linearly independent.

Since any s columns in H are linearly independent, this implies that $rank(H) = s$. Since the subspace S is simply the null space of H, we have $dim(S) = n - s$.

Finally, since $H_n = -\sum_{i \in [n-1]} H_i$ (by construction), we have $H\mathbf{1} = 0$ and thus $\mathbf{1} \in S$.

7.4.2. PROOF OF LEMMA ??

By construction of b_i, we have:

$$Hb_i = H_i + H_{S_i \setminus \{i\}} b_i(S_i \setminus \{i\}) = H_i - H_i = 0 \tag{23}$$

Thus, $b_i \in S$.

Now, if possible, let for some $k \in S_i \setminus \{i\}$, $b_i(k) = 0$. Then, since $b_i \in S$, we have:

$$Hb_i = H_i + H_{S_i \setminus \{i,k\}} b_i(S_i \setminus \{i,k\}) = 0 \tag{24}$$

Consequently, the set of columns $\{j \mid j \in S_i \setminus \{i,k\}\} \cup \{i\}$ is linearly dependent which contradicts H having any s columns being linearly independent (in Lemma **??**). Therefore, we must have every element of $b_i(S_i \setminus \{i\})$ being non-zero.

Now, consider any subset $I \subseteq [n], |I| = n - s$. We shall show that the matrix B_I (corresponding to the rows of B with indices in I) has rank $n-s$ with probability 1. Consequently, the set of vectors $\{b_i \mid i \in I\}$ would be linearly independent. To show this, we consider some $n-s$ columns of B_I, say given by the set $J \subseteq [n], |J| = n-s$, and denote the sub-matrix of columns by $B_{I,J}$. Then, it suffices to show that $\det(B_{I,J}) \neq 0$. Now, by the construction in Algorithm 2, we have: $\det(B_{I,J}) = \text{poly}_1(H)/\text{poly}_2(H)$, for some polynomials $\text{poly}_1(\cdot)$ and $\text{poly}_2(\cdot)$ in the entries of H. Therefore, if we can show that there exists at least one H' with $H'\mathbf{1} = \mathbf{0}$ and $\text{poly}_1(H')/\text{poly}_2(H') \neq 0$, then under a choice of i.i.d. standard gaussian entries of H, we would have:

$$\mathbb{P}\left(\text{poly}_1(H)/\text{poly}_2(H) \neq 0\right) = 1 \tag{25}$$

The remainder of this proof is dedicated to showing that such an H' exists. To show this, we shall consider a matrix $\widetilde{B} \in \mathbb{R}^{n-s \times n}$ such that $\text{supp}(\widetilde{B}) = \text{supp}(B_I)$ and $\det(\widetilde{B}_{:,J}) \neq 0$, where $\widetilde{B}_{:,J}$ corresponds to the sub-matrix of $\widetilde{B}$ with columns in the set J. Given such a $\widetilde{B}$, we shall show that there exists an $s \times n$ matrix H' (with $H'\mathbf{1} = \mathbf{0}$) such that when we run Algorithm 2 with this H', we get a matrix B' s.t. $B'_I = \widetilde{B}$ *i.e.* the output matrix from Algorithm 2 is identical to our random choice $\widetilde{B}$ on the rows in the set I. This suffices to show the existence of an H' such that $\text{poly}_1(H')/\text{poly}_2(H') \neq 0$, since $\text{poly}_1(H')/\text{poly}_2(H') = \det(B'_{I,J}) = \det(\widetilde{B}_J) \neq 0$.

Let us pick a random matrix $\widetilde{B}$ as:

$$\widetilde{B} = B_I^r D \tag{26}$$

where B_I^r is a matrix with the same support as B_I and with each non-zero entry i.i.d. standard gaussian, and D is a diagonal matrix such that $D_{ii} = \sum_{j=1}^{n-s} B_I^r(j,i),\ i \in [n]$. Note that a consequence of the above choice of $\widetilde{B}$ is that the sum of all its rows is the all $\mathbf{1}$s vector. Now, it can be shown that any $(n-s)$ columns of $\widetilde{B}$ form an invertible sub-matrix with probability 1. Let S_i be the support of the i^{th} row of B. The rows of B_I^r have the supports $S_i, i \in I$. Now because of the cyclic support structure in B, any collection $\{i_1, i_2, \ldots, i_k\} (0 \leq k \leq n-s)$ satisfies the property:

$$|\cup_{j=1}^{k} S_{i_j}| \geq s + k \tag{27}$$

Using Lemma 4 in (**?**), this implies that there is a perfect matching between the rows of B_I^r and any of its $(n-s)$ columns . Consequently, with probability 1, any $(n-s)$ columns of B_I^r form an invertible sub-matrix. Also, since every column of B_I^r contains atleast one non-zero (again, owing to the support structure of B), this implies that with probability 1, all the diagonal entries of D are non-zero. Combining the above two observations, we can infer that any $(n-s)$ columns of $\widetilde{B}$ form an invertible sub-matrix with probability 1.

So far, we have shown existence of a matrix $\widetilde{B}$ with the following properties: (i) $\widetilde{B}$ has the same support structure as B_I, (ii) any $(n-s)$ columns of $\widetilde{B}$ form invertible sub-matrix, (iii) the sum of all rows of $\widetilde{B}$ is the all $\mathbf{1}$s vector. Now, for any such $\widetilde{B}$, we shall show that there exists an H' such that $H'\widetilde{B}^T = \mathbf{0}$ such that any s columns of H' form an invertible sub-matrix. This implies that when we run Algorithm 2 with this H', the output matrix would be the same as $\widetilde{B}$ on the rows in the set I. The remainder of the proof then follows from our earlier discussion.

Now, consider any set $Q \subseteq [n], |Q| \leq s$. Suppose we pick any invertible $H'_{:,Q}$, and set $H'_{:,[n]\setminus Q} = -H'_{:,Q} \widetilde{B}^T_{:,Q} (\widetilde{B}^T_{:,[n]\setminus Q})^{-1}$. Then, such an H' satisfies $H'\widetilde{B}^T = 0$ and its columns in the set Q form an invertible sub-matrix. Now, since invertibility on the set Q simply corresponds to $\det(H'_{:,Q}) \neq 0$ (*i.e.* some fixed polynomial being non-zero), if we actually picked a uniformly random H' on the subspace $H'\widetilde{B}^T = 0$, then

$$\mathbb{P}\left(\det(H'_{:,Q}) \neq 0 \,|\, H'\widetilde{B}^T = 0\right) = 1 \tag{28}$$

Taking a union bound over all Qs, we get that

$$\mathbb{P}\left(\text{any } s \text{ columns of } H' \text{ form an invertible sub-matrix} \,|\, H'\widetilde{B}^T = 0\right) = 1 \tag{29}$$

Thus, there exists an H' satisfying $H'\widetilde{B}^T = 0$ with any s of its columns forming an invertible sub-matrix. Also, since the sum of all rows of $\widetilde{B}$ is $\mathbf{1}$, this implies $H'\mathbf{1} = \mathbf{0}$.