# The Benefit of Being Flexible in Distributed Computation 

Linqi Song, Sundara Rajan Srinivasavaradhan, Christina Fragouli<br>University of California, Los Angeles CA 90095, USA<br>Email: \{songlinqi,sundar,christina.fragouli\} @ucla.edu


#### Abstract

In wireless distributed computing, networked nodes perform intermediate computations over data placed in their memory and exchange these intermediate values to calculate function values. In this paper we consider an asymmetric setting where each node has access to a random subset of the data, i.e., we cannot control the data placement. The paper makes a simple point: we can realize significant benefits if we are allowed to be "flexible", and decide which node computes which function, in our system. We make this argument in the case where each function depends on only two of the data messages, as is the case in similarity searches. We establish a percolation in the behaviour of the system, where, depending on the amount of observed data, by being flexible, we may need no communication at all.


## I. Introduction

Wireless distributed computing divides a computation over a set of data into tasks that are solved by a cluster of networked nodes. One way to view such systems is as trading computation with communication cost. Such exchanges may enable to operate with lower latency or lower battery consumption; the decision on how to operate is determined by examining the computation and the communication cost in each case. Accordingly, recent work in the literature has started exploring communication vs. computation trade-off [1].

In this paper, we focus on the case of random data placement, where each node has collected a random subset of the data that it can use for processing - we cannot control the original data placement. This is plausible in wireless, where the processing nodes are often mobile, equipped with sensors, and can collect (potentially overlapping) data: we may want to process these data in a time-critical fashion and upload information to a remote server for portable scientific computing; we may want to combine different views of the data (say videos) for virtual reality applications; or we may want time-sensitive processing for distributed tasks, such as distributed online recommendations [2]. In crowd sourcing applications, privacy concerns may also restrict sharing - nodes may be willing to share function values, but not the raw data.

We operate under the framework of MapReduce [3]. The first stage "Map" is intra-node data processing, followed by the data shuffling which collects intermediate computation results with the same "key" to the same computing node, while the second stage "Reduce" computes each function. As observed in [4], [5] and [6], data shuffling is a limiting factor in the runtime performance of many distributed computing tasks.

The asymmetry (in terms of the function-data relationship) in our framework makes clear the following point: we can
gain significant benefits, in terms of communication, by being "flexible", and leveraging on the fact that in distributed computing, often we not only know in advance what the functions we want to compute are, but also decide which of the nodes in our system computes each function. For instance, we may eventually upload the computation outcome to a centralized server, and thus what is important is to efficiently compute all functions - not which node does what. In the symmetric scenario where each function depends on all data, it makes no difference which node computes which function; in contrast, in asymmetric scenarios, we find a significant difference: depending on the amount of data nodes randomly observe, nodes may not need to communicate at all; in contrast, when functions are allocated independently of the random data realization, this is almost never the case. We establish a percolation in the behaviour of the system, where, depending on the amount of observed data, by being flexible, we may need no communication when the probability to observe a data message is above some threshold or we may need significant communication when this probability is below the threshold, even when allowed to use coded broadcast transmissions.
We mathematically make this argument in the case where each function depends on only two of the data, as is the case in similarity searches. As a concrete example, consider the task of finding "common friends", a task that is regularly computed in social networks, like Facebook, where friendships dynamically change. Consider that we have a social network, where the friendships between users are stored as messages: the message $b_{A}$ is user $A$ 's set of friends $(A,\{B, C, D\})$, the message $b_{B}$ is user $B$ 's set of friends $(B,\{A, D, E\})$, similarly, $b_{C}=(C,\{A, E\}), b_{D}=(D,\{A, B, F\}), b_{E}=$ $(E,\{B, C, F\}), b_{F}=(F,\{D, E\})$. These 6 messages can be distributed among computing nodes, for example four nodes as follows: node 1 has $\left\{b_{A}, b_{C}, b_{E}\right\}, 2$ has $\left\{b_{B}, b_{D}, b_{F}\right\}, 3$ has $\left\{b_{B}, b_{E}, b_{F}\right\}$, and 4 has $\left\{b_{A}, b_{C}, b_{D}\right\}$. Assume we want to find the common friends of the following $K=3$ pairs: $\{\{A, B\},\{B, C\},\{D, E\}\}$. The data shuffling scheme needs to gather the messages needed to compute these $K$ "common friends" to some $K$ nodes. For example, node 4 can broadcast $b_{A}$ and $b_{C}+b_{D}$; then node 1 can decode message $b_{D}$ using the transmission $b_{C}+b_{D}$ and its local data $b_{C}$ and then compute the common friends of $\{D, E\}$; similarly, node 2 can decode message $b_{C}$ and compute the common friends of $\{B, C\}$, and node 3 can decode message $b_{A}$ and compute the common friends of $\{A, B\}$. Such a "common friends" type
tasks pervade a number of computing tasks today; for example, in recommender systems [2], to evaluate the similarity between two users, we need to find those liked videos that are common to the two, or the products bought by both.

The paper is organized as follows. Section II introduces our notation and problem formulation; Section III considers the case where we can flexibly decide which node calculates which function, while Section IV looks at the case where we fix the allocation of functions to nodes.

## II. Problem Formulation

Throughout the paper, we will use $[y]$ (with $y$ a positive integer) to denote the set $\{1,2, \ldots, y\}$ and $|Y|$ to denote the cardinality of set $Y$.

We consider a distributed computing system with $n$ computing nodes $c_{1}, c_{2}, \ldots, c_{n}$, and $m$ data messages, $b_{1}, b_{2}, \ldots, b_{m}$ in some finite field $\mathbb{F}_{q}$. We will interchangeably use $b_{j}$ and message $j \in[m]$ to refer to messages, and similarly $c_{i}$ or node $i$ to refer to nodes. We assume that each node $c_{i}$, $i \in[n]$, has been allocated message $b_{j}, j \in[m]$ with allocation probability $p$ independently and uniformly across nodes and messages. Therefore, the relationship between the $m$ messages and the $n$ nodes can be represented by a random bipartite graph $B(m, n, p)$. The messages allocated to node $c_{i}$ are referred to as side information, and indexed by a set $S_{i} \subseteq[m]$.

We assume we have $K \leq n / 2$ computation functions to compute and each function takes two data messages as input, namely, $f_{1}\left(b_{11}, b_{12}\right), f_{2}\left(b_{21}, b_{22}\right), \ldots, f_{K}\left(b_{K 1}, b_{K 2}\right)$, for two different messages $b_{k 1}, b_{k 2} \in\left\{b_{j} \mid j \in[m]\right\}$, for all $k \in[K]$. The functions and their dependency on messages is known in advance. Since we can discard a message $b_{j}$ if no function relies on it, we always assume that each message is useful to compute some function(s). In distributed computation, each function $f_{k}\left(b_{k 1}, b_{k 2}\right)$ can be calculated as

$$
f_{k}\left(b_{k 1}, b_{k 2}\right)=h_{k}\left(v_{k 1}^{k}, v_{k 2}^{k}\right)=h_{k}\left(l_{k}^{1}\left(b_{k 1}\right), l_{k}^{2}\left(b_{k 2}\right)\right)
$$

where $v_{k 1}^{k}=l_{k}^{1}\left(b_{k 1}\right), v_{k 2}^{k}=l_{k}^{2}\left(b_{k 2}\right)$ are intermediate results of the distributed computation. Given the messages in $S_{i}$, node $i$ can calculate the values $v$ for all functions that take attributes in $S_{i}$. We use $\mathcal{V}_{i}$ to denote the set of these values.

The data shuffling phase ensures nodes collect the intermediate values they need by communicating with each other. We assume that a set of nodes make $T$ broadcast transmissions, $x_{1}, x_{2}, \ldots, x_{T}$, where a broadcast transmission made by node $i \in[n]$ is a linear combination of the messages in $\mathcal{V}_{i}$. We will refer to transmissions as uncoded, if they contain only one message (do not use coding).

We say that a function $f_{k}$ is covered by a node $i$ if $i$ has both of the messages $b_{k 1}$ and $b_{k 2}$ and thus can calculate the values $v_{k 1}^{k}$ and $v_{k 2}^{k}$; we say that a function $f_{k}$ is covered by a node $i$ and a set of $h$ transmissions $x_{t 1}, x_{t 2}, \ldots, x_{t h}$ if $i$ can compute $f_{k}$ given its own messages and the partial results it decodes from the transmissions.

We aim to minimize the number of broadcast transmissions $T$, such that there exist a set of $K$ different computing nodes $i_{1}, i_{2}, \ldots, i_{K}$ that can compute functions $f_{1}, f_{2}, \ldots, f_{K}$,
respectively. That is, we decide which function is computed by which of the nodes so as to minimize the communication overhead. In the random bipartite graph formulation, we ask what is the typical behaviour of the minimum number of transmissions $T$ as the parameters $n, m, K$ tend to infinity.

## III. Flexible Function Assignment

In this section, we consider the case where we can select which node calculates which function. We show a "percolation" property: when the allocation probability $p$ is above a threshold, we do not need data shuffling. That is, we can find a "matching" between nodes and functions, such that each node almost surely has the data needed to compute its allocated function, with no communication.

We first define the threshold property as follows.
Definition 1. The function $p_{t h}(m, n, K)$ is a threshold for the random graph instance $B(m, n, p)$ regarding an increasing property $\mathcal{P}$, if

$$
\begin{array}{r}
\lim _{m, n, K \rightarrow \infty} \operatorname{Pr}\{B(m, n, p) \text { satisfies } \mathcal{P}\} \\
= \begin{cases}0, & p / p_{t h}(m, n, K) \rightarrow 0 \\
1, & p / p_{t h}(m, n, K) \rightarrow \infty\end{cases} \tag{1}
\end{array}
$$

Theorem 1. Given $K \leq n / 2, K \leq\binom{ m}{2}$, and $\sqrt{n}=K^{1-\epsilon} o(1)$ for some fixed $\epsilon>0$, the function $p_{t h}(m, n, K)=\sqrt{\frac{\log (K)}{n}}$ is a threshold for the random graph instance $B(m, n, p)$ regarding the property that no data shuffling is needed.

The theorem follows from Lemmas 1 and 2.

## A. Allocation Probability Above Threshold

We here look at the case where no communication is needed, provided we decide which node calculates which function.
Lemma 1. If the allocation probability $p$ satisfies $p=$ $\sqrt{\frac{\log (K)}{n}} \omega(1)$, then data shuffling is not needed almost surely, i.e., with probability at least $1-\left(\frac{1}{K}\right)^{\frac{1}{2} \omega(1)-1}=1-o(1)$, we can find in polynomial time an allocation of functions to nodes, such that each function is covered.

Proof. First, we can see that the probability that a node $i$ has both messages $b_{k 1}$ and $b_{k 2}$ to compute the function $f_{k}$ is $p^{2}$. Then the probability that the function $f_{1}$ is not covered by any of the $n$ nodes (i.e., $b_{11}$ and $b_{12}$ are not cached simultaneously at any node) is:

$$
\operatorname{Pr}\left\{f_{1} \text { is not covered by } n \text { nodes }\right\}=\left(1-p^{2}\right)^{n} .
$$

Assume we assign $f_{1}$ to one node that covers it. The probability that $f_{2}$ is not covered by $n-1$ nodes (i.e., $b_{21}$ and $b_{22}$ are not cached simultaneously at a given $n-1$ nodes) is:

$$
\operatorname{Pr}\left\{f_{2} \text { is not covered by } n-1 \text { nodes }\right\}=\left(1-p^{2}\right)^{n-1}
$$

Similarly, the probability that $f_{k}$ (for $k=3, \ldots, K$ ) is not covered by $n-k+1$ nodes (i.e., $b_{k 1}$ and $b_{k 2}$ are not cached simultaneously at a given $n-k+1$ nodes) is:
$\operatorname{Pr}\left\{f_{k}\right.$ is not covered by $n-k+1$ nodes $\}=\left(1-p^{2}\right)^{n-k+1}$.

Hence, the probability that some function among all $K$ functions cannot be computed directly is

$$
\begin{aligned}
& \operatorname{Pr}\left\{\exists k \in[K], f_{k} \text { cannot be computed directly }\right\} \\
& \leq \operatorname{Pr}\left\{f_{1} \text { is not covered by } n \text { nodes }\right\}+\cdots+ \\
& \operatorname{Pr}\left\{f_{K} \text { is not covered by } n-K+1 \text { nodes }\right\} \\
& \leq\left(1-p^{2}\right)^{n}\left(1+\frac{1}{1-p^{2}}+\ldots+\frac{1}{\left(1-p^{2}\right)^{K-1}}\right) \\
& \leq K\left(1-p^{2}\right)^{n-K} \leq K e^{-\omega(1)(1-K / n) \log (K)} \\
& \leq\left(\frac{1}{K}\right)^{\frac{n-K}{n} \omega(1)-1} \leq\left(\frac{1}{K}\right)^{\frac{1}{2} \omega(1)-1}=o(1) .
\end{aligned}
$$

To allocate functions to nodes, we can create a bipartite graph that lists all $n$ computing nodes on one side and all $K$ functions on the other side, and connect a function with a computing node if and only if the node has both of the messages required to compute this function. We can then in polynomial time solve the maximum bipartite graph matching problem (see [7] and [8], for example).

## B. Allocation Probability Below Threshold

We here calculate the number of uncovered functions in the case where $p<p_{t h}(m, n, K)$ and provide a lower bound on the amount of communication needed.

Lemma 2. Given $\sqrt{n}=o\left(K^{1-\epsilon}\right)$ for some fixed $\epsilon>0$, if the allocation probability satisfies $p=\sqrt{\frac{\log (K)}{n}} o(1)$, then the number of uncovered functions is $K^{1-o(1)} / 2$ almost surely, i.e., with probability at least $1-\exp \left(-\frac{K^{2-o(1)}}{8 n}\right)=1-o(1)$.

Moreover, if the allocation probability $p \leq \sqrt{\frac{1}{n}}$, then the number of uncovered functions is $\frac{K}{2 e}$ almost surely, i.e., with probability at least $1-\exp \left(-\frac{K^{2}(1-o(1))}{8 e^{2} n}\right)=1-o(1)$.
Proof. We define $Y^{C}$ to be the maximum number of functions that are covered by different nodes, i.e., both $b_{k 1}$ and $b_{k 2}$ for the functions $f_{k}$ are in the side information set of some distinct nodes. We also define $Y=K-Y^{C}$ as the minimum number of uncovered functions. From the proof of Lemma 1, we know that the probability that a function $f_{k}$ is not covered is:

$$
\begin{equation*}
\operatorname{Pr}\left\{f_{k} \text { is not covered by } n \text { nodes }\right\}=\left(1-p^{2}\right)^{n} . \tag{2}
\end{equation*}
$$

Therefore, the average number of functions that are not covered is

$$
\begin{align*}
\mathbb{E} Y & \geq K\left(1-p^{2}\right)^{n} \geq K\left(1-\frac{o(1) \log (K)}{n}\right)^{n}  \tag{3}\\
& \geq K e^{-o(1) \log (K)}\left(1-\frac{o(1) \log ^{2}(K)}{n}\right)=K^{1-o(1)} .
\end{align*}
$$

Next, we show that the number $Y$ is distributed tightly around $\mathbb{E} Y$. To show this, we first use a "vertex exposure" process to form a martingale based on the random graph $B(m, n, p)$ [9], [10]. Specifically, we label the $n$ computing nodes in the random bipartite graph as $1,2, \ldots, n$ and denote by $Z_{l}$ the random variable to indicate whether the vertex $l$ is exposed in the random graph, i.e., $Z_{l}=1$ if the $l$-th node is present in the graph and $Z_{l}=0$ otherwise. Define $Y_{l}=\mathbb{E}\left[Y \mid Z_{1}, Z_{2}, \ldots, Z_{l}\right]$ as a sequence of random variables for $l=1,2, \ldots, n$, then $\left\{Y_{l}\right\}$ is a Doob martingale and $Y_{n}=Y$. Observe that when we add (or remove) a node, the minimum number of
uncovered functions can decrease (or increase) at most 1. In fact, after adding some messages to node $i$ 's side information, if $Y$ decreases by 2 , then without counting on node $i$, the number $Y$ will decrease by at least 1 . This contradicts the fact that $Y$ is the minimum number of uncovered functions. We conclude that the function $Y$ is 1-Lipschitz: if two random graph realization $B_{0}$ and $B_{1}$ differ at one vertex $l$, namely, the subgraph induced by vertices of $B_{0}$ other than $l$ is the same as that induced by vertices of $B_{1}$ other than $l$, and the function $Y$ satisfies $\left|Y\left(B_{0}\right)-Y\left(B_{1}\right)\right| \leq 1$. Therefore, we can use the Azuma's inequality

$$
\begin{equation*}
\operatorname{Pr}\{\mathbb{E}[Y]-Y \geq a\} \leq \exp \left(-\frac{a^{2}}{2 n}\right), \text { for } a>0 \tag{4}
\end{equation*}
$$

to get

$$
\begin{align*}
& \operatorname{Pr}\left\{Y \leq 0.5 K^{1-o(1)}\right\} \leq \operatorname{Pr}\left\{\mathbb{E}[Y]-Y \geq 0.5 K^{1-o(1)}\right\} \\
& \leq \exp \left(-\frac{K^{2-o(1)}}{8 n}\right) . \tag{5}
\end{align*}
$$

The first part of the theorem is proved. For the second part of the theorem, we use similar techniques. The average number of functions that are not covered is

$$
\begin{equation*}
\left.\mathbb{E} Y \geq K\left(1-p^{2}\right)^{n} \geq K\left(1-\frac{1}{n}\right)^{n} \geq K e^{-1}\left(1-\frac{1}{n}\right)\right) \tag{6}
\end{equation*}
$$

Thus, we can also bound the number $Y$ using the martingale construction and Azuma's inequality:

$$
\begin{align*}
& \operatorname{Pr}\left\{Y \leq \frac{K}{2 e}\right\} \leq \operatorname{Pr}\left\{\mathbb{E}[Y]-Y \geq \frac{K}{2 e}(1-o(1))\right\} \\
& \leq \exp \left(-\frac{K^{2}(1-o(1))}{8 e^{2} n}\right) \tag{7}
\end{align*}
$$

Note that for the above lemma, the claim holds for $\sqrt{n}=$ $K^{1-\epsilon} O(1)$, for example, $K=\Theta(n)$.

Next, we show that for $p<p_{t h}(m, n, K)$, we need a nontrivial amount of communication in the data shuffling phase, even if we are allowed to use coding.

Theorem 2. For allocation probability $p \leq p_{t h}(m, n, K)$, the optimal number of (coded) broadcast transmissions is almost surely lower bounded by $\Omega\left(T_{u n}\right)$, where $T_{u n}$ is the minimum number of uncoded broadcast transmissions we would need to make, if after the random data placement realization, we broadcast missing messages $b_{j}, j \in[m]$ so that any $K$ out of the $n$ nodes can compute the $K$ functions.

The proof of this theorem can be found in Appendix A. Intuitively, the reason why coding does not help, is because the side information is limited.

## C. Outage Probability

If the allocation probability is too low, it is probable that none of the cluster nodes observes some of the data, and thus the computation of functions that depend on this data becomes impossible. We show that this outage probability is $p_{\text {out }}=$ $\frac{\log (m)}{n}$, where for $p \geq p_{\text {out }}(1+\epsilon)$, the probability that some message is missing is almost surely 0 , while for $p \leq p_{\text {out }}(1-$ $\epsilon$ ), this probability is almost surely 1 , for any fixed $\epsilon>0$.

The probability that a message $b_{j}$ is not cached in one of the $n$ nodes is $(1-p)^{n}$. Then, the probability that a message $b_{j}$ is not missing is $1-(1-p)^{n}$. So the probability that some message is missing is $p_{\text {miss }} \triangleq 1-\left[1-(1-p)^{n}\right]^{m}$. Hence, we have for $p \leq p_{\text {out }}(1-\epsilon)$,

$$
\begin{align*}
p_{m i s s} & \geq 1-\left[1-e^{-n p}\left(1-n p^{2}\right)\right]^{m} \\
& \geq 1-\left[1-m^{-1+\epsilon}(1-o(1))\right]^{m}  \tag{8}\\
& \geq 1-e^{-m^{\epsilon}(1-o(1))}=1-o(1)
\end{align*}
$$

and for $p \geq p_{\text {out }}(1+\epsilon)$,

$$
\begin{align*}
p_{\text {miss }} & \leq 1-\left[1-m(1-p)^{n}\right] \leq m e^{-n p} \\
& \leq m e^{-(1+\epsilon) \log (m)}=\left(\frac{1}{m}\right)^{\epsilon}=o(1) \tag{9}
\end{align*}
$$

## IV. Benefit wrt Fixed Function Assignment

In this section, we calculate how much we have gained by being flexible, as compared to fixing the function assignment. We compare with the case where each function $f_{k}$ is assigned to some $n_{k}=\Theta(1)$ disjoint computing nodes, independently of the data placement. When $n_{k}=1$ for all $k$, $\mathfrak{t}$ we assign each function to one specific node. Let $C \triangleq \sum n_{k} / K$ to be the average number of computing nodes to compute each function.
Theorem 3. Given a random graph instance $B(m, \Theta(K), p)$, the threshold probability satisfies $p_{t h}(m, \Theta(K), K)=1-$ $\left(\frac{\log (K)}{K}\right)^{1 / C}=1-o(1)$ with respect to the property that data shuffling is not needed for fixed function assignments.

Proof. We now consider the case when $p \leq p_{t h}$. It suffices to prove that data shuffling is almost always needed for this case. We observe that the probability that $f_{k}$ is covered by at least one of its corresponding $n_{k}$ worker nodes is

$$
\operatorname{Pr}\left\{f_{k} \text { is covered }\right\}=1-\left(1-p^{2}\right)^{n_{k}}
$$

Now, the probability that all functions can be computed is

$$
\begin{aligned}
& \operatorname{Pr}\left\{\forall k, f_{k} \text { can be computed directly }\right\} \\
& =\prod_{k=1}^{K} \operatorname{Pr}\left\{f_{k} \text { is covered }\right\}=\prod_{k=1}^{K}\left(1-\left(1-p^{2}\right)^{n_{k}}\right) \\
& \left.\stackrel{(a)}{\leq} 1-\sum_{k=1}^{K} \frac{\left(1-p^{2}\right)^{n_{k}}}{K}\right)^{K} \stackrel{(b)}{\leq}\left(1-\left(1-p^{2}\right)^{C}\right)^{K} \\
& \leq\left(1-\left(1-p_{t h}\right)^{C}\right)^{K} \leq\left(1-\frac{\log (K)}{K}\right)^{K} \leq \frac{1}{K}
\end{aligned}
$$

where (a) follows from Arithmetic Mean - Geometric Mean (A.M-G.M) inequality applied to the product term, and (b) follows, again, from A.M-G.M applied to the summation term inside the parenthesis.

## A. Uncoded transmissions

Let us consider that only one node is designated to each function. We would like to compute the number of uncoded transmissions needed in this case.

Theorem 4. Given a random graph instance $B(m, n, p)$, the threshold function satisfies $p_{\text {th }}(m, n, K)=1-\frac{1}{K^{\epsilon}}$ (for any
fixed $\epsilon>0)$ with respect to the property that $\Theta\left(K^{1-\epsilon}\right)$ uncoded transmissions are needed for a successful computation of all functions.

Proof. Observe that for a function $f_{k}$, we may need to make 2, 1 , or 0 uncoded transmissions. The probabilities are $(1-p)^{2}$, $2 p-p^{2}$, and $p^{2}$, since a node needs to have two messages to compute its function. Therefore the expected number of uncoded transmissions is $K\left(2 p-p^{2}+2(1-p)^{2}\right)=K(2-$ $2 p+p^{2}$ ). Using Chernoff's bound, we get that the probability for the number of transmissions $X$ taking a value less than or equal to $(1-\delta) K\left(2-2 p+p^{2}\right)$ (for $\left.\delta \in(0,1)\right)$ is,

$$
\begin{aligned}
& \operatorname{Pr}\left\{X \leq(1-\delta) K\left(2-2 p+p^{2}\right)\right\} \\
& \quad \leq e^{-K\left(1-p+p^{2} / 2\right) \delta^{2} / 2} \leq e^{-\Theta\left(K^{1-\epsilon}\right)}
\end{aligned}
$$

## B. The Benefit of Being Flexible

We now summarize and compare our results for the flexible and fixed function assignments.

1) $p<p_{\text {out }}$ : in this regime, the distributed computing task is not possible irrespective of function assignments.
2) $p_{\text {out }}<p=\sqrt{\frac{\log (K)}{n}} o(1)$ : in this regime, we may need significant amount of communication, but there is no significant coding gain for flexible function assignments. The uncoded transmissions can be as large as $\Theta(K)$.
3) $\sqrt{\frac{\log (K)}{n}} \omega(1)=p<1-\frac{\log (K)}{K}$ : this regime spans most of the allocation probability region, ranging from $o(1)$ to $1-o(1)$, and underscores the benefit of the extra degree of freedom (in the form of flexibility in function assignment). A fixed function assignment scheme almost surely needs $\Theta(K)$ transmissions for constant allocation probability, while flexibility in function assignment almost surely requires no communication.

## V. Related Work

Recently, as distributed mobile computing is increasingly attracting interest, use of coding techniques in distributed computing is also gaining attention [11], [12], [13], [1], [14], [15]. Closer to ours are works that help improve the communication efficiency in distributed systems. The works of [13], [14], [15] have studied the "master-workers" distributed computing model and show that coding can help to reduce communication cost in data shuffling, for example, in distributed machine learning applications. The work in [12], [1] has established that leveraging of broadcasting and coding can improve the communication efficiency of such systems. However, creating coding opportunities is enabled by allocating data in specific patterns across the nodes [1] (in fact, similar data allocations have also been used for distributed caching [16]). In contrast, in our work we assume random data allocation, and thus the trade-off curve in [1] does not apply.

## References

[1] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "Fundamental tradeoff between computation and communication in distributed computing," in 2016 IEEE International Symposium on Information Theory (ISIT), 2016, pp. 1814-1818.
[2] F. Ricci, L. Rokach, and B. Shapira, Introduction to recommender systems handbook. Springer, 2011.
[3] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," Communications of the ACM, vol. 51, no. 1, pp. 107113, 2008.
[4] M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica, "Managing data transfers in computer clusters with orchestra," in ACM SIGCOMM Computer Communication Review, vol. 41, no. 4, 2011, pp. 98-109.
[5] F. Ahmad, S. T. Chakradhar, A. Raghunathan, and T. N. Vijaykumar, "Tarazu: Optimizing mapreduce on heterogeneous clusters," SIGARCH Comput. Archit. News, vol. 40, no. 1, pp. 61-74 Mar. 2012.
[6] Y. Guo, J. Rao, and X. Zhou, "iShuffle: Improving hadoop performance with shuffle-on-write," in Proceedings of the 10th International Conference on Autonomic Computing (ICAC 13), San Jose, CA, 2013, pp. 107-117.
[7] R. Diestel, Graph theory; 2nd ed. Heidelberg: Springer, 2000.
[8] M. Mucha and P. Sankowski, "Maximum matchings via gaussian elim-ination," in Proceedings. 45th Annual IEEE Symposium on Foundations of Computer Science, 2004, pp. 248-255.
[9] B. Bollob́s, Random graphs. Cambridge Studies in Advanced Mathematics 73, 2001.
[10] N. Alon and J. H. Spencer, The probabilistic method. John Wiley \& Sons, 2004.
[11] Z. Bar-Gossef, Y. Birk, T. Jayram, and T. Kol, "Index coding with side information," IEEE Transactions on Information Theory, vol. 57, no. 3, pp. 1479-1494, 2011.
[12] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "Coded MapReduce," in 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2015, pp. 964971.
[13] K. Lee, M. Lam, R. Pedarsani, D. Papailiopoulos, and K. Ramchandran, "Speeding up distributed machine learning using codes," arXiv preprint
[14] F. Farhat, et al. "Stochastic modeling and optimization of stragglers." IEEE Transactions on Cloud Computing (2016).
[15] L. Song and C. Fragouli, "A pliable index coding approach to data shuffling," to appear in ISIT17, arXiv preprint arXiv:1701.05540, 2017.
[16] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," IEEE Transactions on Information Theory, vol. 60, no. 5, pp. 28562867, 2014.
[17] A. Golovnev, O. Regev, and O. Weinstein, "The minrank of random graphs," arXiv preprint arXiv:1607.04842, 2016.

## Appendix A

## Proof of Theorem 2

In the low probability regime $p \leq p_{t h}(m, n, K)$, we saw from Lemma 2 that some functions cannot be calculated directly after the random data placement. Assume that, given the data placement realization, it was possible to have nodes make uncoded broadcast transmissions of the raw data so that $K$ nodes have all the data needed to compute the $K$ functions without further communication. Let $T_{u n}$ be the minimum number of such uncoded broadcast transmissions that we would need to make. $T_{u n}$ can be calculated using a bipartite graph matching problem: we list the $n$ nodes and the $K$ functions on both sides and connect a node to the functions it can compute. We then interpret $T_{u n}$ as the minimum number of uncoded broadcast transmissions nodes need to make such that there is a matching that covers all $K$ functions. Clearly, if after the random data placement, we are only allowed to exchange intermediate function values, $T_{u n}$ would form a lower bound on the number of uncoded transmissions needed.

We are now going to show that, even if we are allowed to make coded broadcast transmssions, we would not achieve significant transmission savings (in order of magnitude). We first note that, we have at least $T_{u n}$ nodes missing at least $T_{u n}$ different data. The number of coded transmissions we make need to satisfy at least these nodes. Using a similar approach as in index coding [11], we can define a matrix $G \in \mathbb{F}_{q}^{T_{u n} \times T_{u n}}$ that "fits" our data shuffling problem as follows:

- the columns are indexed by $T_{u n}$ missing messages - we assume without loss of generality these to be the first $T_{u n}$, $j=1 \ldots T_{\text {un }}$ out of the $m$ messages;
- the rows are indexed by some $T_{u n}$ computing nodes, $i_{1}, i_{2}, \ldots, i_{T_{u n}}$, with node $i_{j}$ missing data $j$ for some function calculation, i.e., $j \notin S_{i_{j}}$;
- the $(i, j)$ entry of $G, g_{i, j}=1$ for $i=j, g_{i, j}=0$ for $j \in[m] \backslash S_{i}$, and $g_{i, j}$ can be arbitrary element in $\mathbb{F}_{q}$ for $j \in S_{i}$; - node $i_{j}$ can use row $i_{j}$ to decode the message $j$ it needs.

Similarly to [11], let $\mathcal{G}$ denote the set of all matrices that fit our problem, the minimum number of coded broadcast transmissions equals the minimum rank of any matrix in $\mathcal{G}$, denoted as $\operatorname{minrank}(\mathcal{G})$.

To derive a lower bound on minrank $(\mathcal{G})$, we use a similar technique as in [17]. Let the rank of $G$ be $k$. For a principal submatrix $M$ of $G$, we denote the weight (number of nonzero elements) of $M$ by $w^{\prime}$, the minimum weight of a column basis of $M$ by $w_{C}^{\prime}$, and the minimum weight of a row basis of $M$ by $w_{R}^{\prime}$. Let the rank of $M$ be $k^{\prime}$. We define a coding structure $M$ to be a $n^{\prime} \times n^{\prime}$ principal submatrix of $G$ with the following properties [17]:

- the weight and the rank of $M$ satisfy $4 k^{\prime} w^{\prime} \geq n^{\prime 2}, k^{\prime} / n^{\prime} \leq$ $k / T_{u n}, w_{C}^{\prime} \leq 2 w^{\prime} k^{\prime} / n^{\prime}$, and $w_{R}^{\prime} \leq 2 w^{\prime} k^{\prime} / n^{\prime}$;

In [17], it is shown that every matrix $G$ contains such a coding structure $M$ and the number of such principal submatrices is at most $\left(n^{\prime} q\right)^{12 w^{\prime} k / T_{u n}}$.

The probability that a matrix $G$ containing $M$ is a fitting matrix is $p^{w^{\prime}-n^{\prime}}$, since we need that side information covers all off-diagonal non-zero elements of $M$. The number of ways we choose an ordered set of $n^{\prime}$ computing nodes for submatrix $M$ is $n(n-1)(n-2) \ldots\left(n-n^{\prime}+1\right) \leq n^{n^{\prime}}$. The number of ways of choosing $M$ among all possible submatrices of $G$ is $\binom{T_{u n}}{n^{\prime}}$. By setting $k=\Theta\left(T_{u n}\right)$, we can calculate the probability

$$
\begin{aligned}
& \operatorname{Pr}\{\operatorname{minrank}(\mathcal{G}) \leq k\} \\
& \leq \sum_{k^{\prime}, n^{\prime}, w^{\prime}} n^{n^{\prime}}\binom{T_{u n}}{n^{\prime}}\left(n^{\prime} q\right)^{12 w^{\prime} k / T_{u n}} p^{w^{\prime}-n^{\prime}} \\
& \leq \sum_{k^{\prime}, n^{\prime}, w^{\prime}} 2^{n^{\prime} \log (n)+n^{\prime} \log \left(T_{u n}\right)+12 w^{\prime} k / T_{u n} \log \left(n^{\prime} q\right)+\left(w^{\prime}-n^{\prime}\right) \log (p)} \\
& \leq n^{4} 2^{\left(2.25 n^{\prime}+12 w^{\prime} k / T_{u n}\right) \log (n q)-w^{\prime} / 4 \log (n)} \\
& \leq n^{4} 2^{14.25 w^{\prime} k / T_{u n} \log (n q)-w^{\prime} / 4 \log (n)}=2^{-\Omega\left(w^{\prime} \log (n)\right)}
\end{aligned}
$$

where the first three inequalities follow from the properties of $M$ and $p \leq \sqrt{\frac{\log (K)}{n}}$, and the last equality follows from $k=\Theta\left(T_{u n}\right)$ and $w^{\prime} \geq C$ for some large enough fixed constant $C$ according to the properties of $M$.

