The effect of range and bandwidth on the round complexity in the congested clique model*

Florent Becker[1], Antonio Fernández Anta[2], Ivan Rapaport[3], and Eric Rémila[4]

[1] LIFO (EA 4022), Université d'Orléans, Orléans, France
[2] IMDEA Networks Institute, Madrid, Spain
[3] DIM-CMM (UMI 2807 CNRS), Universidad de Chile, Santiago, Chile
[4] Univ. Lyon, UJM Saint-Etienne, Saint-Etienne, France

Abstract. The congested clique model is a message-passing model of distributed computation where k players communicate with each other over a complete network. Here we consider synchronous protocols in which communication happens in rounds (we allow them to be randomized with public coins). In the *unicast* communication mode, each player i has her own n-bit input x_i and may send $k-1$ different b-bit messages through each of her $k-1$ communication links in each round. On the other end is the *broadcast* communication mode, where each player can only broadcast a single message over all her links in each round. The goal of this paper is to complete our Brief Announcement at PODC 2015, where we initiated the study of the space that lies between the two extremes. For that purpose, we parametrize the congested clique model by two values: the *range* r, which is the maximum number of different messages a player is allowed to send in each round, and the *bandwidth* b, which is the maximum size of these messages. We show that the space between the unicast and broadcast congested clique models is very rich and interesting. For instance, we show that the round complexity of the pairwise set-disjointness function PWDISJ is completely sensitive to the range r. This translates into a $\Omega(k)$ gap between the unicast ($r = k-1$) and the broadcast ($r = 1$) modes. Moreover, provided that $r \geq 2$ and $rb/\log r = O(k)$, the round complexity of PWDISJ is $\Theta(n/k \log r)$. On the other hand, we also prove that the behavior of PWDISJ is exceptional: almost every boolean function f has maximal round complexity $\Theta(n/b)$. Finally, we prove that $\min\left(\left\lceil \frac{b'}{\lfloor \log r \rfloor} \right\rceil, \left\lceil \frac{r'}{r-1} \right\rceil \left\lceil \frac{b'}{b} \right\rceil\right)$ is an upper bound for the gap between the round complexities with parameters (b, r) and parameters (b', r') of any boolean function.

* Supported in part by the ANR project QuasiCool (ANR-12-JS02-011-01), MINECO grant TEC2014- 55713-R, Regional Government of Madrid (CM) grant Cloud4BigData (S2013/ICE-2894, co- funded by FSE & FEDER), NSF of China grant 61520106005, EC H2020 grants ReCred and NOTRE, CONICYT via Basal in Applied Mathematics, Núcleo Milenio Información y Coordinación en Redes ICM/FIC RC130003, Fondecyt 1130061.

1 Introduction

In this paper we study a synchronous, message-passing model of distributed computation where the underlying communication network is a complete graph. Therefore, the only obstacle to perform any task is due to *congestion.* In fact, the main theoretical purpose of this model, known as *congested clique,* is to serve as a basic model for understanding the role played by congestion in distributed computation [14, 15, 21, 25, 27, 28]. (Besides this, there are interesting connections between the congested clique model and popular systems such as MapReduce [20].)

The model is defined as follows. There are k players. Each player has her own n-bit input x_i and they all collaborate in order to compute a joint boolean function $f(x_1, \ldots, x_k)$. They communicate with each other in synchronous rounds. More precisely, each of the k players may send up to $k-1$ different b-bit messages through each of her $k-1$ communication links. A protocol that computes f stops when every player knows the output. We use the number of rounds as the goodness metric to be minimized. The absolute minimum of this parameter is what we call *round complexity*. In this paper all protocols are allowed to be randomized with public coins. More precisely, the k players have access to a common infinite string of independent random bits. Protocols may return the wrong answer with probability at most ϵ, for some fixed, small $\epsilon > 0$.

Most work on this (unicast) congested clique model considers the joint input as a graph G by giving to each player i the boolean vector $x_i \in \{0,1\}^n$, which is the indicator function of her neighborhood in G. Note that in this case $n = k$ and, therefore, the total number of bits exchanged in each round is bn^2. Unfortunately, due to the huge number of bits transmitted globally per round (even for $b = 1$), no lower bound is known for this model. Drucker, Kuhn and Oshman gave in [15] an explanation for this difficulty. They proved that in this model it is possible to simulate powerful classes of bounded-depth circuits (and therefore lower bounds in the congested clique would yield lower bounds in circuit complexity). The intrinsic power of the (synchronous) congested clique model has allowed some authors [10, 14, 19, 21] to provide extremely fast protocols for some natural graph problems (assuming always that $b = \log n$, following the spirit of the $\mathcal{CONGEST}$ model [29]).

In the broadcast version of the congested clique model, each player can only broadcast a single b-bit message over all her links in each round [15]. This setting is equivalent to the multi-party, number-in-hand computation model, where communication takes place in a shared blackboard [1, 2, 5–7, 15]. In fact, writing a message $\mathcal{M}$ on the blackboard is equivalent to broadcasting $\mathcal{M}$. In this setting, the number of transmitted bits per round decreases from bn^2 to bn. Therefore, obtaining lower bounds using communication complexity reductions becomes possible. For instance, detecting deterministically a triangle in the input graph G requires $\Omega(n/(e^{\mathcal{O}(\sqrt{\log n})}b)$ rounds [15]. On the other hand, fast protocols are also known in the broadcast congested clique model [1, 2, 18, 23].

There is a particular boolean function that we are going to use throughout this paper. This function, that we call *pairwise set-disjointness,* is defined below.

Definition 1. *Let* $k = 2k'$. *Let* $x = (x_1 \dots x_k) \in (\{0,1\}^n)^k$. *Each* x_i *is the indicator vector of a subset* $X_i \subseteq \{1, \dots, n\}$. *Function* pairwise set-disjointness $\textsc{pwdisj}$ *is defined by:* $\textsc{pwdisj}(x) = 1$ *if* $\forall 1 \leq i \leq k', X_i \cap X_{i+k'} = \emptyset$*; and* $\textsc{pwdisj}(x) = 0$ *otherwise.*

Our goal is to complete the work of [4], where we initiated the study of the round complexity of boolean functions according to two parameters of the model:

- The *range* r: the maximum number of different messages a player can send over her links in one round.
- The *bandwidth* b: the maximum size, in bits, of each of these messages.

By analogy with the notation introduced in [15], we denote this model by $\textsc{clique-rcast}_{r\times b}$. Note that the two extreme cases $r = 1$ and $r = k - 1$, which correspond to the broadcast and the unicast communication modes, are the cases already considered in the literature. More precisely,

$$\begin{aligned} \textsc{clique-rcast}_{(k-1)\times b} &= \textsc{clique-}\mathbf{U}\textsc{cast}_b, \\ \textsc{clique-rcast}_{1\times b} &= \textsc{clique-}\mathbf{B}\textsc{cast}_b. \end{aligned}$$

Note also that, if the available bandwidth b is too small, then having a big range r becomes useless, since the number of possible different messages with a bandwidth b is 2^b. More precisely,

$$\forall r \geq 2^b, \textsc{clique-rcast}_{r\times b} = \textsc{clique-rcast}_{2^b\times b} = \textsc{clique-}\mathbf{U}\textsc{cast}_b.$$

Thus, in the sequel, we will assume that $r \leq 2^b$. We denote by $\textsc{round}_{r\times b}(f)$ the round complexity of function f. That is, $\textsc{round}_{r\times b}(f)$ denotes the minimal number of rounds needed by any k-player protocol in $\textsc{clique-rcast}_{r\times b}$ for computing f. We also denote,

$$\begin{aligned} \mathbf{U}\textsc{round}_b(f) &= \textsc{round}_{(k-1)\times b}(f), \\ \mathbf{B}\textsc{round}_b(f) &= \textsc{round}_{1\times b}(f). \end{aligned}$$

A protocol in $\textsc{clique-rcast}_{r\times b}$ is said to be a *broadcasting protocol* if it consists of every player broadcasting its complete input. Obviously, for any function f, there exists a broadcasting protocol which computes f, and we get the universal bound $\textsc{round}_{r\times b}(f) \leq \mathbf{B}\textsc{round}_b(f) \leq \lceil n/b \rceil$. In order to understand the role played by the range r and the bandwidth b in the round complexity of the congested clique model we define the following ratio.

$$\mathit{\Gamma}^{r\times b}_{r'\times b'}(f) = \frac{\textsc{round}_{r\times b}(f)}{\textsc{round}_{r'\times b'}(f)}.$$

The values above obviously depend on k, n and ϵ. But we omit them in order to avoid heavy notation. Finally, by taking the uniform probability over $\{0,1\}^{\{0,1\}^{kn}}$, we also consider what happens with random boolean functions. For instance, we compute probabilities such us $\Pr\{\mathit{\Gamma}^{r\times b}_{r'\times b'}(f) = \alpha\}$, for fixed α.

1.1 Our results

In Section 2 we compare the broadcast model and the unicast model. For that purpose we consider the pairwise set-disjointness function $\textsc{pwdisj}$. We prove that $\mathbf{U}\textsc{round}_b(\textsc{pwdisj}) = \mathcal{O}(n/kb)$ while $\mathbf{B}\textsc{round}_b(\textsc{pwdisj}) = \Omega(n/b)$. In other words, $\Gamma^{1\times b}_{(k-1)\times b}(\textsc{pwdisj}) = \Omega(k)$. This gives a large gap between the unicast and broadcast congested clique models, that grows at least linearly with k.

In Section 3 we prove that the round complexity of $\textsc{pwdisj}$ is completely sensitive to the range r even in the intermediate values between unicast and broadcast. More precisely, we prove that for k sufficiently large and for $r \geq 2$ such that $rb/\log r = \mathcal{O}(k)$ the following holds: $\textsc{round}_{r\times b}(\textsc{pwdisj}) = \Theta(n/k\log r)$. Then, we give some interpretations to this result. In particular, we conclude that $\Gamma^{r\times \log k}_{r'\times \log k}(\textsc{pwdisj}) = \Theta(\log r'/\log r)$ for every $r' \geq r \geq 2$. Note that the logarithmic bandwidth is the most studied case in the congested clique model, and this result yields a hierarchy of models of different computational power according to the range r for this case.

In Section 4 we prove that almost every boolean function f satisfies that $\mathbf{U}\textsc{round}_b(f) = \mathbf{B}\textsc{round}_b(f) = \lceil n/b \rceil$, provided that k is sufficiently large and that $0 \leq \epsilon \leq 0.2$. In other words, $\Gamma^{1\times b}_{(k-1)\times b}(f) = 1$ for almost every f. This means that the gap we found in Section 2 for function $\textsc{pwdisj}$ is exceptional and that the power given by having $r > 1$ is almost always useless. Nevertheless, as pointed out by Drucker et al. [15], finding for $k = n$ an explicit boolean function f with the behavior $\mathbf{U}\textsc{round}_b(f) = \omega(1)$ is (equivalent to solving) a long-standing open problem in circuit complexity theory.

The goal of Section 5 is to compare models with different combinations of range and bandwidth for arbitrary boolean functions f. For doing this we analyze the ratio $\Gamma^{r\times b}_{r'\times b'}(f)$. We make the following observation: for almost every function f we have $\Gamma^{r\times b}_{r'\times b'}(f) = \Theta(b'/b)$. Moreover, if $r \geq r'$ or $r = 2^b$ then $\Gamma^{r\times b}_{r'\times b'}(f) \leq \lceil b'/b \rceil$ for every boolean function f. The general upper bound we obtain is the following $\Gamma^{r\times b}_{r'\times b'}(f) \leq \min\left(\left\lceil \frac{b'}{\lfloor \log r \rfloor} \right\rceil, \left\lceil \frac{r'}{r-1} \right\rceil \left\lceil \frac{b'}{b} \right\rceil\right)$, for $r \geq 2$.

1.2 Related work: The asynchronous case

The congested clique model with bandwidth $b = 1$ –that is, the multiplayer, number-in-hand, message passing model– was introduced by Dolev and Feder [13]. The main difference with our setting is that the original model was *asynchronous*. Hence, protocols, instead of being designed to minimize the number of rounds, were designed to minimize the number of exchanged bits. The first communication complexity lower bounds were obtained by Duris and Rolim [16].

Recently, new techniques and new results have been developed, and tight bounds for the communication complexity of different functions have been obtained. In [30] the authors introduced the symmetrization technique and were able to prove tight $\Omega(nk)$ lower bounds for several direct-sum-like functions such as coordinate-wise AND or coordinate-wise OR. These lower bounds also apply in the blackboard communication mode, where players write messages on a blackboard, visible to everybody. (Note that, in the asynchronous setting, the communication complexity in the blackboard mode gives stronger lower bounds than the communication complexity in the message-passing, point-to-point mode.) This symmetrization technique has been used and developed by other authors as well [26, 31].

It is important to point out that there exists a strict separation between the blackboard communication mode and the message-passing communication mode. For instance, the communication complexity for computing the multiparty set-disjointness function is $\Theta(n \log k + k)$ in the blackboard communication mode [9] and it is $\Theta(nk)$ in the message-passing communication mode [8]. These results on set-disjointness were obtained by using information complexity, a theory introduced in [11]. Information complexity turned out to be an extremely useful theory for proving communication complexity lower bounds [3, 12, 17].

2 A gap in the round complexity of broadcast versus unicast

The first question we would like to answer is the following: How much do we gain if, instead of broadcasting, we have the possibility of sending at least two different messages in each round? This seems to be a simple question. But it is a fundamental one if we want to understand the role played by the range in the congested clique model. For answering this we use the pairwise set-disjointness function PWDISJ defined in Section 1.

Theorem 1. $\mathbf{U}\text{ROUND}_b(\text{PWDISJ}) = \mathcal{O}(n/kb)$.

Proof. We prove that $\mathbf{U}\text{ROUND}_b(\text{PWDISJ}) \leq \left\lceil \frac{\lceil n/k \rceil}{b} \right\rceil + 1$. The protocol is as follows. Let $T = \left\lceil \frac{\lceil n/k \rceil}{b} \right\rceil$. For every $1 \leq t \leq T$, let

$$w_{i,j,t} = (x_i)_{(j-1)\lceil n/k \rceil + (t-1)b+1}, \dots, (x_i)_{(j-1)\lceil n/k \rceil + tb}.$$

Round $1 \leq t \leq T$**.** Each player i sends to each player j (including itself) the b bits of $w_{i,j,t}$.

Round $T+1$**.** Each player j broadcasts 1 if at all rounds t, all its incoming messages from player $1 \leq i \leq k'$ were disjoints with all its incoming messages from player $i + k'$.

Clearly, after T rounds, player j receives $(x_i)_{(j-1)\lceil n/k \rceil+1}, \dots, (x_i)_{j\lceil n/k \rceil}$ from every i. Hence, $\text{PWDISJ}(x) = 0$ if and only if a 0 is broadcasted by some player in the last round. Therefore, every player will know the answer after the last round. □

Theorem 2. $\mathbf{B}\text{ROUND}_b(\text{PWDISJ}) = \Omega(n/b)$.

Proof. It is well-known that, in the two party case $k = 2$, the round complexity of set-disjointness with error probability ϵ is $\Omega(n/b)$ [22]. If $k > 2$ we get the same bound for (PWDISJ by considering the instance where $x_1 = x \in \{0,1\}^n$ is given to player 1, $x_{1+k'} = y \in \{0,1\}^n$ is given to player 2, and the empty set ϕ, represented by $(0,\ldots,0)^T$, is given to all the other $k-2$ players. □

Corollary 1. *Let $k = n$. Then,* $\mathbf{U}\text{ROUND}_1(\text{PWDISJ}) = 2$ *and* $\mathbf{B}\text{ROUND}_b(\text{PWDISJ}) = \Omega(n/b)$*.*

Corollary 2. $\Gamma^{1\times b}_{(k-1)\times b}(\text{PWDISJ}) = \Omega(k)$*.*

3 A hierarchy of models according to the range

In previous section we proved that the broadcast ($r = 1$) and the unicast ($r = k - 1$) models are fundamentally different in their power to solve one particular problem. These two models are the two ends of the spectrum of values of the range r. In this section we prove that the sensitivity to the range is more general. In particular, we show that the round complexity of PWDISJ is completely sensitive to the range.

Lemma 1. $\text{ROUND}_{r\times b}(\text{PWDISJ}) = \Omega\left(\frac{n}{\min(kb, rb+\lceil\log r\rceil k)}\right)$.

Proof. We use a reduction from the two-party communication problem $\text{DISJ}_{k'n}$, where instances are pairs (x, y) of boolean vectors, each of length $k'n$. The communication complexity (bits to be exchanged) of $\text{DISJ}_{k'n}$ is $\Theta(k'n)$ [22]. We transform an instance of $\text{DISJ}_{k'n}$ into an instance of PWDISJ in the direct way. From (x, y) we define the input $(x_1,\ldots,x_k)$ of function PWDISJ as follows: $x = x_1 \cdots x_{k'}$ and $y = x_{k'+1} \cdots x_k$. Obviously, $\text{DISJ}_{k'n}(x, y) = 1 \iff \text{PWDISJ}(x_1,\ldots,x_k) = 1$.

Let us consider any protocol P that solves PWDISJ in $T_{\mathtt{P}}$ rounds. If we group players 1 to k' into a global player A and players $k' + 1$ to k into a global player B, protocol P would yield a protocol for solving $\text{DISJ}_{k'n}$. So the question is the following: How many bits are exchanged between A and B? Let us derive an upper bound for this.

Consider a player i in A. Player i sends one message of length b to each player in B, thus he sends $k'b$ bits. However, since $r \leq 2^b$, the messages sent by player i to players in B can be compressed as follows (see Figure 1). Since player i can send up to r different messages, one can consider that she sends to each player $j \in B$ a message numbered from the set $\{0, 1,\ldots, r-1\}$ that identifies the message $m(i, j)$ sent to player j. These numbers, of $\lceil\log r\rceil$ bits each, can be used to obtain the actual message from a table that contains the r messages, of b bits each, sent by i. Hence, the total number of bits sent by i to B is upper bounded by the length of the k' numbers, $\lceil\log r\rceil k'$ bits, and the size of the message table, br bits; a total of $rb + \lceil\log r\rceil k'$ bits.

Let us define $\beta = \min(bk', rb + \lceil\log r\rceil k')$. In each round, the number of bits exchanged between A and B is upper bounded by $k\beta$. Therefore, considering

Destination	message number
1	00
2	01
3	00
4	11
5	01

Message number	contents
00	11101010
01	01010100
10	00101100

Fig. 1. Summarizing the information sent by a player i: sending the 34 bits represented in black is enough, as opposed to 40 bits for the concatenation of the messages of each of the 5 destination players.

that the communication complexity of $\text{DISJ}_{k'n}$ is $\Theta(k'n)$, it follows that $T_{\mathsf{P}} k\beta = \Omega(k'n)$. Therefore, $\text{ROUND}_{r\times b}(\text{PWDISJ}) = \Omega(\frac{n}{2\beta})$, as claimed. □

Lemma 2. $\text{ROUND}_{r\times b}(\text{PWDISJ}) \le \left\lceil \frac{n}{k\lfloor \log r\rfloor} \right\rceil + 1$.

Proof. Consider the same protocol used in the proof of Theorem 1 but with messages of $\lfloor \log r\rfloor \le b$ bits. □

Putting these together, we get the following theorem.

Theorem 3. *For k sufficiently large and for $r \ge 2$ such that $\frac{rb}{\log r} = O(k)$,*

$$\text{ROUND}_{r\times b}(\text{PWDISJ}) = \Theta\left(\frac{n}{k\log r}\right).$$

Proof. The upper bound follows from the previous lemma. For the lower bound, it follows from Lemma 1 that

$$\text{ROUND}_{r\times b}(\text{PWDISJ}) \ge \frac{n}{k\min(b, \lceil \log r\rceil(1+\frac{2rb}{\lceil \log r\rceil k}))} \ge \frac{n}{k\lceil \log r\rceil(1+\frac{2rb}{\lceil \log r\rceil k})},$$

where the last inequality follows from $\lceil \log r\rceil \le b$ and $\frac{2rb}{\lceil \log r\rceil k} > 0$. Since $\frac{rb}{\log r} = O(k)$, we deduce that, for k sufficiently large, there is a constant $\Delta > 0$ such that $\frac{2rb}{\lceil \log r\rceil k} \le \Delta$, and hence

$$\text{ROUND}_{r\times b}(\text{PWDISJ}) \ge \frac{n}{k\lceil \log r\rceil(1+\Delta)} = \Omega\left(\frac{n}{k\lceil \log r\rceil}\right).$$

□

The natural way to interpret Theorem 3 is to parametrize everything by k. Following the spirit of the $\mathcal{CONGEST}$ model [29], we are going to restrict both the bandwidth and the range by taking $b = \log k$ and varying r from 2 to $k-1$. Observe that, when $b = \log k$ and $r \le k-1$, it always holds that $\frac{rb}{\log r} = O(k)$. Hence, the next corollaries are direct consequences of Theorem 3.

Corollary 3. *For every n and every constant integer $c \geq 2$, we have*

$$\text{ROUND}_{\log k \times \log k}(\text{PWDISJ}) = \Theta\left(\frac{n}{k \log\log k}\right) \quad \textit{and} \quad \text{ROUND}_{c \times \log k}(\text{PWDISJ}) = \Theta\left(\frac{n}{k}\right)$$

$$\textit{In other words,} \quad \Gamma^{c\times\log k}_{\log k \times \log k}(\text{PWDISJ}) = \Theta\left(\log\log k\right).$$

In general, we can state the following corollary.

Corollary 4. *For every n and every $r' \geq r \geq 2$, we have*

$$\text{ROUND}_{r' \times \log k}(\text{PWDISJ}) = \Theta\left(\frac{n}{k \log r'}\right) \quad \textit{and} \quad \text{ROUND}_{r \times \log k}(\text{PWDISJ}) = \Theta\left(\frac{n}{k \log r}\right)$$

$$\textit{In other words,} \quad \Gamma^{r\times\log k}_{r' \times \log k}(\text{PWDISJ}) = \Theta\left(\frac{\log r'}{\log r}\right).$$

4 Most functions have maximal round complexity

From the results presented in the previous sections one may be tempted to conclude that, in general, increasing the range r increases the power of the protocols. In particular, one may conclude that the unicast congested clique model has much more power than the broadcast congested clique model (even if in the first we restrict the bandwidth to $b = 1$ while in the latter we allow it to be $b = o(n)$). We show here that this fact, which holds for function PWDISJ, holds for very few other functions. More precisely, we are going to prove that for almost every boolean function f, the broadcasting protocol is optimal. We start by considering deterministic decision protocols that compute functions f correctly (i.e., they make no mistake). (The proofs omitted can be found in the Appendix.)

Lemma 3. *The number of T-round deterministic decision protocols in the unicast congested clique model* CLIQUE-**U**CAST$_b$ *is at most $2^{N(T)}$, where*

$$N(T) = 2^{T(k-1)b+n}\left(1 + \frac{(k+1)(k-1)b}{2^{(k-1)b}}\right).$$

Now, we still consider deterministic protocols, but now we allow them to make mistakes. We say that a deterministic protocol P computes f with error $\epsilon \geq 0$ if it outputs $f(x)$ for at least $(1-\epsilon)2^{nk}$ of the inputs x of f.

Lemma 4. *Let P be a deterministic decision protocol and let $P(x)$ denote the output of P with input $x \in \{0,1\}^{nk}$. Let $M_\epsilon(P)$ be the number of functions f which are computed by P with an error $\epsilon > 0$. We have,*

$$M_\epsilon(P) \leq \left(\frac{2e}{\epsilon}\right)^{\epsilon 2^{nk}} = 2^{\log(\frac{2e}{\epsilon})\epsilon 2^{nk}}.$$

We show now that a deterministic protocol P that computes a function f chosen uniformly at random with error ϵ requires the maximal number of rounds $\lceil n/b \rceil$ with high probability. Let us extend our notation, so that $\textbf{U}\text{ROUND}_b^\epsilon(f)$ is the round complexity of function f when protocols are deterministic and error ϵ is allowed.

Theorem 4. *For k sufficiently large and for every n, and $\epsilon > 0$ such that $1 - \log(\frac{2e}{\epsilon})\epsilon > 0$, we have*

$$\Pr\{\textbf{U}\text{ROUND}_b^\epsilon(f) = \lceil n/b \rceil\} \geq 1 - 2^{-2^{kn}(\frac{1-\log(\frac{2e}{\epsilon})\epsilon}{2})}.$$

For $\epsilon = 0$ (i. e. the case without error), we have

$$\Pr\{\textbf{U}\text{ROUND}_b^0(f) = \lceil n/b \rceil\} \geq 1 - 2^{-2^{kn}\,0.5}.$$

Proof. Since there are $2^{2^{kn}}$ different functions $f : \{0,1\}^{kn} \to \{0,1\}$, we have

$$\Pr\{\textbf{U}\text{ROUND}_b^\epsilon(f) \leq T\} \leq \frac{2^{N(T)} \max_P M_\epsilon(P)}{2^{2^{kn}}}.$$

From Lemmas 3 and 4, for $\epsilon > 0$, we have

$$\begin{aligned}\Pr\{\textbf{U}\text{ROUND}_b^\epsilon(f) \leq T\} &\leq 2^{2^{T(k-1)b+n}(1+\frac{(k+1)(k-1)b}{2^{(k-1)b}})} 2^{\log(\frac{2e}{\epsilon})\epsilon 2^{nk}} 2^{-2^{kn}} \\ &\leq 2^{-2^{kn}(1-\log(\frac{2e}{\epsilon})\epsilon - 2^{T(k-1)b+n-kn}(1+\frac{(k+1)(k-1)b}{2^{(k-1)b}}))}.\end{aligned}$$

For k sufficiently large, the quantity $1 + \frac{(k+1)(k-1)b}{2^{(k-1)b}}$ can be upper bounded (by 2 for example). Now let us take $T = \lceil n/b \rceil - 1$. Then, we have $Tb - n \leq -b$ and, thus

$$2^{T(k-1)b+n-kn} = 2^{(k-1)(Tb-n)} \leq 2^{-b(k-1)} \leq 2^{-k}$$

Thus, for k sufficiently large, the term, $2^{T(k-1)b+n-kn}(1 + \frac{(k+1)(k-1)b}{2^{(k-1)b}})$ can be upper bounded by any positive value, in particular by $\frac{1-\log(\frac{2e}{\epsilon})\epsilon}{2}$. Thus, we get that

$$\Pr\{\textbf{U}\text{ROUND}_b^\epsilon(f) \leq \lceil n/b \rceil - 1\} \leq 2^{-2^{kn}(\frac{1-\log(\frac{2e}{\epsilon})\epsilon}{2})}.$$

which is the result since, for any f, one trivially has $\textbf{U}\text{ROUND}_b^\epsilon(f) \leq \textbf{U}\text{ROUND}_b^0(f) \leq \lceil n/b \rceil$.

For $\epsilon = 0$ we proceed on the same way, after noticing that $\max_P M_0(P) = 1$. □

Theorem 5. *For k sufficiently large, for every n, and $0 \leq \epsilon \leq 0.2$, there exists a positive constant $c(\epsilon) > 0$ such that $\Pr\{\textbf{U}\text{ROUND}_b^\epsilon(f) = \lceil n/b \rceil\} \geq 1 - 2^{-2^{kn}c(\epsilon)}$.*

Recall that $\textbf{U}\text{ROUND}_b(f)$ is the round complexity of computing function f with randomized protocols, which may use public coins, with success probability $1 - \epsilon$. From the previous results we can prove that most functions have round complexity $\lceil n/b \rceil$.

Corollary 5. *For k sufficiently large and for every n, and $0 \leq \epsilon \leq 0.2$, there exists a positive constant $c(\epsilon) > 0$ such that,*

$$\Pr\{\mathbf{U}\text{ROUND}_b(f) = \lceil n/b \rceil\} \geq 1 - 2^{-2^{kn} c(\epsilon)}.$$

Proof. The result follows from Theorem 5, and Theorem 3.20 at [24] using the uniform distribution as the distribution μ of the inputs. □

The following bound is obvious for any function f.

$$\Gamma^{1\times b}_{(k-1)\times b}(f) = \frac{\mathbf{B}\text{ROUND}_b(f)}{\mathbf{U}\text{ROUND}_b(f)} \geq \min_f \Gamma^{1\times b}_{(k-1)\times b}(f) \geq 1.$$

Next corollary, which is a direct consequence of Corollary 5, says that previous inequality is in fact an equality for almost every boolean function.

Corollary 6. *For k sufficiently large and for every n, and $0 \leq \epsilon \leq 0.2$, there exists a positive constant $c(\epsilon) > 0$ such that,*

$$\Pr\{\Gamma^{1\times b}_{(k-1)\times b}(f) = 1\} \geq 1 - 2^{-2^{kn} c(\epsilon)}.$$

5 Comparing models with different combinations of range and bandwidth for arbitrary boolean functions

In this section we explore the relative round complexities of different modes of the congested clique model with various combinations of range and bandwidth $\Gamma^{r\times b}_{r'\times b'}(f)$ for arbitrary boolean functions f. The first result shows that for *most* boolean functions f, $\Gamma^{r\times b}_{r'\times b'}(f) = \Theta(b'/b)$.

Theorem 6. *For k sufficiently large and for every n, there is a positive constant $c(\epsilon) > 0$ such that*

$$\Pr\left\{\Gamma^{r\times b}_{r'\times b'}(f) = \lceil n/b \rceil \,/\, \lceil n/b' \rceil\right\} \geq 1 - 2^{-2^{kn} c(\epsilon)+1}.$$

Proof. From Corollary 5, a function f *simultaneously* satisfies $\text{ROUND}_{r\times b}(f) = \lceil n/b \rceil$ and $\text{ROUND}_{r'\times b'}(f) = \lceil n/b' \rceil$ with probability at least $1 - 2^{-2^{kn} c(\epsilon)+1}$. □

Now, we show that in fact the typical case shown in the previous theorem is not far from the worst case, studied in the following sequence of results.

Theorem 7. *Let r be such that $r \geq r'$ or $r = 2^b$. Then, for every function f, $\Gamma^{r\times b}_{r'\times b'}(f) \leq \lceil b'/b \rceil$.*

Proof. Let `P'` be a T-round protocol in $\text{CLIQUE-RCAST}_{r'\times b'}$. From `P'` we construct the protocol `P` in $\text{CLIQUE-RCAST}_{r\times b}$ as follows. Consider the message $m_t(i,j)$ sent by player i to player j in round t of `P'`. For each $1 \leq \ell \leq \lceil b'/b \rceil$, let $\text{block}^{\ell}_t(i,j)$ be the ℓ^{th} block of length b of $m_t(i,j)$. The last block is padded with 0s. For each ℓ and i, we have: $|\{\text{block}^{\ell}_t(i,j), 1 \leq j \leq k \in N\}| \leq \min\{r', 2^b\} \leq r$.

Then, during round number $(t-1)\lceil b'/b \rceil + \ell$ of P, player i sends to player j the b bits of $\text{block}_t^i(u,v)$. The inequalities above ensure that P is a well-defined protocol in CLIQUE-RCAST$_{r\times b}$. Since P knows the bandwidth b' it can discard the padding bits. The total number of rounds executed by P is $T\lceil b'/b \rceil$. □

Theorem 8. *Let $b \le b' \le n$, and k sufficiently large. Then, there exists a function f such that: $\Gamma^{r\times b}_{r'\times b'}(f) = \lceil b'/b \rceil$.*

Proof. Let $b' = n$. In this case, every function $f : (\{0,1\}^n)^k \to \{0,1\}$ can be solved in one round in the model CLIQUE-RCAST$_{r'\times b'}$. On the other hand, from Corollary 5, almost every function $f : (\{0,1\}^n)^k \to \{0,1\}$ satisfies $\text{ROUND}_{r\times b}(f) = \lceil n/b \rceil = \lceil b'/b \rceil$. When $n > b'$, let us define $n' = b'$. From Corollary 5, almost every function $f' : (\{0,1\}^{n'})^k \to \{0,1\}$ satisfies $\text{ROUND}_{r\times b}(f') = \left\lceil \frac{n'}{b} \right\rceil$. Let us take one such function f', and define a new function $f : (\{0,1\}^n)^k \to \{0,1\}$ as follows: $f(x_1, x_2 .... x_k) = f'(y_1, y_2, ..., y_k)$, where each y_i is the vector formed with the n' first bits of x_i. Hence, $\text{ROUND}_{r\times b}(f) = \text{ROUND}_{r\times b}(f') = \lceil n'/b \rceil = \lceil b'/b \rceil$ while $\text{ROUND}_{r'\times b'}(f) = 1$. □

Remark 1. When $b|b'$ is a multiple of b and $b'|n$, we have $\lceil n/b \rceil / \lceil n/b' \rceil = (n/b)/(n/b') = b'/b = \lceil b'/b \rceil$. When $n = b'$, we also have $\lceil n/b \rceil / \lceil n/b' \rceil = \lceil b'/b \rceil$. Thus, in the previous cases, for $r \ge r'$ of $r = 2^b$, the maximal value $\lceil b'/b \rceil$ for the value of $\Gamma^{r\times b}_{r'\times b'}(f)$ is reached with high probability. On the other hand, in some cases, there exists a small but intriguing gap between the maximal value $\lceil b'/b \rceil$ and the value $\lceil n/b \rceil / \lceil n/b' \rceil$ reached with high probability. For example, take $b = 2$, $b' = 3$. For $n = 4$, we have $\lceil b'/b \rceil = 2$ and $\lceil n/b \rceil / \lceil n/b' \rceil = 1$.

Note that Theorem 7 holds when $r' \le r$ or $r = 2^b$. Without this hypothesis we only get the following weaker, general bound.

Theorem 9. *Let $r \ge 2$ and $r' \ge 1$. Then, for every function f,*

$$\Gamma^{r\times b}_{r'\times b'}(f) \le \min\left(\left\lceil \frac{b'}{\lfloor \log r \rfloor} \right\rceil, \left\lceil \frac{r'}{r-1} \right\rceil \left\lceil \frac{b'}{b} \right\rceil \right).$$

Observe that the two values of the minimum are complementary, since none implies the other.

References

1. Ahn, K.J., Guha, S., McGregor, A.: Analyzing graph structure via linear measurements. In: Proc. of SODA 2012. pp. 459–467 (2012)
2. Ahn, K.J., Guha, S., McGregor, A.: Graph sketches: Sparsification, spanners, and subgraphs. In: Proc. of PODS 2012. pp. 5–14 (2012)
3. Bar-Yossef, Z., Jayram, T.S., Kumar, R., Sivakumar, D.: An information statistics approach to data stream and communication complexity. In: Proc. of FOCS 2002. pp. 209–218 (2002)

4. Becker, F., Fernández Anta, A., Rapaport, I., Rémila, E.: Brief Announcement: A hierarchy of congested clique models, from broadcast to unicast. In: Proc. of PODC 2015. pp. 167–169 (2015)
5. Becker, F., Kosowski, A., Nisse, N., Rapaport, I., Suchan, K.: Allowing each node to communicate only once in a distributed system: shared whiteboard models. In: Proc. of SPAA 2012. pp. 11–17 (2012)
6. Becker, F., Matamala, M., Nisse, N., Rapaport, I., Suchan, K., Todinca, I.: Adding a referee to an interconnection network: What can (not) be computed in one round. In: Proc. of IPDPS 2011. pp. 508–514 (2011)
7. Becker, F., Rapaport, I., Montealegre, P., Todinca, I.: The simultaneous number-in-hand communication model for networks: private coins, public coins and determinism. In: Proc. of SIROCCO 2014. pp. 83–95 (2014)
8. Braverman, M., Ellen, F., Oshman, R., Pitassi, T., Vaikuntanathan, V.: A tight bound for set disjointness in the message-passing model. In: Proc. of FOCS 2013. pp. 668–677 (2013)
9. Braverman, M., Oshman, R.: On information complexity in the broadcast model. In: Proc. of PODC 2015. pp. 355–364 (2015)
10. Censor-Hillel, K., Kaski, P., Korhonen, J.H., Lenzen, C., Paz, A., Suomela, J.: Algebraic methods in the congested clique. In: Proc. of PODC 2015. pp. 143–152
11. Chakrabart, A., Shi, Y., Wirth, A., Yao, A.: Informational complexity and the direct sum problem for simultaneous message complexity. In: Proc. of FOCS 2001. pp. 270–278. IEEE (2001)
12. Chattopadhyay, A., Mukhopadhyay, S.: Tribes is hard in the message passing model. In: Proc. of STACS 2009. pp. 224–237 (2015)
13. Dolev, D., Feder, T.: Multiparty communication complexity. In: Proc. of FOCS 1989. pp. 428–433 (1989)
14. Dolev, D., Lenzen, C., Peled, S.: Tri, Tri again: Finding triangles and small subgraphs in a distributed setting. In: Proc. of DISC 2012. pp. 195–209 (2012)
15. Drucker, A., Kuhn, F., Oshman, R.: On the power of the congested clique model. In: Proc. of PODC 2014. pp. 367–376 (2014)
16. Ďuriš, P., Rolim, J.D.: Lower bounds on the multiparty communication complexity. Journal of Computer and System Sciences 56(1), 90–95 (1998)
17. Gronemeier, A.: Asymptotically optimal lower bounds on the NIH-multi-party information complexity of the AND-function and disjointness. In: Proc. of STACS 2009. pp. 505–516 (2009)
18. Guha, S., McGregor, A., Tench, D.: Vertex and hyperedge connectivity in dynamic graph streams. In: Proc. of PODS 2015. pp. 241–247 (2015)
19. Hegeman, J.W., Pandurangan, G., Pemmaraju, S.V., Sardeshmukh, V.B., Scquizzato, M.: Toward optimal bounds in the congested clique: Graph connectivity and mst. In: Proc. of PODC 2015. pp. 91–100 (2015)
20. Hegeman, J.W., Pemmaraju, S.V.: Lessons from the congested clique applied to MapReduce. In: Proc. of SIROCCO 2014. pp. 149–164 (2014)
21. Hegeman, J.W., Pemmaraju, S.V., Sardeshmukh, V.B.: Near-constant-time distributed algorithms on a congested clique. In: Proc. of DISC 2014. pp. 514–530
22. Kalyanasundaram, B., Schintger, G.: The probabilistic communication complexity of set intersection. SIAM Journal on Discrete Mathematics 5(4), 545–557 (1992)
23. Kari, J., Matamala, M., Rapaport, I., Salo, V.: Solving the induced subgraph problem in the randomized multiparty simultaneous messages model. In: Proc. of SIROCCO 2015. pp. 370–384 (2015)
24. Kushilevitz, E., Nisan, N.: Communication Complexity. Cambridge U. Press (2006)

25. Lenzen, C.: Optimal deterministic routing and sorting on the congested clique. In: Proc. of PODC 2013. pp. 42–50 (2013)
26. Li, Y., Sun, X., Wang, C., Woodruff, D.P.: On the communication complexity of linear algebraic problems in the message passing model. In: Proc. of DISC 2014. pp. 499–513 (2014)
27. Lotker, Z., Pavlov, E.: MST construction in $O(\log \log n)$ communication rounds. In: Proc. of SPAA 2003. pp. 94–100 (2003)
28. Patt-Shamir, B., Teplitsky, M.: The round complexity of distributed sorting. In: Proc. of PODC 2011. pp. 249–256 (2011)
29. Peleg, D.: Distributed computing: a locality-sensitive approach. Society for Industrial and Applied Mathematics (2000)
30. Phillips, J.M., Verbin, E., Zhang, Q.: Lower bounds for number-in-hand multiparty communication complexity, made easy. In: Proc. of SODA 2012. pp. 486–501
31. Woodruff, D.P., Zhang, Q.: An optimal lower bound for distinct elements in the message passing model. In: Proc. of SODA 2014. pp. 718–733 (2014)

A Proof of Lemma 3

Lemma 3 *The number of T-round deterministic decision protocols in the unicast congested clique model* CLIQUE-UCAST$_b$ *is at most $2^{N(T)}$, where*

$$N(T) = 2^{T(k-1)b+n}\left(1 + \frac{(k+1)(k-1)b}{2^{(k-1)b}}\right).$$

Proof. Let $1 \le t \le T-1$. At round t, player i can send at most $2^{(k-1)b}$ different messages. The choice of player i depends on its input together with the sequence of messages received during the previous $t-1$ rounds; this makes $2^{(t-1)(k-1)b+n}$ possibilities. Hence, the behavior of player i at round t is a function from a set of cardinality $2^{(t-1)(k-1)b+n}$ to a set of cardinality $2^{(k-1)b}$. Considering the k players acting independently, the number of possible behaviors of the protocol at round t is at most $2^{k(k-1)b\,2^{(t-1)(k-1)b+n}}$. Now we can count all the possible behaviors during the $T-1$ first rounds. This gives $2^{\sum_{t=1}^{T-1} k(k-1)b\,2^{(t-1)(k-1)b+n}}$ and we can bound the exponent from above as follows.

$$\begin{aligned}\sum_{t=1}^{T-1} k(k-1)b\,2^{(t-1)(k-1)b+n} &= k(k-1)b\,2^n \frac{2^{(k-1)b(T-1)}-1}{2^{(k-1)b}-1} \\ &\le k(k-1)b\,2^{(k-1)b(T-1)+n}.\end{aligned}$$

For the last round, one only needs to know the message received by some fixed player, for instance player 1. In fact, the output of the protocol depends on what is received in round T by player 1 (the output computed by the other players are all the same). Therefore, the behavior of player i in the last round T is a function $\{0,1\}^{(T-1)(k-1)b+n} \to \{0,1\}^b$ Hence, globally, the number of possible behaviors in the last round is at most $2^{(k-1)b\,2^{(T-1)(k-1)b+n}}$.

The final decision of player 1 is a function $\{0,1\}^{T(k-1)b+n} \to \{0,1\}$. The number of possible decisions is therefore $2^{2^{T(k-1)b+n}}$. Putting it all together, it follows that the number of T-round protocols is at most

$$\begin{aligned}2^{k(k-1)b2^{(k-1)b(T-1)+n}}\,2^{(k-1)b2^{(T-1)(k-1)b+n}}\,2^{2^{T(k-1)b+n}} &= 2^{2^{T(k-1)b+n}(1+\frac{k(k-1)b+(k-1)b}{2^{(k-1)b}})} \\ &= 2^{2^{T(k-1)b+n}(1+\frac{(k+1)(k-1)b}{2^{(k-1)b}})}.\end{aligned}$$

□

B Proof of Lemma 4

Lemma 4 *Let P be a deterministic decision protocol and let $P(x)$ denote the output of P with input $x \in \{0,1\}^{nk}$. Let $M_\epsilon(P)$ be the number of functions f which are computed by P with an error $\epsilon > 0$. We have,*

$$M_\epsilon(P) \le \left(\frac{2e}{\epsilon}\right)^{\epsilon 2^{nk}} = 2^{\log(\frac{2e}{\epsilon})\epsilon 2^{nk}}.$$

Proof. Let $s = \lfloor 2^{nk}\epsilon \rfloor$. Consider a set S of possible inputs of P, (i.e. $S \subset \{0,1\}^{nk}$) of size s. Given such a set, let F_S be the set

$$F_S = \{f | f(x) = P(x) \text{ for } x \notin S\}$$

If $f \in F_S$ then f can be computed by P with error ϵ. Conversely, if f can be computed by P with error ϵ, then there exists a set S such that $f \in F_S$. Thus

$$M_\epsilon(P) \leq 2^s \binom{2^{nk}}{s}.$$

We use the classical bound $\binom{p}{q} \leq (\frac{pe}{q})^q$ to get

$$M_\epsilon(P) \leq 2^s \left(\frac{2^{nk}e}{s}\right)^s \leq 2^s \left(\frac{e}{\epsilon}\right)^{\epsilon 2^{nk}} \leq \left(\frac{2e}{\epsilon}\right)^{\epsilon 2^{nk}}.$$

□

C Proof of Theorem 5

Theorem 5 *For k sufficiently large, for every n, and $0 \leq \epsilon \leq 0.2$, there exists a positive constant $c(\epsilon) > 0$ such that* $\Pr\{\mathbf{U}\text{ROUND}_b^\epsilon(f) = \lceil n/b \rceil\} \geq 1 - 2^{-2^{kn}c(\epsilon)}$.

Proof. For $\epsilon > 0$, let $g(\epsilon) = \log(\frac{2e}{\epsilon})\epsilon = \epsilon\log(2e) - \epsilon\log(\epsilon)$. We have $g'(\epsilon) = \log(2e) - \log(\epsilon) - 1$ thus g' is decreasing and $g'(1) = \log(2e) - 1 > 0$, from which we deduce that g' is positive on $]0,1]$ and, therefore, g is increasing in $]0;1]$.

Moreover, $g(0.2) \simeq 0.95292462715 < 1$. Thus, for $0 < \epsilon \leq 0.2$, $g(\epsilon) < 1$. Thus, for $0 < \epsilon \leq 0.2$, Theorem 4 applies and we get the result. □

D Proof of Theorem 9

Theorem 9 *Let $r \geq 2$ and $r' \geq 1$. Then, for every function f,*

$$\Gamma^{r\times b}_{r'\times b'}(f) \leq \min\left(\left\lceil \frac{b'}{\lfloor \log r \rfloor}\right\rceil, \left\lceil \frac{r'}{r-1}\right\rceil \left\lceil \frac{b'}{b}\right\rceil\right).$$

Proof. First, we have that $\text{ROUND}_{r\times b}(f) \leq \mathbf{U}\text{ROUND}_{\lfloor \log r \rfloor}(f)$. Hence, from Theorem 7 and since $r \leq 2^b$, we have that $\Gamma^{r\times b}_{r'\times b'}(f) \leq \Gamma^{2^{\lfloor \log r \rfloor}\times \lfloor \log r \rfloor}_{r'\times b'}(f) \leq \left\lceil \frac{b'}{\lfloor \log r \rfloor}\right\rceil$.

Second, we can assume that $r \leq r'$ (because, otherwise, we have $\left\lceil \frac{r'}{r-1}\right\rceil = 1$ and we can apply Theorem 7). We will make a refinement of the simulation used in the proof of Theorem 7. Now, in protocol `P`, player i sends the messages $\text{block}_t^\ell(i,j)$ in $\lceil \frac{r'}{r-1} \rceil > 1$ rounds. To do so, player i first sorts lexicographically the (up to) r' different messages in the set $\{\text{block}_t^\ell(i,j), 1 \leq j \leq k\}$ (recall that all these messages have length b.)

Let L be the list of sorted messages. In a first round, player i extracts the first (up to) $r-1$ messages from L, excluding the message 1^b if it is among them. Then, she sends the extracted messages to their recipients, and 1^b to the other players. Since 1^b cannot be among the extracted messages, players that receive it know it is a flag meaning "no message".

For rounds 2 to $\lceil \frac{r'}{r-1} \rceil$, player i extracts the first (up to) $r-1$ messages from L not sent yet, sends them to their recipients, and 0^b to the other players. In these rounds, it is the message 0^b the flag which means "no message" (observe that if 0^k was initially in L it had to be sent in the first round).

After $\lceil \frac{r'}{r-1} \rceil$ rounds, all the messages in L, and hence in $\text{block}_t^\ell(i,\cdot)$, have been sent. Repeating this process for every block ℓ, $1 \leq \ell \leq \lceil \frac{b'}{b} \rceil$, after $\left\lceil \frac{r'}{r-1} \right\rceil \left\lceil \frac{b'}{b} \right\rceil$ rounds protocol `P` has completed the simulation of one round of protocol `P'`, as claimed. □