Brief Announcement: A Hierarchy of Congested Clique Models, from Broadcast to Unicast*

Florent Becker
Université d'Orléans
Orléans, France
florent.becker@univ-orleans.fr

Antonio Fernández Anta
IMDEA Networks Institute
Madrid, Spain
antonio.fernandez@imdea.org

Ivan Rapaport
DIM-CMM (UMI 2807 CNRS)
Universidad de Chile
Santiago, Chile
rapaport@dim.uchile.cl

Eric Rémila
Université de Lyon - GATE
LSE (UMR 5824 - CNRS)
St-Etienne, France
eric.remila@univ-st-etienne.fr

ABSTRACT

The $\mathcal{CONGEST}$ model is a synchronous, message-passing model of distributed computation in which each node can send (possibly different) messages of $\mathcal{O}(\log n)$ bits along each of its incident communication links in each round, where n is the number of computing nodes in the system. In the particular case where the communication network is a complete graph, we have the unicast *congested clique model.* On the other end is the broadcast version of the congested clique model, in which each node can only broadcast a single message over all its links in each round. In this paper we explore the space, in terms of round complexity, that lies between these two congested clique models. Hence, we parametrize the congested clique model with the *range* r, the maximum number of different messages a node can send over its incident links in one round. Additionally, we study the effect of the *bandwidth* b, the maximum size in bits 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 a problem (especially designed for this work) takes $\Omega(n/\log n)$ rounds in the broadcast model ($r = 1$), while it can be solved in two rounds if two messages can be sent ($r = 2$). Other gaps are found in other parts of the spectrum of values of r. We do this by providing techniques to simulate protocols with different parameters. Therefore, we conclude that, with respect to their power to solve certain problems, there is a strict hierarchy of congested clique models.

Categories and Subject Descriptors

D.1.3 [**Software**]: Programming Techniques—*Concurrent programming*; F.1.1 [**Theory of Computation**]: Computation by Abstract Devices—*Models of computation*

General Terms

Theory, Algorithms

Keywords

$\mathcal{CONGEST}$ model, congested clique model, message passing, distributed computation

1. INTRODUCTION

The $\mathcal{CONGEST}$ model is a synchronous, message-passing model of distributed computation in which each node can send (possibly different) messages of $\mathcal{O}(\log n)$ bits along each of its incident communication links in each round [12], where n is the number of computing nodes in the system. In the particular case where the communication network is a complete graph, all the information distributed in the nodes becomes local and 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 [6, 7, 8, 9, 10, 11], is to serve as a basic model for understanding the role played by congestion in distributed computation.

In the much more restricted broadcast version of the congested clique model, each node can only broadcast a single $\mathcal{O}(\log n)$-bit message over all its links in each round [7]. This setting is equivalent to the multi-party, number-in-hand computation model, where communication takes place in a shared whiteboard [1, 2, 3, 4, 5, 7]: writing a message $\mathcal{M}$ on the whiteboard is equivalent to broadcasting $\mathcal{M}$.

In this paper we explore the space, in terms of round complexity, that lies between the model in which a different message can be sent over each link, and the model in which all the links of the node carry the same message. Hence, we parametrize the congested clique model with the *range* r,

*Partially supported by CONICYT via Basal in Applied Mathematics (I.R.), Núcleo Milenio Información y Coordinación en Redes ICM/FIC RC130003 (I.R.), Fondecyt 1130061 (I.R.), Regional Government of Madrid (CM) grant Cloud4BigData (S2013/ICE-2894, cofunded by FSE & FEDER, and grant ANR Quasicool (ANR-12-JS02-011-01).

PODC'15, July 21–23, 2015, Donostia-San Sebastián, Spain.

ACM 978-1-4503-3617-8 /15/07 ...$15.00.
http://dx.doi.org/10.1145/2767386.2767447..

the maximum number of different messages a node can send over its outgoing links in one round. Additionally, we study the effect of the *bandwidth* b, the maximum size in bits of these messages.

The congested clique model with n nodes, and parameters r and b, will be denoted by $\mathrm{RCAST}_n(r \times b)$. Observe that the two extreme cases $r = 1$ and $r = n-1$, which correspond to the broadcast and the unicast communication modes, are the cases already considered in the literature [7]. More precisely, $\mathrm{RCAST}_n((n-1) \times b) = \text{CLIQUE-UCAST}_{n,b}$, and $\mathrm{RCAST}_n(1 \times b) = \text{CLIQUE-BCAST}_{n,b}$.

One important type of problem solved in these models has to do with computing some function or identifying some property of an input graph. In this case, the joint input to the n nodes in the congested clique model is an undirected weighted graph $G = (V, E, \omega)$. The only information each node u initially has is its own ID (the ID of each node is a unique number between 0 and $n-1$), the network size n, the list of IDs of its neighbors v in G, and the weights $\omega(uv)$ of its incident edges. At the end of the protocol every node must know the solution to some particular problem (for instance, a minimum spanning tree of G [10]).

Related Work.

Broadcast. In [4] it was proved that, if the degeneracy m of the input graph G is bounded and known in advance, then it is possible to reconstruct G with a one-round protocol of $\mathcal{O}(\log n)$ message size. Drucker, Kuhn and Oshman [7] gave an upper bound to the round complexity of the subgraph detection problem. They made the following remark: the degeneracy of H-free graphs can be upper bounded in terms of the Turán number $ex(n, H)$, which is the maximal number of edges of an n-node graph which does not contain a subgraph isomorphic to H. Plugging this into the reconstruction protocol introduced by Becker *et al.* [4], they designed a randomized protocol that solves the H detection problem in $\mathcal{O}(ex(n, H) \log^2 n/(nb) + \log^3 n/b)$ rounds with high probability (where b is the number of bits each player can broadcast in each round). Ahn, Guha and McGregor [1, 2] introduced a powerful technique that allows to decide in one round whether G is connected using messages of size $\mathcal{O}(\log^3 n)$, with high probability.

Some negative results have also been obtained. For instance, deciding deterministically in one round whether a graph has a triangle requires messages of size $\Theta(n)$ [4]. On the other hand, if instead of bounding the number of rounds we bound the message size b, then the best known result is the following: detecting deterministically a triangle requires $\Omega(n/(e^{\mathcal{O}(\sqrt{\log n}}b))$ rounds [7].

In [5], the authors consider three variants of the broadcast congested clique model: randomized protocols with public coins, randomized protocols with private coins and deterministic protocols. They showed that this choice affects the message size complexity of some problems. More precisely, they introduced a problem called TRANSLATED-TWINS. They proved that if only one round is allowed, then the message size complexity is $\Theta(n)$ in the deterministic case and $\mathcal{O}(\log n)$ in the randomized, public coin case. For the private coins setting, the message size complexity is lower bounded by $\Omega(\sqrt{n})$ and upper bounded by $\mathcal{O}(\sqrt{n} \log n)$.

Unicast. No lower bounds are known for the general, unicast congested clique model, where nodes may send different messages to each of its neighbors. Drucker, Kuhn and Oshman gave in [7] a possible explanation for the difficulty of finding such bounds. In fact, 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 model has allowed some authors to provide extremely fast protocols for some natural problems: $\mathcal{O}(1)$-round protocols for routing and sorting [9, 11], a $\mathcal{O}(n^{(d-2)/d}/\log n)$- round protocol for finding a particular d-vertex subgraph [6], and a $\mathcal{O}(\log \log \log n)$-round protocol (in expectation) for finding a 3-ruling set [8]. Finally, Dolev, Lenzen, and Peled describe in [6] a protocol for reconstructing deterministically any graph in $\mathcal{O}(|E|/n)$ rounds, which is efficient for sparse graphs.

2. OUR RESULTS

We first prove that there exists a clear round complexity difference between the broadcast model and the model where two different messages are possible. To prove these results, we introduce a useful problem, called edge translation problem and denoted EDGE-TRANS_k.

DEFINITION 1. *Assume that n is even and that $n = 2n'$. Let k be a positive integer, and let the input graph $G = (V, E, \omega)$ be a clique such that $\omega(e) \in \{0, 1\}^k$ for each $e \in E$. (The nodes have IDs in $\{0, \ldots, n-1\}$). The problem* EDGE-TRANS_k *consists in deciding whether for all $0 \leq i, j < n'$, $\omega(\{i, j\}) = \omega(\{i + n', j + n'\})$.*

As mentioned, we use the number of rounds as the complexity measure of a problem.

DEFINITION 2. *Let* PROBLEM *be a problem defined on undirected, simple, weighted graphs. The* round complexity, *denoted $ROUND^n_{r \times b}$(PROBLEM), is the minimal number of rounds for solving* PROBLEM *on any n-node graph with a protocol in $RCAST_n(r \times b)$.*

Round Complexity of Broadcast versus Two Messages.

The first question we would like to answer is how much can be gained if, instead of broadcasting, we have the possibility of sending two different messages in each round? This seems like a simple but fundamental question in order to evaluate the power of the different congested clique models we have defined. To answer this question, at least partially, we first show that, in the model $\mathrm{RCAST}_n(2 \times 1)$, the problem EDGE-TRANS_1 can be solved with a two-round protocol.

THEOREM 1. $ROUND^n_{2\times 1}(\text{EDGE-TRANS}_1) \leq 2$.

PROOF. Let $K_n = (V, E)$ be the n-node clique and let $\omega : E \to \{0, 1\}$. The protocol is as follows.

`Round 1.` Each node i, with $0 \leq i < n'$ sends to each node j, with $n' \leq j < 2n'$, the bit $\omega(\{i, j - n'\})$, and, say, 0 to the other nodes.

`Round 2.` Each node i, with $0 \leq i < n'$, broadcasts the bit 1. Each node j, with $n' \leq j < 2n'$, broadcasts also only one bit, as follows.

- 1 if j received $\omega(\{i, j - n'\}) = \omega(\{i + n', j\}))$ from all $0 \leq i < n'$,

- 0 otherwise.

Clearly, EDGE-TRANS$_1$ is satisfied if and only if no 0 is broadcasted by any node in the second round. Therefore, every node will know the answer after the second round. $\square$

On the other hand, we show that any protocol that solves EDGE-TRANS$_1$ in the model RCAST$_n(1 \times b)$ needs $\Omega(n/b)$ rounds. In other words, in order to solve EDGE-TRANS$_1$ in the broadcast model in a constant number of rounds, the bandwidth b needs to be linear in the number of nodes n. Hence, there is in fact a gain by having the possibility of sending two different packets instead of one.

A Hierarchy of Congested Clique Models.

We continue our study of the power of the congested clique models with various values of range and bandwidth. First, we give a universal bound on the number of rounds to solve EDGE-TRANS$_k$, for any integer $k > 0$, in the model RCAST$_n(r \times b)$ where $r = 2^k$ and $b = \log n$.

THEOREM 2. *For any integer $k > 0$, $r = 2^k \leq n - 1$, $b = \log n \geq k$, $ROUND^n_{r\times b}$(EDGE-TRANS$_k$) ≤ 2.*

Then, we give a lower bound on the number of rounds to solve the same problem EDGE-TRANS$_k$ in RCAST$_n(r \times b)$ when $r = k$ and $b = \log n$.

THEOREM 3. *For any integer $k > 0$, $r = k \leq n-1$, $k > 2$, $b = \log n$, $ROUND^n_{r\times b}$(EDGE-TRANS$_k$) $\geq \frac{(n-2)k}{4k\log n + 2\lceil \log k\rceil n} = \Omega\left(\frac{k}{\log k}\right)$.*

These two bounds combined show the existence of a hierarchy among the classical congested clique models (those with $b = \log n$), since the complexity in number of rounds to solve the same problem can vary significantly as a function of the range r. In particular, we have the following result.

COROLLARY 1. *For every integer $k = \omega(1)$ such that $2^k \leq n - 1$, there is a gap of $\Omega\left(\frac{k}{\log k}\right) = \omega(1)$ on the round complexity of problem* EDGE-TRANS$_k$ *between the congested clique models $RCAST_n(2^k \times \log n)$ and $RCAST_n(k \times \log n)$.*

This corollary shows that there is a non-constant gap of complexity $\omega(1)$ in the number of rounds required by the model RCAST$_n(k \times \log n)$ with respect to the model RCAST$_n(2^k \times \log n)$ to solve the same problem EDGE-TRANS$_k$, as long as $k = \omega(1)$.

Simulation between Congested Clique Models.

Finally, we give techniques to simulate protocols designed for a model RCAST$_n(r' \times b')$ in another model RCAST$_n(r \times b)$, where $2 \leq r \leq r' < n$. These techniques yield upper bounds on the complexity gap between two models for any problem. In particular, our results show that this gap between models RCAST$_n(r' \times b')$ and RCAST$_n(r \times b)$ with $2 \leq r \leq r' < n$ is upper bounded by $\min\left(\left\lceil \frac{b'}{\lfloor \log r\rfloor}\right\rceil, \left\lceil \frac{r'}{r-1}\right\rceil \left\lceil \frac{b'}{b}\right\rceil\right)$. This bound becomes $\left\lceil \frac{b'}{b}\right\rceil$ for the special case in which $r = 2^b$ and $r = r'$.

3. REFERENCES

[1] K. J. Ahn, S. Guha, and A. McGregor. Analyzing graph structure via linear measurements. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 459–467, 2012.

[2] K. J. Ahn, S. Guha, and A. McGregor. Graph sketches: Sparsification, spanners, and subgraphs. In *Proceedings of 31st Symposium on Principles of Database Systems (PODS 2012)*, pages 5–14, 2012.

[3] F. Becker, A. Kosowski, N. Nisse, I. Rapaport, and K. Suchan. Allowing each node to communicate only once in a distributed system: shared whiteboard models. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architecture (SPAA 2012)*, pages 11–17, 2012.

[4] F. Becker, M. Matamala, N. Nisse, I. Rapaport, K. Suchan, and I. Todinca. Adding a referee to an interconnection network: What can (not) be computed in one round. In *Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2011)*, pages 508–514, 2011.

[5] F. Becker, I. Rapaport, P. Montealegre, and I. Todinca. The simultaneous number-in-hand communication model for networks: private coins, public coins and determinism. In *Proceedings of the 21st International Colloquium on Structural Information and Communication Complexity (SIROCCO 2014)*, pages 83–95, 2014.

[6] D. Dolev, C. Lenzen, and S. Peled. Tri, tri again: Finding triangles and small subgraphs in a distributed setting. In *Proceedings of the 25th Symposium on Distributed Computing (DISC 2012)*, pages 195–209, 2012.

[7] A. Drucker, F. Kuhn, and R. Oshman. On the power of the congested clique model. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC 2014)*, pages 367–376, 2014.

[8] J. W. Hegeman, S. V. Pemmaraju, and V. B. Sardeshmukh. Near-constant-time distributed algorithms on a congested clique. In *Proceedings of the 27tth Symposium on Distributed Computing (DISC 2014)*, pages 514–530, 2014.

[9] C. Lenzen. Optimal deterministic routing and sorting on the congested clique. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC 2013)*, pages 42–50, 2013.

[10] Z. Lotker and E. Pavlov. MST construction in $O(\log \log n)$ communication rounds. In *Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architecture (SPAA 2003)*, pages 94–100, 2003.

[11] B. Patt-Shamir and M. Teplitsky. The round complexity of distributed sorting. In *Proceedings of the 30th ACM Symposium on Principles of Distributed Computing (PODC 2011)*, pages 249–256, 2011.

[12] D. Peleg. *Distributed computing: a locality-sensitive approach.* Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.