# Tighter Worst-Case Bounds on Algebraic Gossip 

Bernhard Haeupler, MIT, haeupler@mit.edu


#### Abstract

Gossip and in particular network coded algebraic gossip have recently attracted attention as a fast, bandwidthefficient, reliable and distributed way to broadcast or multicast multiple messages. While the algorithms are simple, involved queuing approaches are used to study their performance. The most recent result in this direction shows that uniform algebraic gossip disseminates $k$ messages in $O(\Delta(D+k+\log n))$ rounds where $D$ is the diameter, $n$ the size of the network and $\Delta$ the maximum degree.

In this paper we give a simpler, short and self-contained proof for this worst-case guarantee. Our approach also allows to reduce the quadratic $\Delta D$ term to $\min \{3 n, \Delta D\}$. We furthermore show that a simple round robin routing scheme also achieves $\min \{3 n, \Delta D\}+\Delta k$ rounds, eliminating both randomization and coding. Lastly, we combine a recent non-uniform gossip algorithm with a simple routing scheme to get a $O\left(D+k+\log ^{O(1)}\right)$ gossip information dissemination algorithm. This is order optimal as long as $D$ and $k$ are not both polylogarithmically small.


## I. Introduction

Broadcast and multicast are fundamental communication primitives with many practical applications like the maintenance of distributed databases or content distribution networks. With modern networks increasing in size they often become less reliable and behave more like distributed systems. This increases the need for reliable ways to disseminate information without knowing the network topology.

Gossip protocols in which nodes forward information to randomly chosen neighbors have been developed as a powerful new paradigm in this direction. Their randomized approach has been shown to disperse information quickly in many networks while keeping the total number of messages and the congestion on any particular link small.

When more information than can be fit in a packet is to be distributed selecting the right piece of information to forward can be hard. Algebraic gossip, an adaptation of random linear network coding, has been proposed as a solution and it was shown that in a complete network it outperforms routing [1]. Subsequent works like [2], [3], [4] extend this idea to general networks. These works also follow the approach of [1] to understand the performance of algebraic gossip by analyzing the queuing behavior of innovative packets. The most recent work in this line [3], [4] use Jackson's theorem to study the worst-case performance of algebraic gossip in terms of network diameter and maximum degree. This will also be the main focus of this paper.

We show that the projection analysis technique of [5] provides easier means to understand algebraic gossip. With Theorem 1 we give a simple, short and self-contained proof for the main result in [4], which itself generalizes the main result in [3]. In Theorem 2 we then further tighten Theorem 1. Next, we observe that the worst-case performance of algebraic
gossip can also be matched by uniform gossip without coding (Theorem 4) or even a simple deterministic round robin routing protocol (Theorem 3). Appealing to the optimality of algebraic gossip [6] leads then to yet another alternative proof for Theorems 1 and 2] Lastly, we follow the idea of [4] to bridge the gap to more structured broadcast protocols by giving an efficient non-uniform gossip protocol. Our Theorem 5 shows that this protocol achieves an optimal performance in any network up to an additive polylogarithmic term. All previous algorithms had gaps of up to $\Theta(n / \log n)$.

## II. Gossip Algorithms and the Multicast Problem

As in [4] we assume a network to be modeled by an undirected graph $G=(V, E)$ with $n=|V|$ nodes, node degree $\Delta_{u}$ for $u \in V$, maximum degree $\Delta=\max _{u} \Delta_{u}$ and diameter $D$. During a gossip protocol nodes communicate in synchronous rounds in which each node can initiate one bidirectional exchange of a packet with one neighbor. The goal of the information dissemination, rumor spreading or $k$-message multicast task is to spread $k$ initially distributed messages to all nodes as fast as possible and with high probability (whp), that is, with probability at least $1-1 / n$. We assume that the messages have equal size and that a packet can fit one message plus some header information (which becomes negligibly small for large message/packet sizes). A gossip solution to the multicast problem specifies both, which node contacts which in each round and what is sent in each transmitted packet. The most studied instantiations for these choices are uniform gossip, that is, each node contacts a uniformly random neighbor in each round, and algebraic gossip, that is, nodes perform random linear network coding on the messages to create new packets. In random linear network coding nodes send out random linear combinations of messages (over some finite field $G F(q)$ ) together with a vector containing all coefficients used. Each new packet is then simply created as a uniformly random linear combination of already received packets or initially known messages; to decode, nodes perform Gaussian elimination.

## III. A Simpler and Stronger Proof for Uniform Algebraic Gossip

The main result in [4] is the following worst-case guarantee on the performance of algebraic gossip in any network:

Theorem 1. Uniform algebraic gossip over $G F(2)$ disseminates $k$ messages whp in no more than $O(\Delta(D+k+\log n))$ time.

The proof is somewhat involved and uses Jackson's queuing theorem. Here, we show that one can alternatively obtain the same result using the projection analysis of [5]. In particular,
one can essentially get it by setting the min-cut probability $\gamma$ in [5, Lemma 6] to be $1 / \Delta$. Next, we give a simple, short and completely self-contained proof based on [5] Theorem 3]:

Proof: We say a node knows a coefficient vector $\mu \in$ $\{0,1\}^{k}$ if it has received a packet with a coefficient vector that is non-perpendicular to $\mu$ (over $G F(2)$ ). We claim that for any non-zero vector $\mu$ the probability that any fixed node $v$ does not learn $\mu$ within $O(\Delta(D+k+\log n))$ rounds is at most $2^{-(k+2 \log n)}$. Then, a union bound over all nodes and all $2^{k}$ coefficient vectors shows that whp all nodes know all vectors. From this it is easy to conclude that all nodes can decode.

To prove this claim we look at a shortest path $P$ from $v$ to a node that initially knows $\mu$ (that is, that starts with a message with a non-zero coefficient in $\mu$ ). At any round $t$ let node $u$ be the closest node to $v$ on $P$ that knows $\mu$. There is a $1 / \Delta_{u}$ chance that $u$ contacts the next node on the path and independently a chance of $1 / 2$ that $u$ sends out a packet with a coefficient vector that is non-perpendicular to $\mu$. Thus, in any round independently with probability at least $\frac{1}{2 \Delta}$ knowledge of $\mu$ makes progress on $P$. A Chernoff bound shows that the probability that less than $D$ progress is made in $16 \Delta(D+k+$ $\log n)$ rounds when $8(D+k+\log n)$ successes are expected is at most $2^{-(k+2 \log n)}$ as claimed.

Note, that even for constant $k$ Theorem 1 results in a quadratic bound of $\Delta D=\Theta\left(n^{2}\right)$ for many networks. Next, we tightening the bound of Theorem 1 and show that small number of messages (e.g., $k=O(\log n)$ ) never take more than $O(n \log n)$ time:
Theorem 2. Uniform algebraic gossip disseminates $k$ messages whp in at most $O(\min \{n, \Delta D\}+\Delta(k+\log n))$ time.

Proof: In the proof of Theorem 1 let $X_{1}, X_{2}, \ldots$ be the number of rounds spent for knowledge about $\mu$ to successfully make the $i$ th step along $P$ towards $v$ and let $X=\sum_{i} X_{i}$ be the total number of rounds needed. Note that each $X_{i}$ is an independent geometric random variable with success probability $\frac{1}{2 \Delta_{u}}$ and thus $E[X]=\sum_{i} E\left[X_{i}\right]=\sum_{u \in P} 2 \Delta_{u} \leq$ $2 D \Delta$. Since $P$ is a shortest path, the sum of node degrees along $P$ is at most $3 n$ because every node can be adjacent to at most three consecutive nodes on $P$ without creating a shortcut. We thus also have $E[X] \leq 6 n$. By applying Markov's inequality to $(1-1 / \Delta)^{-2 X}$ one can show that $P[X \geq 2(E[X]+t)]<(1-1 / \Delta)^{-t}$. Using this with $t=\Delta(k+\log n)$ replaces the Chernoff bound argument in the second paragraph of the proof of Theorem 1 .

## IV. Gossip without Coding or Randomization

Among the main motivations for using randomized gossip is that its randomization adds robustness and often avoids bottlenecks in the network. In regular, well-connected networks this allows for much faster dissemination times than the bound given in Theorem 2. Adding coding on top of this (in particular, network coded algebraic gossip) allows mixing and efficient diffusion of information which often significantly increases throughput [5].

In this section, we show that if one just wants to achieve the worst-case running time of Theorem 2 neither randomization
nor coding is needed. Instead the following simple deterministic round robin routing achieves the same performance (with no probability of failure): Each node $u$ repeatedly contacts its neighbors such that every neighbor is contacted at least every (order) $\Delta_{u}$ rounds. A node furthermore can forward any packet as long as the same packet is not forwarded twice to the same node. One simple protocol that does this efficiently is prioritized round robin routing. In this protocol each node fixes a cyclic order of its neighbors and simply contacts the next node in its list in every round. Each node $u$ furthermore keeps a counter how often it has sent out each packet and sends out the lexicographically first packet it has not sent out $\Delta_{u}$ often already (if such a packet exists).

Theorem 3. (Prioritized) round robin routing disseminates all $k$ messages in at most $\min \{3 n, \Delta D\}+\Delta k$ time.

Proof: Throughout the proof we use the following notion of distance between any two nodes $u$ and $v$ : Let $\operatorname{dist}(u, v)$ be the smallest number such that there is a path $P$ between $u$ and $v$ with $\sum_{w \in P-\{v\}} \Delta_{w} \leq \operatorname{dist}(u, v)$. Note that $\operatorname{dist}(u, v) \leq$ $\min \{3 n, \Delta D\}$ as shown in the proof of Theorem 2. We claim that for any nodes $u, v$, any message $m$ and any $i \geq 0$ if node $u$ knows message $m$ initially then at time $t=\operatorname{dist}(u, v)+\Delta i$ node $v$ will either know message $m$ or at least $i+1$ messages smaller than $m$. Note that this claim also implies the theorem since at time $\min \{3 n, \Delta D\}+\Delta k$ any node has either received every message $m$ or $k$ (smaller) messages.

We prove the claim by induction on $t$. Nothing needs to be shown for $t=0$. For the induction step consider any $u, v, m$ and $t=\operatorname{dist}(u, v)+\Delta i$ and let $v^{\prime}$ be the neighbor of $v$ on a shortest distance path from $u$ to $v$. By induction hypothesis we get that for any $j \leq i$ node $v^{\prime}$ at time $d\left(u, v^{\prime}\right)+\Delta j$ either knows $m$ or $j+1$ messages smaller than $m$. From this it is clear that at time $d\left(u, v^{\prime}\right)+\Delta_{v}^{\prime}=d(u, v)$ node $v^{\prime}$ has sent either message $m$ or one message smaller than $m$ to $v$ and similarly for any $j<i$ at time $d\left(u, v^{\prime}\right)+\Delta j+\Delta_{v}=d(u, v)+\Delta j$ node $v$ has sent either $m$ or $j+1$ messages to $v$ that are smaller than $m$. For $j=i$ this is exactly the claim that needed to be proved.

We remark that a similar performance guarantee also holds for randomized uniform gossip. It can be obtained along the same line as the proof for Theorem 3. We state the result next but omit the proof.
Theorem 4. Prioritized uniform gossip whp routes all messages to all nodes in $O(\min \{n, \Delta D\}+\Delta(k+\log n))$ time.

Given that these two simple routing schemes achieve a good worst-case performance one would expect that adding algebraic gossip on top of it should not harm the performance. Indeed we can make this formal and appeal to the optimality of algebraic gossip [6] to give yet another simple and alternative proof for Theorem 2:

Proof of Theorem 2: In contrast to before we require the field size used for coding to be at least $q=n^{2}$ which gives rise to $2 \log n$-size coefficients. In [6] it is shown that with probability $1-n / q$ algebraic gossip completes in exactly optimal time, that is, at the first time it is possible
in hindsight to route all messages to each node individually. Furthermore, Theorem 4 shows how to route the messages whp in $O(\min \{3 n, \Delta D\}+\Delta(k+\log n))$ rounds via the uniform gossip exchanges. Thus with probability $1-n / q-1 / n=$ $1-2 / n$ uniform algebraic gossip also completes in this time.

## V. FASTER NON-UNIFORM GOSSIP

Lastly, [4] shows that fast non-uniform gossip protocols can be obtained by first running a 1-message broadcast gossip algorithm and then performing algebraic gossip along an induced spanning-tree. More specifically, for any broadcast algorithm with running time $B$ (and given a leader) this leads to a $O(B+k+\log n)$ solution to the $k$-message multicast problem considered here. Instantiating this with the broadcast protocol from [7] it is shown that an $\Theta\left(D+k+\Phi^{-1} \log n\right)$ algebraic gossip protocol is possible for graphs with weak conductance $\Phi$. Unfortunately, there are graphs for which $\Phi^{-1} \log n=\Theta(n)$ even so the network diameter is logarithmic (for example, a balanced binary tree). For these graphs a gap of up to $\Theta(n / \log n)$ between the upper bound and the $\Omega(D+k)$ lower bound remains. The following lemma improves upon this using [8], a recent strengthening of [7]:

Theorem 5. There is a gossip algorithm that whp routes $k$ messages to all nodes in $O\left(D+k+\log ^{O(1)}\right)$ time in any network.

Proof: In [8] a broadcast gossip algorithm with time $O\left(D+\log ^{O(1)}\right)$ is given. We first run this algorithm with each node using its ID as a message. After completion every node declares the node it first learned the smallest ID from as its parent (breaking ties arbitrarily). This induces a spanning tree. In all future rounds, each node initiates a bidirectional exchange with its parent forwarding any message not sent by it before (if such a message exists). A standard pipelining proof (similar to Theorem 3) shows that after $O\left(k+D^{\prime}\right)$ time all nodes know about all messages, where $D^{\prime}$ is the diameter of the tree. Since a message can travel at most one step per round during the initial broadcast we have $D^{\prime}=O\left(D+\log ^{O(1)}\right)$ which completes the proof.

## VI. Conclusion

We have given simple, short and tighter proofs for the worstcase performance of algebraic gossip using the projection analysis technique of [5]. This improves over the recent results in [4]. We could furthermore show that this worstcase performance is also met by a simple round robin routing scheme that eliminates both randomness and coding. A similar routing scheme has been studied intensively under the term quasirandom rumor spreading [9] and has been shown to perform surprisingly well in many topologies. Lastly, bridging the gap to more structured broadcast protocols we have shown a non-uniform gossip protocol that achieves order optimal (up to an additive polylogarithmic term) distributed information spreading. All these results do not take any unreliability or changes in the network topology into account. This is unfortunate since fault-tolerance with respect to dynamic or
unreliable topologies is one of the main motivations behind pursuing randomized gossip approaches. Some work in this direction can be found in [5] but any further approaches to formalize and study dynamic topologies and the reliability of the algorithms presented here is an important open problem of interest.

## REFERENCES

[1] S. Deb, M. Médard, and C. Choute, "Algebraic gossip: a network coding approach to optimal multiple rumor mongering," IEEE/ACM Trans. Netw., vol. 14, no. SI, pp. 2486-2507, June 2006.
[2] D. Mosk-Aoyama and D. Shah, "Information dissemination via network coding," in IEEE International Symposium on Information Theory. IEEE, 2006, pp. 1748-1752.
[3] M. Borokhovich, C. Avin, and Z. Lotker, "Tight bounds for algebraic gossip on graphs," in IEEE International Symposium on Information Theory, 2010, pp. 1758-1762.
[4] C. Avin, M. Borokhovich, K. Censor-Hillel, and Z. Lotker, "Order optimal information spreading using algebraic gossip," in PODC, 2011, pp. 363372.
[5] B. Haeupler, "Analyzing network coding gossip made easy," in STOC, 2011, pp. 293-302.
[6] B. Haeupler, M. Kim, and M. Medard, "Optimality of Network Coding with Buffers," in ITW, 2011, pp. 533-537.
[7] K. Censor-Hillel and H. Shachnai, "Fast information spreading in graphs with large weak conductance," in SODA, 2011, pp. 440-448
[8] K. Censor-Hillel, B. Haeupler, J. Kelner, and P. Maymounkov, "Global Computation in a Poorly Connected World: Fast Rumor Spreading with No Dependence on Conductance," in STOC, 2012.
[9] B. Doerr, T. Friedrich, and T. Sauerwald, "Quasirandom rumor spreading," in ACM-SIAM Symposium on Discrete Algorithms, 2008, pp. 773781.

