

# Ant Routing scalability for the Lightning Network

Cyril Grunspan, Gabriel Lehericy <sup>\*,1</sup> and Ricardo Pérez-Marco <sup>†,2</sup>

<sup>1</sup>De Vinci Research Center, Paris-La Défense, France

<sup>2</sup>CNRS, IMJ-PRG, Université Sorbonne, Paris, France

February 4, 2020

## Abstract

The ambition of the Lightning Network is to provide a second layer to the Bitcoin network to enable transactions confirmed instantly, securely and anonymously with a world scale capacity using a decentralized protocol. Some of the current propositions and implementations present some difficulties in anonymity, scaling and decentralization. The Ant Routing algorithm for the Lightning Network was proposed in [?] for maximal decentralization, anonymity and potential scaling. It solves several problems of current implementation, such as channel information update and centralization by beacon nodes. Ant Routing nodes play all the same role and don't require any extra information on the network topology beside for their immediate neighbors. The goal of LN transactions are completed instantaneously and anonymously. We study the scaling of the Ant Routing protocol. We propose a precise implementation, with efficient memory management using AVL trees. We evaluate the efficiency of the algorithm and we estimate the memory usage of nodes by local node workload simulations. We prove that the number of transactions per second that Ant Routing can sustain is of the order of several thousands which is enough for a global payment network.

**Keywords.** Bitcoin, Lightning Network, Ant Routing, AVL Trees

## 1 Introduction

The Lightning Network was introduced in order to address Bitcoin scalability [?]. It consist on a “layer 2”-network set on top of Bitcoin's network which facilitates micro-payments between bitcoin users without needing to broadcast the transaction on the blockchain.

The Lightning Network is not a complete graph, but has the property of channel composition. This allows any user on the network to make secure payments to any other node in the same connected component of the network, even when no direct

---

\*Emails: [cyril.grunspan@devinci.fr](mailto:cyril.grunspan@devinci.fr), [gabriel.lehericy@devinci.fr](mailto:gabriel.lehericy@devinci.fr),

†Email: [ricardo.perez.marco@gmail.com](mailto:ricardo.perez.marco@gmail.com)

payment channel exists. For example, assume that Alice and Bob are both LN users, but have no direct payment channel between them. Now assume that there is a third user, Charlie, who has a direct channel with Alice and a direct channel with Bob. Then Alice can pay to Bob via Charlie: Alice sends the payment to Charlie, who will then forward it to Bob. All these payments are realized in a secured way by the protocol of channel composition. This allows any LN node to send money to any other LN node when the network is connected.

Obviously this requires finding a payment route from Alice to Bob, hence the need for a good routing algorithm. Ideally, one wants to have a routing algorithm which is safe, decentralized, anonymous, and only requires minimum knowledge of the network from all participants. In [?], the authors proposed a new routing algorithm, called *Ant Routing*. Ant Routing is a completely decentralized routing algorithm which does not require beacon nodes nor routing tables. Instead, each node of the network executes the exact same task which requires no knowledge of the topology of the network. Moreover, it is worth noting that using Ant Routing on the lightning network would have a positive impact on fees. Indeed, unlike the current routing mechanism in use, ant routing does not allow for a node to ask for fees proportional to the transaction amount. This also strengthens the network's decentralization.

The goal of this article is to estimate the performance of Ant Routing, and propose more precise and efficient implementation. In particular, we investigate if this routing algorithm is compatible with the original ambition of the Lightning Network to solve Bitcoin scalability at planetary level, hence providing an instantaneous decentralized global payment network. In practice, the question is whether Ant Routing algorithm can allow the network to sustain several thousands of transactions per second. This article gives a positive answer to this question. The conservative estimates presented in Sections ?? and ?? show that Ant Routing can sustain more than 10.000 transactions per second.

The paper is organized as follows. In Section ??, we give an overview of routing protocols which have been proposed for the Lightning Network and of algorithms based on ant behavior. Section ?? explains the ant routing protocol. We start by briefly recalling the rough ideas of ant routing as presented in [?]. We then give a more detailed description of the Ant Routing protocol. In particular, we give a precise list of data that must be kept in memory and a list of tasks performed by nodes. In Section ?? we estimate the duration for seeds to be kept in memory. In section ?? we give a general argument to determine network capacity for networks with a similar structure than Bitcoin. This argument applies to Ant Routing provided that the local process time for transactions in the nodes is neglectable or of the same order of magnitude than the propagation time across the network. We propose efficient local algorithms and run simulations. In Section ??, we review AVL trees, which is the data structure used for the Ant Routing implementation. In Section ??, we propose a concrete implementation of Ant Routing. Finally, in Section ?? we estimate the computing time and memory usage. Putting together these results, we conclude by an estimate on the number of transactions per second that the proposed implementation of Ant Routing can sustain.

## 2 Overview of routing proposals and existing ant algorithms

The white paper for the Lightning Network is sketchy about routing and no specific routing protocol is proposed [?]. The authors believe that routing tables are necessary, and in fact, it seems every routing protocol proposed for the Lightning Network prior to [?] makes use of routing tables. We review in this section some of these proposals. A review of routing solutions for the lightning network can also be found in [?]. One should keep in mind that, in accordance to Bitcoin philosophy, one wishes the Lightning Network to remain decentralized and to preserve confidentiality of payments.

One major constraint for routing in the Lightning Network is to find a path with sufficient capacity. It seems natural to try to solve this problem via a distributed max-flow approach which computes the path with the largest capacity [?, ?, ?]. However, max-flow routing performs poorly for large network, as its run time grows quadratically on the number of channels [?]. Another possibility is to use *beacons* instead. A beacon is a special node whose role is to enhance the other nodes' visibility of the network and help them in route computation. One example of an algorithm which uses beacons is *Flare* [?]. Another classical approach to routing is *landmark routing*. Landmark routing has been adapted to payment network in algorithms such as *SilentWhispers* [?]. The idea of landmark routing is to compute a route from the source to the destination via a previously chosen node, called a *landmark*, which is usually a highly connected node. The main drawback of having a system rely on landmarks or beacons is the potential threat to decentralization, as well as its fragility to attacks.

Another approach to routing is *embedding-based routing*. The idea is to embed nodes into a vector space such that the hop distance between nodes is reflected by their distance in the vector space. For routing, the source node chooses the neighbor whose vector representation is the closest to the vector representation of the destination node. Each node in the path then repeats this strategy. This requires regularly updating information about the network. An example of such routing algorithm is *Speedy Murmurs*, which is based on *Voute* [?, ?].

Recent attention has been given to payment routing solutions that keep channels balanced. Keeping channels balanced is considered important in the current conception of the Lightning Network<sup>1</sup>: if a channel becomes too unbalanced, then payment can only flow in one direction, and this restricts payment routes that can no longer use that channel. In [?], the authors introduce an interesting idea which considers channels' balance for route selection called *Spider*. More precisely, *Spider* splits payments into several smaller payments, each of which can be sent through a different route. The route for each of these payment is chosen to optimize the re-balancing of unbalanced channels.

Note that all of these routing strategies either require giving some nodes a special role (as a landmark or a beacon) or require each node to keep some knowledge of the network topology. This knowledge is a vector of attack for the Lightning Network. One of the reasons for the proposal of Ant Routing is to obfuscate the topology of the network: It requires no landmark and requires no knowledge on the topology, while

---

<sup>1</sup>A network with only one directional channels is theoretically more limited, but in practical terms can be useful since most of the channels are used unidirectionally.

at the same time preserves anonymity and decentralization. The decentralization of the algorithm is achieved by making every node play exactly the same role in the routing process and using only knowledge about its neighbors.

The algorithm is inspired by the behavior of ants. Ant colonies have raised the interest of ethologists by their ability to optimize their route from nest to food source. Although each ant individually seems to follow a random motion, their collective behavior finds efficiently the shortest path from their nest to a food source. This is achieved through a “*stigmergic*” communication of the ants with their environment through pheromones [?]. More precisely, each walking ant leaves a trail of pheromone, which has the attracts other ants to its path.

The author of [?] and [?] describes the following experiment to explain this remarkable observed ant optimization behavior. We study an ant colony connected to a food source by two paths, a short and a longer one. At the outset of the experiment, each ant walks in a random direction, choosing each path randomly with equal probability. As they walk, the ants leave a pheromone trail behind. Because ants will travel faster on the shorter path than on the longer one, pheromones accumulate faster on the shorter path. The concentration of pheromone on the shorter path increases, so does the probability that ants to choose this path. This leads to a positive feedback, and eventually few ants will choose the longer path.

The author of [?] had the idea to simulate this ant behavior with a multi-agent stochastic learning model to solve optimization problems. An application to network routing problems is given in chapter 6 of [?].

The Ant Routing algorithm is inspired from these ant behavior models, but there are some major differences with the routing propositions in [?] and [?]. Our algorithm is simpler and deterministic, and not probabilistic. It only uses the idea of leaving pheromones to trace a path, but does not consider accumulation of pheromones on a path. Eventually, the model could be enhanced taking into account pheromone accumulation to improve the efficiency. The paths are not chosen randomly. The pheromones are flooding the network. Also, apparently the applications of ant algorithms to routing problems in [?] and [?] use routing tables, which we want to avoid.

### 3 Ant routing protocol

We briefly recall the main ideas of ant routing as presented in [?]. We assume that Alice and Bob both have a node on the Lightning Network and they have a direct communication channel but not necessarily a payment channel. Alice wants to pay Bob via the Lightning Network, so they need to find a payment route in the Lightning Network.

Alice and Bob start by agreeing on a random number  $S$ , chosen big enough to avoid collisions. Alice creates the number  $S(0) = 0 \frown S$  and Bob creates the number  $S(1) = 1 \frown S$ .  $S(0)$  and  $S(1)$  are called *pheromone seeds*. Alice and Bob send their respective seeds to the network through their neighbors. Pheromone seeds propagate on the network being successively forwarded by all nodes to their neighbors. When  $S(0)$  and  $S(1)$  meet at some node we have a match. Each pheromone seed is forwarded with a counter associated. This counter is increased by 1 at each hop. If a node receives the same pheromone seed several times, then only the pheromone seed with lower counter is stored and forwarded. The nodes keep track from whom

they received the stored seeds. When we have a match, the node where it occurs creates matched seeds  $M(0)$  and  $M(1)$ , and sends back to the respective neighbors from whom he received  $S(0)$  and  $S(1)$ . Nodes receiving a matched seed forward them in the same way, and after a finite number of steps, bounded by the counters, the matched seeds reach Alice and Bob.

After some time, Alice will receive several matched seeds, each corresponding to a payment route. Alice is free to choose any of these matched seed as a payment route. Note that, in order to preserve anonymity, it is desirable for Alice to choose a route with at least two intermediary nodes. She selects one, by a selection algorithm that she is free to decide. Creates a confirmed seed  $C$  and sends  $C$  back to her neighbor from whom she received the matched seed. Nodes forward confirmed seed they receive as for matched seeds. Hence, the confirmed seed will reach Bob through a unique path. Once Bob has received  $C$ , he informs Alice and she can proceed with the payment. The payment is then done via this path.

Notice that this “path discovering algorithm” finds a connecting path without anyone knowing this path. The only knowledge the participating nodes have is about their neighbors, and they don’t even know if these are Alice or Bob.

During the pheromone phase, a node may be tempted to cheat with the counter in order to make the path more attractive. For this reason, we require a verification phase after the confirmation phase and before the payment. The point of this verification phase is to check that no node on the path indicated by  $C$  has cheated with the counter. This can be done in the following way: during the confirmation phase, each node, starting from Alice neighbor, is asked to generate a random number. This random number is appended to a list  $l$ , which lists all the random numbers from previous nodes. Then the nodes forwards  $l$  along with  $C$  to the next node. When  $C$  and  $l$  reach Bob, Bob sends  $l$  back and communicates the list of random numbers to Alice. The nodes forward it back only if their random number is in the list. Then  $l$  can only reach back Alice through the nodes in the path if nobody has cheated on the counter.

We make precise the task done by the nodes in the protocol. It is important to note that all nodes, apart from Alice and Bob in their payment, execute the same task. Moreover, the nodes have no knowledge about the topology of the network. The only topological information a node needs is who are its direct neighbors.

We list in the next section which data the nodes must keep in memory, and which must be relayed (Section ??). We then give the algorithms that the nodes perform in each phase of the ant routing protocol: the pheromone phase, the match phase, the confirmation phase and the counter check phase (Section ??).

### 3.1 Memory data

Each node allocates some memory space for the routing task. This memory space contains three trees. The first one is used to store pheromone seeds, the second one for matched seeds, and the third one for confirmation seeds. When she starts her routing task, Alice creates a fourth tree in her memory, which we call the **special match tree**. This tree will be destroyed once Alice has completed her payment. This tree is used to store the matched seeds created from her own pheromone seed. These matches are therefore stored separately from other matched seeds.

Each seed comes with some data. More precisely, for each pheromone seed  $P$ , a node stores:

- (1) The pheromone seed  $P$ .
- (2) A counter  $c$ .
- (3) The id  $s$  of the node from which it received the seed .
- (4) The remaining fees  $f$ : at the start, Alice chooses a maximal amount  $f_{max}$  of fees which she is willing to pay for the transaction. If  $g$  denotes the sum of all the fees of all nodes through which the seed has traveled, then the remaining fees  $f$  is  $f = f_{max} - g$ .
- (5) The amount  $a$  of the transaction.

All this information is required for the functioning of the ant routing algorithm. The id  $s$  is needed for the “match” phase, when the matched seed will have to trace the way back to Alice. The amount of the transaction is used by the relaying nodes to select channels with sufficient funds. The remaining fees are used by Alice at the start of the confirmation phase in her choice of the matched seed. Finally, the counter is used to tame the flood of data on the network (see section ??).

The node keeps  $s$  to itself to preserve anonymity. The nodes also relay one extra piece of information which is not stored in memory: a timestamp  $t$ , which is decided upon by Alice at the creation of the seed. This timestamp allows the node to know the age of the seed, which is necessary information so that the the node knows when to remove the seed from its memory. Pheromone seeds are therefore forwarded as messages of the form:  $(P, c, f, a, t)$  (pheromone seed, counter, remaining fees, payment amount and timestamp).

For each matched seed  $M$ , the node stores the following items:

- (1) A “match identifier”  $Id$ .
- (2) The “target” of the match, which is the name of the next node in the path from Alice to Bob which  $M$  indicates.

The number  $Id$  is a random number generated at the creation of the match. Its purpose is to distinguish between different matches which may have been created from the same pheromone seed. Note that it is not necessary to store  $M$  itself, although  $M$  needs to be forwarded. There are also four extra pieces of information which need to be forwarded but do not have to be kept in memory. These are:

- (1) Two counters  $c$  and  $C$ . The counter  $C$  gives the number of intermediate nodes from Alice to Bob in the path indicated by  $M$ . The counter  $c$  is decreased by 1 at each hop. Its purpose is to be compared with the counter of the pheromone seed associated to  $M$  (see section ?? below).
- (2) The total remaining fees  $F$  associated to  $M$ . We have  $F = 2f_{max} - G$ , where  $G$  denotes the sum of all fees from Alice to Bob in the path indicated by  $M$ . The purpose of  $F$  is to be transmitted to Alice, so that she can choose the path with the lowest fees.
- (3) A timestamp  $t$ , which is the timestamp of the corresponding pheromone seed.

Therefore, matched seeds are forwarded to other nodes via a message of the form  $(M, Id, c, C, F, t)$ . Note that, in Alice’s special match tree,  $F$  and  $C$  are also stored.

A confirmation seed is simply a random number  $Id$ . For each confirmation seed, the node stores:

- (1) An identifier  $Id$ . This is the identifier of the matched seed chosen by Alice.
- (2) The “target”, i.e. the next node in the path from Alice to Bob.
- (3) An integer  $check$ . This is a randomly generated integer meant for the final counter check phase.

Confirmed seeds are forwarded in messages of the form  $(Id, l, t)$ , where  $l$  is a list of integers and  $t$  is the timestamp of the corresponding pheromone seed. More precisely,  $l$  is the list of the “ $check$ ” integers of all preceding nodes.

**Remark 3.1:** The node does not store the timestamp of a given seed. However, the location of a given seed in memory depends on its timestamp  $t$ . Indeed, we will see in Section ?? that seeds are stored in several trees, each of which corresponds to a time interval to which the seed’s timestamp belongs. It follows that, in order to find a seed in its memory, the node needs to know the associated timestamp. This explains why the timestamp also has to be forwarded with each type of seed.

## 3.2 Algorithms

The ant routing protocol has four phases: the “pheromone” phase, the “match” phase, the “confirmation” phase and the “counter check” phase. We now describe these phases one by one in this order.

Alice chooses two random numbers  $S$  and  $c_0$ , the timestamp  $t$ , and the maximal amount of fees  $f_{max}$  which she is willing to pay, and communicates to Bob this data. The number  $S$  will be used to form the pheromone seeds and  $c_0$  will be used the starting value for the counter<sup>2</sup>. The purpose of choosing a random number as the starting value of the counter is to preserve Alice’s and Bob’s anonymity. If we start the counter at 0, then the neighbors of Alice and Bob will know that they are originating the transaction. The timestamp  $t$  serves as the time at which Alice and Bob broadcast their pheromone seeds.

Assume that  $S, c_0$  and  $t$  have been agreed upon by Alice and Bob, and that Alice has chosen  $f_{max}$ . Alice then creates  $P(0) := 0 \frown S$  and stores the information  $(P(0), c_0, 0, f_{max}, a)$  in her pheromone tree, where  $a$  denotes the amount of money which she wants to send Bob. The “0” indicates that the seed has no sender, so the seed originates from Alice. Similarly, Bob creates  $P(1) := 1 \frown S$  and stores the information  $(P(1), c_0, 0, f_{max}, a)$  in his pheromone tree. At time  $t$ , Alice sends the message  $(P(0), c_0, f_{max}, a, t)$  to all her neighbors, and Bob sends the message  $(P(1), c_0, f_{max}, a, t)$  to all his neighbors.

After this initial step, the nodes will keep forwarding  $P(0)$  and  $P(1)$  until matches occur, increasing the counter by 1 at each hop. In general, the same node will receive the same pheromone seed several times coming from different paths. However, after a pheromone seed has been received the first time, then the node will only broadcast new arrivals of the same seed if the counter of the newly arrived seed is smaller than the counter of the previously received seed. This reduces the amount of information sent through the network and makes sure that Alice only receives the shortest route proposals.

---

<sup>2</sup>The value of  $c_0$  can be chosen between  $2^6 = 64$  and  $2^7 = 128$ . By doing this, we only dedicate 1 Byte to the counter, but still allow for 128 intermediary nodes between Alice and Bob, which is enough.

Each node subtracts its own fees from the remaining fees before forwarding the seed. A node only has knowledge of remaining fees, and does not know  $f_{max}$ . This makes it extremely difficult for intermediary nodes to deduce from the fees any information about the sender or receiver of the payment, thus preserving Alice's and Bob's anonymity.

Here is an algorithm describing the work of a node upon receiving a pheromone seed from node  $s$  in the form of a message  $(P, c, f, a, t)$ . We denote by  $g$  the fees of the node performing the task.  $\bar{P}$  denotes the conjugate of  $P$ , i.e  $\bar{P}(0) = P(1)$  and  $\bar{P}(1) = P(0)$ .

```

Look for  $P$  in memory;
if  $P$  is not in memory and  $f - g \geq 0$  then
  | Insert  $(P, c, s, f, a)$  in memory;
else if  $(P, c', s', f', a)$  is in memory then
  | if  $c' \leq c$  then
  |   | Exit program;
  | else if  $f - g \geq 0$  then
  |   | Replace  $(P, c', s', f', a)$  by  $(P, c, s, f, a)$  in memory.
Look for  $\bar{P}$  in memory.
if  $\bar{P}$  is not in memory then
  | Send  $(P, c + 1, f - g, a, t)$  to all neighbors with channel balance at least  $a$ , except  $s$ .
else
  | create_and_send_match( $P$ ).

```

**Algorithm 1:** treatment of pheromone seeds

The function *create\_and\_send\_match*() used above starts the match phase of the ant routing protocol. A matched seed has the form  $M(0) := 0 \frown P(0)$  or  $M(1) := 0 \frown P(1)$ , where  $P(0)$  and  $P(1)$  are pheromone seeds.  $M(0)$  is sent back to Alice, and  $M(1)$  is sent back to Bob. Assume  $(P(0), c, s, f, a)$  and  $(P(1), c', s', f', a)$  are both in memory, and denote by  $g$  the fees of the node where the match occurs. Then the function *create\_and\_send\_match*( $P$ ) does the following:

```

if  $f + f' - g \geq 0$  then
  | Set  $M(0) := 0 \frown P(0)$ ,  $M(1) := 0 \frown P(1)$ .
  | Generate a random number  $Id$ ;
  | Set  $F := f + f' - g$ ;
  | Set  $C := c + c' + 1$ ;
  | Store  $(Id, s')$  in the match tree;
  | Send  $(M(0), Id, c, C, F, t)$  to  $s$ ;
  | Send  $(M(1), Id, c', C, F, t)$  to  $s'$ ;

```

**Algorithm 2:** Match creation

Matched seeds are forwarded as messages of the form  $(M(\epsilon), Id, c, C, F, t)$ ,  $\epsilon \in \{0, 1\}$ . The counter  $c$  is decreased by 1 at each hop. When a node receives a matched seed  $(M(\epsilon), Id, c, C, F, t)$ , it looks for the corresponding  $P(\epsilon)$  in its memory and checks that the counter associated to  $P(\epsilon)$  is equal to  $c - 1$ . If this is the case, then the node retrieves the sender  $s$  of  $P(\epsilon)$  and forwards  $M(\epsilon)$  to  $s$ . If the counter associated to  $P(\epsilon)$  does not agree with  $c - 1$ , this means that, before receiving  $M$ , the node has received a new instance of  $P(\epsilon)$  with lower counter. But then the information associated to  $P(\epsilon)$  in the node's memory do not correspond to the



information associated to  $M(\epsilon)$ , and thus  $M(\epsilon)$  must be discarded. However, the node will eventually receive a new instance of  $M(\epsilon)$  with new information.

Here is the algorithm for the treatment of matched seeds. The input is a message of the form  $(M(\epsilon), Id, c, C, F, t)$  received from node  $y$ , and  $g$  denotes the fees of the node performing the task:

```

if  $\epsilon = 0$  then
  Look for  $P(0)$  in memory.
  if  $(P(0), c', s, f, a)$  is in memory then
    if  $c' \neq c - 1$  then
      | Exit program;
    if  $s = 0$  then
      | Store  $(Id, y, C, F)$  in the special match tree
    else
      | Store  $(Id, y)$  in the general match tree;
      | Send  $(M(\epsilon), Id, c - 1, C, F)$  to  $s$ .
if  $\epsilon = 1$  then
  Look for  $P(1)$  in memory.
  if  $(P(1), c', s, f, a)$  is in memory then
    if  $c' \neq c - 1$  then
      | Exit program;
    Store  $(Id, s)$  in memory.
    if  $s \neq 0$  then
      | Send  $(M(\epsilon), Id, c - 1, C, F, t)$  to  $s$ .

```

**Algorithm 3:** Treatment of matched seeds

After some time, Alice will have received several matched seeds, stored in her special match tree. She now chooses one (say, the one with the lowest fees), which is stored as  $(Id, y, C, F)$ . Note that Alice can recover the amount of fees she will have to pay by computing  $2f_{max} - F$ . This number is bounded by  $2f_{max}$ , but not by  $f_{max}$ . If Alice is not willing to pay up to  $2f_{max}$  in fees, then she can choose to replace  $f_{max}$  by  $\frac{f_{max}}{2}$  at the beginning of the pheromone phase.

To launch the confirmation phase, Alice generates a list of random numbers  $l_0$ , which will serve for the final counter check. Alice then sends the confirmation seed  $(Id, l := l_0, t)$  to  $y$ . This starts the confirmation phase of the ant routing protocol.

Here is the algorithm performed by a node after receiving the confirmation seed  $(Id, l, t)$ :

```

Look for the match identifier  $Id$  in the match tree;
if  $(Id, s)$  is in the match tree and  $s \neq 0$  then
  | Generate a random integer check.
  | Append check to  $l$ .
  | Store  $(Id, s, check)$  in the confirmation tree;
  | Send  $(Id, l, t)$  to  $s$ ;

```

**Algorithm 4:** Treatment of confirmed seeds

Remember that the “ $s$ ” in the matched seed  $(Id, s, t)$  refers to the sender of the pheromone seed from which  $Id$  was made, and that  $s = 0$  only for Bob and Alice. Therefore, the “ $s \neq 0$ ” clause in algorithm ?? states that the algorithm is executed by every intermediary node but not by Bob. When Bob receives  $(Id, l, t)$ , then Bob sends  $l$  to Alice. Alice then checks that the number of random numbers which were

appended to  $l$  is equal to  $C - 2c_0$  (note that the number of intermediary nodes between Alice and Bob is  $C - 2c_0$ ). If this is not the case, this means that one of the nodes on the route indicated by  $Id$  is a cheater. Alice then chooses another route: she chooses another “ $Id$ ” seed in her special match memory and starts a new confirmation phase with this new “ $Id$ ”.

If the number of random numbers appended to  $l$  matches  $C$ , then Alice removes  $l_0$  from  $l$ . She appends a new list of random numbers  $l_1$  at the end of  $l$  and sends  $(Id, l, t)$  to  $y$ . Then each node on the path checks that the first number in  $l$  is the one that they generated, i.e each node follows the following algorithm:

```

Look for  $Id$  in the confirmation tree;
if  $(Id, s, check)$  is in the confirmation tree and  $s \neq 0$  then
    | if  $check = l[0]$  then
    | | Remove  $check$  from  $l$ ;
    | | Send  $(Id, l, t)$  to  $s$ ;

```

**Algorithm 5:** counter check round

When Bob receives  $(Id, l, t)$ , he sends a message to Alice telling her that she can proceed with the payment. Note that the list  $l_1$  that has been appended to  $l$  at the start of the counter check round is here so that nodes don't know the length of the payment route. In particular, no node knows that Alice is paying Bob.

**Remark 3.2:** One way of reducing the workload on the network would be to use a variant of the ant routing algorithm which we just presented. In this variant, only Alice, and not Bob, sends pheromone seeds. Bob then waits until he receives the pheromone seed from Alice. Bob may receive several seeds, each indicating a different route. Bob can then communicate the information he has on each of these routes to Alice, who then chooses one of them. Bob then creates the match himself from the pheromone seed chosen by Alice, and sends the matched seed back to Alice through the chosen route.

Note that this variant is only viable if Alice trusts Bob. Indeed, since Bob knows the remaining fees for each seed he receives, he could cheat Alice by taking all the remaining fees for himself.

## 4 Robustness of ant routing

Payment routing depending on knowledge of the network threatens decentralization. If knowledge of the network is required for payment, some nodes may choose to delegate the task of gathering knowledge to other nodes. This can lead to the emergence of “hub” nodes which centralize most routing tasks, with most nodes at the periphery. In particular, the use of landmarks encourages this process. This type of network is easy to disrupt: if all payments depend on a small group of nodes, then one only needs to neutralize these nodes to incapacitate the whole network.

An important advantage of the ant routing algorithm is its total decentralization. All nodes in the network perform the exact same task and obey the exact same rules. To have the same protocol rules for all nodes is an important key idea of decentralization. Ant routing requires no knowledge of the geometry of the network, which prevents the emergence of “hub” nodes. This makes the network particularly resilient to attacks: even if an attacker neutralizes a larger part of the network, the remaining part can still function normally provided it remains connected.

It is worth noting that the main reason for the high failure rate of the current routing algorithm in the lightning network is the fact that nodes who compute payment routes only know the capacities of the channels on the network, but they do not know their balances. As a consequence, the paying node often ends up choosing a route which cannot relay the payment due to insufficient available funds in one direction [?]. Ant routing solves this problem by making sure that every channel on the path has sufficient funds in the desired direction. Indeed, during the pheromone phase, a node will forward the seed to a neighbor only if the funds available on their channel is sufficient. As a consequence, Alice knows that, for every match which she receives, the channels on the path indicated by the match have enough funds to forward her payment.

Note also that ant routing addresses some concerns about the privacy of lightning transactions which have been raised in [?]. The authors of [?] noticed that many LN transactions only have one intermediary node between payer and payee, which means that the intermediary node knows their identities. Ant routing solves this problem by offering several possible path to Alice. Thanks to the counter, Alice knows the length of each path indicated by a match, and she can choose a path of length at least 3 if she wants to preserve her privacy.

A malicious node may attempt to disturb the network by cheating on the counter. Indeed, a malicious node may add a negative counter to a pheromone seed to make the path more attractive. Setting very small fees also increases the chances of Alice choosing this path. The malicious node could then refuse to transmit the payment, thus disturbing the network. Note however that, thanks to the counter check phase which we added at the end of ant routing, Alice will discover the presence of a malicious node before sending the payment. Because the ant routing algorithm usually returns several matched seeds to Alice, she can then choose another route, avoiding the malicious node. If Alice sees that the confirmation phase, the counter check phase or the payment itself did not terminate successfully, she can send a message to the other nodes of the path to warn them of the presence of a malicious node. If a node notices frequent problems with attempted payments through one of its neighbors, it can decide to stop forwarding pheromone seeds through him or even to close the payment channel. In this way, the malicious node will eventually be de facto isolated and rejected from the network, and will not be able to disturb it any longer. The same will happen with a denial of service attack by nodes that respect all the rules of the protocol but in the last moment refuse to collaborate in establishing the payment. Therefore, each node has interest in keeping historical data and statistics on the performance of its neighbors to avoid being connected to dishonest nodes. In this sense, the network behaves more like an ant colony since the traces left by older payments will reinforce the best neighbors.

## 5 Life time of seeds

The seeds are only useful for the routing task for a transaction, and must therefore be kept only a short time in memory. This raises the question of the life time of seeds: how long should the seeds be kept in memory? They need to be kept long enough to allow the routing task to be completed, but as short as possible in order to reduce memory usage.

To determine the optimal seed life time, we need an estimate of how long it

will take to Alice to receive matched seeds. This depends on the bandwidth for data propagation on the Lightning Network. We can have an approximate estimate using the Bitcoin network as a proxy for the lightning network. Both networks are topologically of the same nature, more precisely, they are well-connected networks. Also seeds in ant routing do propagate faster through the network than bitcoin transaction on the Bitcoin network. This is due to the fact that pheromone seeds are much smaller than bitcoin transactions, but also because the local node process of transactions in Bitcoin nodes is much heavier than the process of pheromone seeds. It follows that the speed of transaction propagation on the Bitcoin network gives us a generous upper bound on the speed of pheromone propagation on the Lightning Network.

We found on [?] that, in 2014, it took around 0.8 seconds for a transaction to reach 50% of the nodes in the Bitcoin network. This is in accordance with [?] (See their fig. 16). It follows that we make the conservative assumption that in average a pheromone seed reaches 50% of nodes in 0.8 seconds. Assuming the initial independence of propagation of Alice and Bob seed, for a given node, the probability to have a match after 0.8 seconds is  $0.5^2 = 0.25$  and the probability that there is a match somewhere in the network after 0.8 seconds is  $1 - 0.75^N$  where  $N$  is the number of nodes. For  $N \geq 32$ , this probability is more than 99.99%. Now, if we assume that both Alice's pheromone and Bob's pheromone have reached 10% of the nodes in 0.5 seconds, then, the probability to have a match at this date is  $1 - 0.99^N$  which is greater than 99.99% when  $N \geq 1000$  which is already the case today. The exact distribution of the time it takes to have a first match depends on the network topology, but in view of these numbers it is reasonable to assume the expected unmatched time cycle of a pheromone seed to be about 0.5 seconds. Obviously, this may be higher when Alice or Bob do have a bad connectivity to the network, but this situation is neglectable if the average node is well connected as we assume. Note also that this estimate implies that the expected duration time for the whole routing task to be completed is less than 1 second.

## 6 A general maximum capacity estimate

We can develop the previous argument using the Bitcoin network as a proxy for estimating the maximum network capacity of an arbitrary network with similar characteristics. We first estimate in an elementary way the maximum network capacity of the Bitcoin network. We assume instant propagation of the data across the network and that the local nodes workload time for the propagation of a transaction is neglectable. The estimate remains valid if there is no lag caused by the local workload of nodes. We consider several random variables, for  $t \geq 0$ ,

- $\mathbf{N}(t)$  number of blocks mined by the network at time  $t$ ,  $\mathbf{N}(0) = 0$ .
- $\mathbf{L}(t)$  number of transactions created by the network at time  $t$ ,  $\mathbf{L}(0) = 0$ .
- $\mathbf{M}(t)$  mempool size in Bytes at time  $t$ ,  $\mathbf{M}(0) = M_0$ .
- $\mathbf{B}(b)$  size in Bytes of the block number  $b$ .
- $\mathbf{m}(x)$  size in Bytes of transaction  $x$ .
- $\mathbf{t}(x)$  time when transaction  $x$  was created.

Note that  $\mathbf{L}, \mathbf{M}, \mathbf{N}$  are random process, and  $\mathbf{L}, \mathbf{N}$  are Poisson processes [?]. The distribution of  $\mathbf{M}(t)$  depends on the network condition. At  $t = 0$  we start counting

blocks  $B_1, B_2, \dots$  and transactions created  $x_1, x_2, \dots$  from this moment and the initial condition  $M(0) = M_0$ . We have

$$M(t) = M_0 + \sum_{k=1}^{\mathbf{L}(t)} m(x_k) - \sum_{k=1}^{\mathbf{N}(t)} B(b_k) .$$

In the steady regime of maximum capacity,  $\mathbf{M}(t)$  is constant,  $\mathbf{M}(t) = M_0$ . Taking expected values in the precedent equation, we get, using Wald's Theorem,

$$0 = \mathbb{E}[\mathbf{L}(t)]\mathbb{E}[\mathbf{m}] - \mathbb{E}[\mathbf{N}(t)]\mathbb{E}[\mathbf{B}]$$

thus

$$\mathbb{E}[\mathbf{L}(t)] = \frac{\mathbb{E}[\mathbf{N}(t)]\mathbb{E}[\mathbf{B}]}{\mathbb{E}[\mathbf{m}]}$$

We have,  $\mathbb{E}[\mathbf{N}(t)] = t/\tau_0$  where  $\tau_0$  is the average inter-block time. At steady maximum capacity regime  $\mathbb{E}[\mathbf{B}]$  is constant and equal to the maximum block size,  $\mathbb{E}[\mathbf{B}] = B_{\max}$ . Let  $m_0 = \mathbb{E}[\mathbf{m}]$  the expected size of a transaction. These calculations show the following result:

**Theorem 6.1**

With the previous assumptions, at steady maximum capacity regime for typical transactions we have

$$l_0 = \frac{B_{\max}}{m_0\tau_0}$$

where  $l_0$  is the maximum number of transactions per second,  $B_{\max}$  is the maximum block size,  $m_0$  is the average size of a transaction, and  $\tau_0$  is the average inter-block time.

The arguments given to derive the above formula are quite general and apply to any network with similar characteristics as the Bitcoin network.

**Application (Bitcoin network):** We can apply the precedent formula to the Bitcoin network. We have  $\tau_0 = 10$  min (slightly smaller when difficulty is increasing). Also, before segwit improvement and neglecting the block header size of  $80B$ ,  $B_{\max} = 1$  MB. Typical Bitcoin transactions have a size of  $m_0 = 250$  B. We remind that one bitcoin transaction can contain several financial transactions and the size depends sensitively in the number of inputs and outputs. From this we get

$$l_{\text{BTC}} = 6.67 \text{ tx/sec}$$

for the maximum number of typical transactions per second in the Bitcoin network. The minimal size of a transaction is 61 B (or 63 B for a coinbase transaction, but this is neglectable, see [?]) then we get

$$l_{\text{BTC}} = 27.3 \text{ tx/sec}$$

**Application (Monero network):** Our formula applies directly. The maximum block size is set to  $M_0 = 1$  MB, the current average size of a transaction is  $m_0 = 550$  B and the average interblock time of  $\tau_0 = 120$  sec. Our formula gives

$$l_{\text{XMR}} = 15.15 \text{ tx/sec}$$

for the maximum number of transactions per second in the Monero network.

**Application (Ethereum network):** We apply the precedent ideas to the Ethereum network. The Ethereum protocol imposes no size limit on blocks or transactions, but the miners set a gas limit  $G_{\max}$  for the blocks, currently set at  $G_{\max} = 10^7$  units of gas, that is the total sum of gas for the transactions in the block. The gas of a transaction roughly measures the complexity of the computations in the smart contracts in Ethereum transactions. Indeed one line of code costs 1 gas unit, hence there is some correlation between the gas of a transaction and its size. But the proper metric to measure the network capacity is by using our formula where we replace the  $B_{\max}$  by  $G_{\max}$ , and  $m_0$  by the average gas  $g_0$  required by a transaction. The formula for the maximum number of transactions per second is

$$l_{ETH} = \frac{G_{\max}}{g_0 \tau_0}$$

We currently have  $\tau_0 = 15$  sec. Observe that the full capacity of the network can be reached by a maximum of basic transactions with minimal gas, or by a few transactions with maximal gas. For  $g_0 = g_{\min} = 21.000$  we have

$$l_{ETH} = 31.75 \text{ tx/sec}$$

For  $g_0 = g_{\max} = 200.000$  we have

$$l_{ETH} = 3.33 \text{ tx/sec}$$

Currently, the observed historical maximum block size, according to the block explorer Etherscan, is 36 630 B. Note that this represents 1.46 MB size per 10 min. The average Ethereum transactions size is about  $m_0 = 500$  B (not counting data volume necessary for the smart contract). with these figures, applying our formula we get

$$l_{ETH} = 4.88 \text{ tx/sec}$$

for the maximum number of transaction per seconds assuming a maximum size block of 36 630B. Obviously this size limit is not set by the protocol, but the practical limit has the same order of magnitude. Ethereum has similar scaling problems than Bitcoin, and this is reflected by the current heavy size of its full blockchain.

**Application (Ant Routing):** The same argument leading to the formula applies to the Lightning Network running the Ant Routing protocol. This time we denote by  $M_0$  is the maximum size of the local mempool at each node. Let  $\mu_0$  the average size of data propagated through all the network for each transactions, essentially unmatched and matched pheromone seeds so  $\mu_0 \simeq 100$  B. Let  $\eta$  the average life time of most data in the Lightning network, we can take the conservative figure of  $\eta = 2$  sec. We have,

$$l_{AR} = \frac{M_0}{\tau_0 \eta}$$

for the number of transactions per second of the Lightning Network running on Ant Routing. For the conservative figures choosen, we have

$$l_{AR} = \frac{M_0}{200}$$

For example, for a modest mempool of  $M_0 = 20$  MB we get

$$l_{AR} = 10.000 \text{ tx/sec}$$

As we will see in the next sections, the numerical simulations of the local process time of seeds at nodes allow for a flow of  $10^6$  tx/sec with no impact on the propagation and matching time, thus the network can process an order of  $10^6$  tx/sec.

## 7 AVL trees

In this section we describe the structure we choose in later sections to implement the ant routing protocol: AVL trees. AVL trees were introduced in [?]

We first fix some notations. If  $T$  is a binary tree, we denote by  $h(T)$  the height of  $T$ . If  $x$  is a node of  $T$ , we denote by  $T_x$  the subtree of  $T$  generated by  $x$ .  $T_x^l$  denotes the subtree generated by the left child of  $x$ , and  $T_x^r$  denotes the subtree of  $T$  generated by the right child of  $x$ . If  $x$  has no left child (respectively, right child), we say that its left child (respectively, right child) is NULL.

A **binary search tree** is a binary tree  $T$  such that:

1. Every node of  $T$  is labelled with an integer.
2. For every nodes  $x, y$  of  $T$  respectively labeled by the integers  $n$  and  $m$ , we have  $y \in T_x^l \Rightarrow m < n$  and  $y \in T_x^r \Rightarrow m > n$ .

If  $T$  is a binary search tree, we say that  $T$  is **balanced** if for every  $x \in T$ ,  $|h(T_x^l) - h(T_x^r)| \leq 1$ . An **AVL tree** is a balanced binary search tree.

The following is a look-up algorithm for a binary search tree. It takes an integer  $n$  as entry, and returns the node of  $T$  labeled by  $n$  if it exists, and returns NULL if no node of  $T$  is labeled by  $n$ .

```

input : Integer  $n$ , tree  $T$ 
output:  $lookup(n, T)$  := the node of  $T$  labeled by  $n$  if it exists, "NULL" otherwise
Set  $node := root(T)$ ;
while  $node \neq NULL$  do
  | if  $n < label(node)$  then
  |   |  $node := left\_child(node)$ ;
  | else if  $n > label(node)$  then
  |   |  $node := right\_child(node)$ ;
  | else
  |   | return  $node$ ;
end
Return  $node$ ;

```

**Algorithm 6:** Look-up algorithm for binary search tree.

Once the lookup is done, insertion in a binary search tree can be done fast. Here

is an algorithm which inserts  $n$  in  $T$  if  $n$  is not already in  $T$ :

```

input : Integer  $n$ , tree  $T$ 
Set  $node = root(T)$ ;
while  $true$  do
  if  $n < label(node)$  then
    if  $left\_child(node) = NULL$  then
       $left\_child(node) := create\_node()$ ;
       $label(left\_child(node)) := n$ ;
      Exit program;
    else
       $node := left\_child(node)$ ;
    else if  $n > label(node)$  then
      if  $right\_child(node) = NULL$  then
         $right\_child(node) := create\_node()$ ;
         $label(right\_child(node)) := n$ ;
        Exit program;
      else
         $node := right\_child(node)$ ;
    else
      Exit program;
  end
end

```

**Algorithm 7:** Insertion algorithm for binary search tree.


Now we discuss the run time of these algorithms. Let  $N$  be the number of nodes in  $T$ . If  $T$  is not assumed to be balanced, then in the worst case algorithms ?? and ?? could run in time  $O(N)$ . However, if  $T$  is an AVL tree, then algorithms ?? and ?? are guaranteed to run in time  $O(\log_2(N))$ .

Algorithm ?? is a naive insertion algorithm. This naive algorithm could be problematic, as it takes no account of the height of the newly generated tree and could thus result in an unbalanced tree if repeated enough times. This means we cannot use algorithm ?? for insertion in AVL trees. Fortunately, there is a workaround to this problem. More precisely, we can slightly change the insertion algorithm so that the tree remains balanced. We now describe the insertion algorithm for AVL trees. We first need to explain how an unbalanced tree can be rebalanced.

Re-balancing a tree is done via operations on the tree called rotations. A rotation at a node  $z$  is one of the two operations shown in Figure ??.




**Figure 2:** Rotations at node  $z$ .



Assume that  $T$  is balanced, and assumed that we now insert  $w$  in  $T$  following algorithm ???. Assume further that, after this insertion,  $T$  has become unbalanced. We can re-balance  $T$  by performing rotations on certain nodes of  $T$ . Starting from  $w$ , move up towards the root of  $T$ . We denote by  $z$  the first unbalanced node which we encounter on this path. Let  $y$  be the heigher child of  $z$  and  $x$  the heigher child of  $y$ . There are four possible configurations for the relative positions of  $x, y$  and  $z$ , as shown by Figure ??. To each configuration, we associate a sequence of rotations, shown in Figure ??. These rotations transform  $T_z$  into a balanced tree. It is easy to show that the new tree obtained from  $T_z$  by these rotations has the same height as  $T_z$  had before inserting  $w$ . Since  $T$  was balanced before inserting  $w$ , it follows that  $T$  is again balanced once the rotations have been performed.

**Figure 3:** Rebalancing the tree



Here is a modified version of algorithm ?? which gives the insertion algorithm for AVL trees:

```

input : Integer  $n$ , tree  $T$ 
Set  $node = root(T)$ ;
while true do
  if  $n < label(node)$  then
    if  $left\_child(node) = NULL$  then
       $left\_child(node) := create\_node()$ ;
       $label(left\_child(node)) := n$ ;
    else
       $node := left\_child(node)$ ;
    else if  $n > label(node)$  then
      if  $right\_child(node) = NULL$  then
         $right\_child(node) := create\_node()$ ;
         $label(right\_child(node)) := n$ ;
      else
         $node := right\_child(node)$ ;
    else
      Exit program;
  end
 $height(node) := 1 + \max(height(left\_child(node)), height(right\_child(node)))$ 
 $balance := height(left\_child(node)) - height(right\_child(node))$ 
while  $-1 \leq balance \leq 1$  do
  if  $node = root$  then
    return;
   $node := parent(node)$ ;
   $height(node) := 1 + \max(height(left\_child(node)), height(right\_child(node)))$ ;
   $balance = height(left\_child(node)) - height(right\_child(node))$ ;
end
if  $balance > 1$  then
  if  $n < label(left\_child(node))$  then
     $node = right\_rotate(node)$ ;
  else
     $left\_child(node) = left\_rotate(left\_child(node))$ ;
     $node = right\_rotate(node)$ ;
if  $balance < -1$  then
  if  $n > label(right\_child(node))$  then
     $node = left\_rotate(node)$ ;
  else
     $right\_child(node) = right\_rotate(right\_child(node))$ ;
     $node = left\_rotate(node)$ ;

```

**Algorithm 8:** AVL insertion

Now let us compute the running time of this algorithm. The algorithm starts with the naive insertion, which runs in time  $O(\log_2(N))$ . We then have to move back up the tree to update the heights of each node until we find the first unbalanced node. In the worst case, this is done in time  $O(\log_2(N))$ . Finally, the rotation operations are done in  $O(1)$  time. Therefore, this algorithm runs in time  $O(\log_2(N))$ . This shows that AVL trees allow for fast look-up and insertion; more precisely, these operations are performed in logarithmic time.

## 8 Implementation of ant routing

We saw earlier that the main task of a node is to manage a set of seeds. The node must frequently operate a look-up on its set of stored seeds and sometimes insert a new seed. It is therefore important to use a data structure allowing fast look-up and insertion. For this reason, we chose AVL trees as the base structure for storing seeds.

A node must store three types of seeds: pheromone, matched and confirmed. Each of these is stored in a separate tree.

We focus on how pheromone seeds are stored and managed. Let  $\eta$  be the life time of seeds (in seconds), and set  $k := \frac{\eta}{0.1}$ . Pheromone seeds are stored in a tree  $T$ . The root of  $T$  has  $k + 1$ -many children. Each of these children is the root of an AVL tree. We denote by  $T_0, T_1, \dots, T_k$  these subtrees of  $T$ . Each  $T_i$  corresponds to a time interval  $I_i := [t_0 + i * 0.1, t_0 + (i + 1) * 0.1[$ . A seed with timestamp  $t$  is stored in the tree  $T_i$  such that  $t \in I_i$ . If the node receives a seed with a timestamp outside of the time interval  $[t_0, t_0 + (k + 1) * 0.1[$ , then the seed is discarded. At time  $t_1 := t_0 + (k + 1) * 0.1$ , the node erases the tree  $T_0$ , which contains seeds older than  $\eta$ , and creates a new tree  $T_{k+1}$ . This new tree will contain the seeds with timestamp in  $[t_1, t_1 + 0.1[$ . At time  $t_1 + 0.1$ , the node will erase  $T_1$ , create  $T_{k+2}$ , and the process continues.

Each  $T_i$  is an AVL tree labeled by seeds. Each node of the tree contains the following fields:

- An integer *seed*. This is the random number  $S$  chosen by Alice and Bob. This also serves as the label of the node.
- An integer *amount*, indicating the amount of the transaction.
- A boolean *pheromone0*. The value is true if and only if  $P(0)$  has been received.
- A boolean *pheromone1*. The value is true if and only if  $P(1)$  has been received.
- An integer *fees0*, giving the remaining fees associated to  $P(0)$  if  $P(0)$  was received, and set to 0 otherwise.
- An integer *fees1*, giving the remaining fees associated to  $P(1)$  if  $P(1)$  was received, and set to 0 otherwise.
- An integer *counter0*. This is the counter associated to  $P(0)$ , set to 0 if  $P(0)$  was never received.
- An integer *counter1*. This is the counter associate to  $P(1)$ , set to 0 if  $P(1)$  was never received.
- An integer *sender0*. This is the node which sent  $P(0)$ , set to 0 if  $P(0)$  was never received.
- An integer *sender1*. This is the node which sent  $P(1)$ , set to 0 if  $P(1)$  was never received.
- A pointer *left\_child*.
- A pointer *right\_child*.

This describes the pheromone tree. The match tree is built in a similar way. The only difference between pheromone and match trees resides in the structure of the nodes. In the match tree, each tree  $T_i$  is labeled with a match identifier  $Id$ . Each node of  $T_i$  contains the following fields:

- An integer *Id* (match identifier).
- An integer *target*.
- A pointer *left\_child*.
- A pointer *right\_child*.

The confirmation tree is similar to the pheromone and the match tree, except for the structure of nodes. Each node of the confirmation tree contains the following fields:

- An integer *Id* (match identifier).
- An integer *target*.
- An integer *check*.
- A pointer *left\_child*.
- A pointer *right\_child*.

We now give an estimate of the memory space used by the seeds with this implementation. The following table gives the size in Bytes of a node of each tree:

	seed	counter	amount	fees	sender/target	children	check	total
pheromone	8	1	4	4	1	16	0	34
match	8	0	0	0	1	16	0	25
confirmation	8	0	0	0	1	16	8	33

Note also that biggest nodes have around 100 neighbors, which is why 1 Byte is enough to store the sender or target. Pointers typically have size 8 Bytes, which is why a tree node needs 16 Bytes to store the addresses of its children. We can check if the length of 8 Bytes for seeds is sufficient to avoid collisions. We can estimate the probability of collisions occurring by using an approximation formula of the birthday problem. If  $n$  is the number of seeds simultaneously present in the network and  $N$  the number of all possible seeds, then the probability  $p$  of having a collision at a given time is  $p \simeq \frac{n^2}{2N}$ . We have  $n = \lambda\eta$ , where  $\lambda$  is the rate of incoming transactions and  $\eta$  the life time of seeds. If we take  $\lambda = 10000$  and  $\eta = 2$ , then we get  $n = 20000$ . For seeds of length 8 Bytes we have  $N = 2^{64}$ . Therefore, the probability of having a collision at a given time is  $p \simeq \frac{(20000)^2}{2^{65}}$ . The probability of having a collision in 100 years is then  $1 - (1 - \frac{(20000)^2}{2^{65}})^{3600 \times 24 \times 31 \times 12 \times 100}$ , which is approximately 3%.

If one wishes to improve these odds, one can increase the length of seeds to 9 Bytes. In that case, the probability of having a collision in 100 years is  $1 - (1 - \frac{(20000)^2}{2^{82}})^{3600 \times 24 \times 31 \times 12 \times 100}$ , which is approximately  $2.10^{-5}$ .

Note that the simulations whose results we present in Section ?? were done with seeds of size 8 Bytes. However, increasing the length of seeds to 9 Bytes will have a negligible effect on the computation time needed for the treatment of seeds.

Now let us compute the total memory  $M$  taken in one node for the routing task. For one routing task, a node stores at most one pheromone seed and one confirmed seed, but potentially several matched seeds. Therefore, The memory taken in a node for one routing task is upper-bounded by  $34 + 25r + 33$  Bytes, where  $r$  is the number of matched received. Let  $\lambda$  be the rate of transactions, i.e. the average number of

transactions performed on the network at each second. Then the number of seeds kept in memory by a node is at most  $\lambda\eta$ , where  $\eta$  is the life time of seeds, so the maximal amount of memory taken is  $\lambda\eta(34 + 25r + 33)$ .

We argued above that we can expect to set  $\eta := 2$ . Setting  $\lambda = 10000$ , we get a  $M = 20000 \times (67 + 25r)$ . For smaller, less connected nodes,  $r$  should take small values. For  $r \leq 8$ , we get  $2\text{MB} \leq M \leq 4\text{MB}$ . For bigger nodes,  $r$  may take bigger values, in which case we get (for  $r > 8$ )  $M \simeq 500000r = 0.5r\text{MB}$ .

Finally, we want to estimate the necessary bandwidth for the propagation of pheromones. Remember that pheromone seeds are relayed as messages of the form  $(P, c, f, a, t)$  (see Section ??). The size in Bytes of one such message is given by the following table:

pheromone seed	counter	fees	amount	timestamp	total
8	1	4	4	1	16

With a transaction rate of  $\lambda = 10000$ , this means we need a bandwidth of  $16 \times 10000$  Bytes per second, i.e 160kB/s. The parameter amount above is the amount of the transaction. By comparison with the Bitcoin network, 16 Bytes is half the size of the hash of a Bitcoin transaction that is sent to nodes in an “inventory message” to announce a new transaction. See [?, ?] for details on bitcoin transactions propagation.

**Remark 8.1:** Note that a standard timestamp like the unix timestamp is 4 Bytes long. However, we propose here to use timestamps of length 1 Byte. Here is how we think timestamps should work in ant routing: each timestamp indicates a time interval of 0.1 seconds *modulo 20 seconds*. Since there are only 200 such time intervals within 20 seconds, 1 Byte is enough to express every possible timestamp. We argue that these small timestamps are enough to make ant routing function. Because seeds are stored in trees each corresponding to a time interval of 0.1 seconds, it is indeed enough to have timestamps with precision 0.1 second. Moreover, we argued in section ?? that the life time of seeds can be chosen to be 2 seconds. This means that every seed can be removed after staying more than 2 seconds in memory. It follows that counting timestamps modulo 20 seconds is enough to manage seeds.

## 9 Estimating scalability

We address now the following question: is the ant routing protocol compatible with the goal of solving bitcoin’s scalability problem? In other words, how many routing requests per second can the network process using ant routing? In view of the arguments given in section ??, we answer this question by computing the time taken by a node to process all incoming routing requests. This is given by formula (??) below, which depends on the rate  $\lambda$  of transactions. In section ??, we give the numerical values for the parameters appearing in formula (??), which we obtained experimentally. This allows us to give a numerical estimate of the maximal transaction rate which the network can sustain.

## 9.1 Formula

We now want to estimate the time taken by a node to perform the routing task. We showed in Section ?? that look-up and insertion in an AVL tree both run in logarithmic time. The operation of deleting a tree is linear in  $N$ . We respectively denote by  $\alpha$ ,  $\beta$  and  $\gamma$  the numbers such that, for an AVL tree  $T$  with  $N$ -many nodes, look-up in  $T$  runs in  $\alpha \log_2(N)$  seconds, Insertion in  $T$  runs in  $\beta \log_2(N)$  seconds, and deleting the whole tree takes  $\gamma N$  seconds.

Let us estimate the time taken by a node to process all pheromone seeds for one routing task. We denote by  $\lambda$  the rate of transactions, i.e. the number of transactions performed per second. In our implementation of ant routing given in Section ??, the tree containing pheromone seeds is split into  $k + 1$ -many subtrees. We can assume that the routing tasks arrive at a constant rate, so that each subtree has the same number of nodes. In that case, each subtree has  $\frac{\lambda}{10}$ -many nodes. The node will only have to perform one insertion of pheromone seed, but may have to perform several look-ups. Let  $p$  be the number of look-ups to be performed. Then the total time taken to process the pheromone seeds of one routing task is  $p\alpha \log_2(\frac{\lambda}{10}) + \beta \log_2(\frac{\lambda}{10})$ .

Now let us consider the match phase. Now the trees containing the matches are labelled by the identifier  $Id$  of the match. In general, a node may receive several matches produced by the same pheromone seed. Denote by  $m$  the average number of matches created by one pheromone seed and received by a node. The node will need to perform  $m$ -many insertions (one for each match received) on the match tree, but only one look-up (the look-up will be performed during the confirmation phase). The total time taken to process matched seeds is on average  $\alpha \log_2(m \frac{\lambda}{10}) + m\beta \log_2(m \frac{\lambda}{10})$ .

Finally, the confirmation seed only needs one insertion and one look-up (the look-up is actually performed during the counter check phase). Note that, for a given transaction, only very few nodes will actually have to process a confirmation seed. Let  $c$  denote the probability that a node receives a confirmation seed. The time to process confirmed seeds is then on average  $c(\alpha \log_2(c \frac{\lambda}{10}) + \beta \log_2(c \frac{\lambda}{10}))$ .

The total time to process one routing task is therefore on average:

$$T(\lambda) := (p\alpha + \beta) \log_2(\frac{\lambda}{10}) + (\alpha + m\beta) \log_2(m \frac{\lambda}{10}) + c(\alpha + \beta) \log_2(c \frac{\lambda}{10}) \quad (1)$$

Now let us estimate the total time that a node must dedicate to the routing task to process all the routing demands that arise in the span of one second. The node has to process on average  $\lambda$ -many routing tasks, so the time taken to process seeds is  $\lambda T(\lambda)$ . We must also take into account the time taken to clean old seeds. This is done by deleting a tree every 0.1 seconds. The time dedicated to cleaning the mempool is therefore  $\gamma \lambda(1 + m + c)$ . Thus, the total time dedicated to the ant routing algorithm is

$$\hat{T}(\lambda) := \lambda(T(\lambda) + \gamma(1 + m + c)) \quad (2)$$


Now we just need to estimate the value of the coefficients  $\alpha, \beta, \gamma, p, m, c$ .

## 9.2 Experimental result

We do determine  $\alpha$ ,  $\beta$  and  $\gamma$  experimentally. We have implemented the work of a node with a C program following Section ?. Our program treats seeds stored

in an AVL tree. We use a look-up and insertion algorithm in our program that is a recursive Implementation of algorithms ?? and ?. Our program first randomly generates a mempool of a given size  $N$ . The program then inserts a new randomly generated seed in the tree and records the time it takes to the program to perform this insertion.

We ran the program for different values of  $N$  in the range from 100 to 100000. For each  $N$ , we ran a Montecarlo simulation executing the program with 1000 trials. We also did the same for the tasks of look-up and cleaning of the mempool. The experimental results are shown in Figure ?? below. This allowed us to estimate the values of  $\alpha$ ,  $\beta$  and  $\gamma$  appearing in formula (??). Experimentally, we find  $\alpha = 0.7 \times 10^{-6}$ ,  $\beta = 1.1 \times 10^{-6}$  and  $\gamma = 8.2 \times 10^{-8}$ .



**Figure 4:** Insertion, look-up and cleaning time depending on the size of the mempool

Now we come back to formula (??). We can now determine the maximum number of routing tasks which a node can perform per second. In other words, the maximum  $\lambda$  such that  $\widehat{T}(\lambda) < 1$ . The value of this maximal  $\lambda$  depends on parameters  $p$ ,  $m$  and  $c$ . Note that these parameters essentially depend on the centrality and connectivity of the node.

For most nodes, we can assume that  $c$  is negligible, and that  $m = 1$  on average, so it



remains to determine  $p$ . It is established in [?] that the average number of channels per node is 7. This means that  $p$  is bounded by 8. In the worst case, we thus have  $T(\lambda) = (8\alpha + 2\beta) \log_2(\frac{\lambda}{10}) = 7.8 \log_2(\frac{\lambda}{10})10^{-6}$  and  $\widehat{T}(\lambda) = \lambda(7.8 \log_2(\frac{\lambda}{10})10^{-6} + 8.2 \times 10^{-8} \times 2)$ . We find  $\lambda_{\max} \simeq 12500$ .

**Remark 9.1:** • This result means that an average node can process up to 12500 tx per second. However, it is important to note that it is not necessary that all nodes process all incoming routing tasks. Indeed, since the network is well-connected, a routing task should still be completed successfully even if only 50 percent of the network is performing the task. This means that we can hope that the network could sustain up to 25000 tx per second.

- The above estimate is valid for nodes with low or average connectivity. However, for big nodes,  $p$  can be much higher, the value of  $c$  might be non-negligible, and  $m$  might take an average value higher than 1. It would take a deeper analysis of the network to estimate the capacity of these nodes to process all incoming routing tasks.

The C code used in this section for estimating the capacity of Ant Routing algorithm is available at : <https://github.com/gabrielLehericy/Ant-routing-simulation>

**Acknowledgement** The two first authors were partially supported by the FUI Moneytrack project joined between INRIA and Pôle Universitaire Léonard de Vinci.

## References

- [1] G.M. Adelson-Velski, E.M. Landis, *An algorithm for the organization of information*, Dokl. Akad. Nauk SSSR, 146:2, 263–266, 1962.
- [2] B. Awerbuch, *Reducing complexities of the distributed max-flow and breadth-first-search algorithms by means of network synchronization*. Networks, 1985.
- [3] F. Beres and I.A. Seres and A.A. Benczur *A Cryptoeconomic Traffic Analysis of Bitcoins Lightning Network*, preprint, Cryptoeconomic systems, MIT Press, 2020.
- [4] M. Dorigo and T. Stützle *Ant Colony Optimization*, Bradford Company, 2004.
- [5] M. Dorigo, V. Maniezzo and A. Colorni, *Ant system: optimization by a colony of cooperating agents* in IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), vol. 26, no. 1, pp. 29-41, 1996.
- [6] E. Giorgiadis, *How many transactions per second can bitcoin really handle ? Theoretically*, Cryptology ePrint Archive, Report 2019/416, 2019.
- [7] C. Grunspan and R. Perez-Marco, *Double spend races*, International Journal of Theoretical and Applied Finance, Vol. 21, 2018.
- [8] C. Grunspan and R. Perez-Marco, *Ant routing algorithm for the Lightning Network*, arxiv:1807.00151, 2018.
- [9] Lester R Ford and Delbert R Fulkerson. Maximal Flow Through a Network. Canadian Journal of Mathematics, 8(3), 1956.
- [10] G. Malavolta, P. Moreno-Sanchez, A. Kate, and M. Maffei. *Silentwhispers: Enforcing security and privacy in decentralized credit networks*. ISOC Network and Distributed System Security Symposium - NDSS, 2017.

- [11] G. Naumenko, G. Maxwell, P. Wuille, A. Fedorova, I. Beschastnikh, *Bandwidth-efficient transaction relay in Bitcoin*, CCS'19, Proceeding of the ACM SIGSAC Conference on Computer and Communication Security, 2019.
- [12] S. Bakshi, B. Bhattacharjee, S. Delgado-Segura, J. Litton, A. Miller, A. Pachulski, C. Pérez-Solà, TxProbe: Discovering Bitcoin's Network Topology Using Orphan Transactions, Financial Cryptography and Data Security, 2019.
- [13] Shivanand C Kohalli, *Coordinate Routing in the Lightning Network*, Master's Thesis, Delft University of Technology, 2019.
- [14] J. Poon and T. Dryja, *The bitcoin Lightning Network: scalable off-chain instant payment*, Online, <http://lightning.network/docs/> , 2016.
- [15] P. Prihodko, S. Zhigulin, M. Sahnó, A. Ostrovskiy, and O. Osuntokun *Flare: An approach to routing in Lightning Network*. Online, <https://bitfury.com> , 2016.
- [16] S. Roos, M. Beck, and T. Strufe. Anonymous addresses for efficient and resilient routing in f2f overlays. In Computer Communications, IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on, pages 1–9. IEEE, 2016.
- [17] S. Roos, P. Moreno-Sanchez, A. Kate, and I. Goldberg. *Settling payments fast and private: Efficient decentralized routing for path-based transactions*. ISOC Network and Distributed System Security Symposium - NDSS, 2018.
- [18] I.A. Seres, L. Gulyas, D.A. Nagy, P. Burcsi, *Topological analysis of Bitcoin's Lightning Network*, arxiv:1901.04972, 2019.
- [19] <http://bitcoinstats.com/network/propagation/>
- [20] V. Sivaraman, S. Bojja Venkatakrishnan, M. Alizadeh, G. Fanti, P. Viswanath, *Routing cryptocurrency with the spider network*, arXiv:1809.05088, 2018.
- [21] H. S. Wilf. Algorithms and complexity. AK Peters/CRC Press, 2002.
- [22] Imbalance measure and proactive channel rebalancing algorithm for the Lightning Network, M. Nowostawski, R. Pickhardt, arXiv:1912.09555, 2019.