



About decentralized swarms of asynchronous distributed cellular automata using inter-planetary file system's publish-subscribe experimental implementation

Vincent Manuceau, Article Info

► To cite this version:

Vincent Manuceau, Article Info. About decentralized swarms of asynchronous distributed cellular automata using inter-planetary file system's publish-subscribe experimental implementation. International Journal of Informatics and Communication Technology (IJ-ICT), 2022, 10.11591/ijict.v11i1.pp32-44 . hal-03593611v2

HAL Id: hal-03593611

<https://hal.science/hal-03593611v2>

Submitted on 7 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

About Decentralized Swarms of Asynchronous Distributed Cellular Automata Using Inter-Planetary File System's Publish-Subscribe Experimental Implementation.

Vincent Manuceau¹

¹Department of Research and Development, Makis Research, Remalard, Orne, France

Article Info

Article history:

Received Sep 07, 2021

Revised Dec 24, 2021

Accepted Jan 4, 2022

Keywords:

Cellular Automata

IPFS

Publish-Subscribe

Distributed Cellular

Automata

Grid'5000

ABSTRACT

This article describes the simple implementation of asynchronous distributed cellular automata and decentralized swarms of asynchronous distributed cellular automata built on top of Inter-Planetary File System's Publish-Subscribe (IPFS pubsub) experimentation. Various Publish-Subscribe models are described. As an illustration example, two distributed versions and a decentralized swarm version of a 2D elementary cellular automaton are thoroughly detailed to highlight the simplicity of implementation with IPFS and the inner workings of these kinds of cellular automata. Both algorithms were implemented, and experiments were conducted throughout five datacenters of Grid'5000 testbed in France to obtain preliminary performance results in terms of network bandwidth usage. This work is prior to implementing a large-scale decentralized epidemic propagation modelling and prediction system based upon asynchronous distributed cellular automata with application to the current epidemic of SARS-CoV-2 (COVID-19).

This is an open access article under the [CC BY-SA](#) license.



Corresponding Author:

Vincent Manuceau,

Department of Computer Science Research and Development,

Makis Research c/o Creole Shop (La Boutique Creole),

7 Les Docks Saint-Marc, ZA Saint-Marc Sud, 61110 Remalard-en-Perche, France.

Email: vincent@manuceau.net

1. INTRODUCTION

Decentralization and distribution of computing and communication systems represent an ever-increasing topic of interest, with a wide array of problems to solve and an even wider array of applications from free-speech protection [1] to climate change modeling [2]. This article describes a simple implementation of an asynchronous distributed cellular automata and a decentralized swarm of asynchronous distributed cellular automata built on top of Inter-Planetary File System's [3] Publish-Subscribe experimentation. After briefly describing notions of Cellular Automata, IPFS and Publish-Subscribe protocols, two distributed versions and a decentralized swarm version of a simple 2D cellular automaton are detailed in order to highlight the simplicity of implementation with IPFS and the inner workings of these kinds of cellular automata. The algorithms described in this paper are intentionally straightforward to focus on the simplicity of such asynchronous and decentralized pubsub systems. Implementation and preliminary results were obtained with experiments conducted across five datacenters throughout France with Grid'5000 testbed. This paper is motivated as preliminary work

to build a very large scale distributed computing application that models population dynamics and predicts demography at city/region/country level, with multi-level and inter-level interactions, in the scope of the current SARS-CoV-2 (COVID-19) epidemic.

2. DEFINITIONS

2.1. Cellular Automata

A Cellular Automaton is a dynamical system composed of a finite lattice of cells with local and straightforward communication capabilities. Each cell has a finite number of states evolving in discrete steps through time, and each cell state depends on the state of its neighbourhood. A transition function processes state changes. Originating from Ulam's works on dynamical systems[4], Cellular Automata directly contributed to Von Neumann's theory of self-replicating automata[5]. Cellular Automata and Distributed Cellular Automata represent an extensive field of study and notably interesting modelling and simulation tools for contemporary topics ranging from submicroscopic physics[6] to large scale macroscopic phenomena[7] such as climate change[2]. This paper will implement asynchronous distributed and decentralized swarms of cellular automata through Inter-Planetary File System's Publish-Subscribe capabilities.

2.2. Inter-Planetary File System

Inter-Planetary File System (IPFS) is a multipurpose, distributed, peer-to-peer, version-controlled file system with no single point of failure[3]. IPFS engenders a global Merkle Directed Acyclic Graph data structure (Merkle DAG)[8], with a content-addressing storage block architecture driven by Distributed Hash Tables (DHT), a block exchange system and a self-certifying namespace[3]. Stored content inside IPFS is accessible via CID hyperlinks (Content Identifiers). Currently IPFS is used for a wide variety of applications, such as distributed web applications and serverless applications [9], telecommunication, cloud data storage networks[10], content delivery networks (CDN) and blockchains [8], and even a crypto-currency based large scale decentralized file storage network called Filecoin[11][12]. Its trustless and decentralized structure may lead to the development of a censorship-resistant and permanent web[1].

2.3. PubSub: IPFS and Libp2p implementations

2.3.1. About PubSub

Publish-Subscribe models (PubSub) consists of asynchronous distributed and independent nodes, where each node can Publish events or Subscribe to related topics or contents through an overlay communication infrastructure [13]. Focusing on topic-based PubSub systems, when a node publishes an event to a related topic, all nodes that subscribed to this topic receive it asynchronously. Publishers and subscribers preserve their anonymity, as they do not interact directly and do not have to know each other to communicate. Specifically, the section below describes three topic-based PubSub implementations: IPFS FloodSub, Libp2p GossipSub and Libp2p EpiSub.

2.3.2. IPFS FloodSub

FloodSub, also known as DumbSub or PubSub-Flood, is the first and the most simple Publish-Subscribe implementation experiment in IPFS / Libp2p. FloodSub is based upon message routing by network flooding with no CastTree forming [14], and ambient peer discovery by the use of external Distributed Hash Tables (DHT) [15]. Concerning message broadcasting on small networks, FloodSub has low latency and is ideal for instant messaging applications. However, it cannot scale to more significant sized networks due to its high overhead and bandwidth consumption.

2.3.3. Libp2p GossipSub

GossipSub is a gossip-based Publication-Subscribe protocol[16] relying on a mesh construction and a score function [17]. The network structure uses two types of bidirectional peering: Full-message peerings, where peers send entire messages to sparsely connected peers called mesh members. The second type is Metadata-only peerings, where peers gossip about message availability and maintain full-message peerings[18]. In a GossipSub network, any peer can change their peering type from full-message to metadata-only (Pruning) and conversely (Grafting). Additionally, each node scores its peers based upon each peer behaviour, and then this scoring limits message transmission only

to peers reaching a certain score threshold[19]. Nodes can also publish to unsubscribed topics via Fan-out peering. In this unidirectional mechanism, they send their message to 3 randomly picked peers that subscribed to the topic, called fan-out peers, and then redistribute the message to the network[18]. Thus by design, GossipSub is highly scalable, reliable, fast and efficient, resilient and attack-resistant[17].

2.3.4. Libp2p EpiSub

EpiSub is a proximity aware mono-source multicast optimized Publish-Subscribe protocol[20] implementing Epidemic Broadcast Trees Plumtree protocol[21] (Gossip-based spanning tree construction, tree repair and optimization). HyParView membership protocol[22] manages EpiSub peers. Each node has two different views in this protocol: a small active one that engenders a message-passing overlay. A larger passive one maintains active view network resilience in case of node failure. EpiSub implements GoCast proximity-aware overlay scheme[23] which dynamically maintains near neighbourhood connectivity, i.e. low latency neighbours, by a near degree node confinement method during add and drop connection phases. Libp2p EpiSub comprises two main protocols: the broadcast protocol (publish) that uses lazy multicast tree construction through Plumtree's epidemic broadcast[21]. The second protocol is membership management (subscribe), which maintains active and passive peers lists for a related topic. Eager and lazy active peers are distinguished, eager ones actively disseminating new messages at the edges of the multicast tree, and lazy ones only gossiping about message summaries and maintaining the multicast tree[20]. Multicast tree and peer proximity optimization[23] constantly optimize transmission latency and propagation latency. An implementation of EpiSub is currently under active development in the scope of this project. An upcoming paper will describe its performance results.

3. PROPOSED ALGORITHMS

3.1. Simple Distributed cellular automata

3.1.1. Used Cellular Automata and Conventions

The cellular automata used will be Conway's Game of Life: two states, grade IV[24], totalistic, Moore's Neighborhood driven[25], elementary cellular automaton. This CA will run on a 2-dimensional circular grid. Each cell runs independently and asynchronously, as subscription events trigger its transition function: its own cell for the neighbour publishing version and its eight neighbouring cells for the neighbour subscribing version. A generic Publish-Subscribe mechanism will be used below in the form of two functions `pubsub.pub` as publish function and `pubsub.sub` as subscribe function. Publish function takes for argument the name of the topic and the message to publish. Subscribe function takes for argument the name of the topic to subscribe and a callback function that processes received messages. These two generic functions can be adapted to match FloodSub, GossipSub and EpiSub versions. Each published message reflects the cell state or a start action request. Message can be 0 (dead cell), 1 (alive cell) or 2 (start request). Any cell of the CA can receive a start request, then the triggered cell broadcasts its state to its neighbours and effectively starts the CA. Let $N \geq 3$, the CA will run on a $N \times N$ circular grid, and let C a cell of coordinates $(x, y) \in N \times N$. Thus the `pubsub` topic of C will be "cell- x - y ". For performance comparison purposes, two Distributed cellular automata (DCA) versions, a Neighbour publishing version (NP) and a Neighbour subscribing version (NS), will be detailed below and then implemented.

3.1.2. Neighbour publishing version

In this version, on a $N \times N$ CA, $N \geq 3$, each cell has one topic subscription, thus N^2 total subscriptions, and at each round, a cell publishes to its eight neighbours as shown in figure 1a, thus $8 \times N^2$ publications per round. Algorithm 1 shows a single cell class, instantiating a cell object member of the CA. The initialization parameter `coord` is in the form $(x, y) \in N \times N$, the state parameter is either 0 or 1 and the length parameter is N . The `neighb` variable is an array of the cell neighbour coordinates, thus an array of length 8 as Moore's neighbourhood [25] is used. Variable `alive_neighb` is the number of alive neighbours for this current round, and `current_neighb` is the number of messages received by its neighbourhood. Variable `subscribe` keeps the `pubsub` mechanism, processes each event received by updating the cell and broadcasting each cell state update to its neighbourhood.

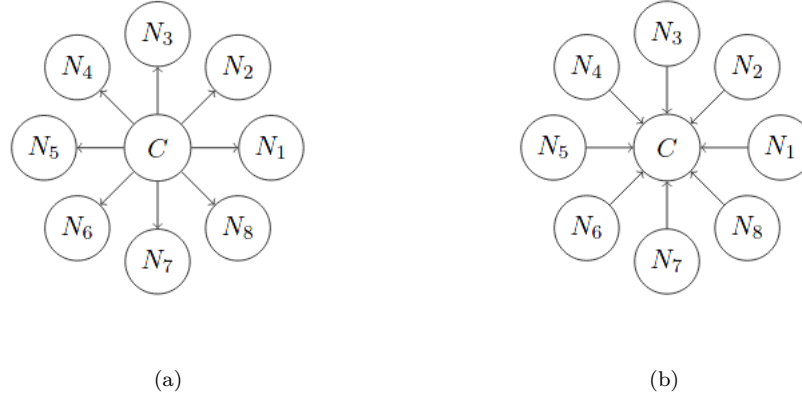


Figure 1. Cellular Automata publishing/subscribing schemes,
(a) Cell C publishing to neighbours N_i , (b) Cell C subscribing to neighbours N_i

Algorithm 1 Cell Object : Neighbour publishing version

```

class Cell_NP(coord,state,length)
  coord ← coord
  state ← state
  neighb ← neighb_list(coord,length)
  alive_neighb ← 0
  current_neighb ← 0
  subscribe ← cell_subscribe(this,neighb_publish)
end class

```

3.1.3. Neighbour subscribing version

In this version, on a $N \times N$ CA, $N \geq 3$, each cell has 8 topics subscriptions as shown in figure 1b thus $8 \times N^2$ total subscriptions, and at each round a cell publishes to its own topic, thus N^2 publications per round. Algorithm 2 shows a single cell class, instantiating a cell object member of the CA. The initialization parameter *coord* is in the form $(x, y) \in N \times N$, the state parameter is either 0 or 1 and the length parameter is N . Variable *alive_neighb* is the number of alive neighbours for this current round, and *current_neighb* is the number of messages received by its neighbourhood. Variable *subs* keeps the pubsub mechanism as an array of neighbourhood subscriptions that processes each event received, updates the cell and broadcasts the cell state update to its current topic.

3.1.4. Publish-Subscribe functions

Publish-Subscribe functions of both CA versions used above can be described as follows : On a Neighbour publishing scheme, each cell uses the publish function to send a message to its neighbours, and the subscribe function listens to its topic. In Algorithm 3, these are *neighb_publish* and *cell_subscribe* functions. On a Neighbour subscribing scheme, each cell uses the publish function to update its topic and the subscribe function is used to listen to its neighbours' topics. In Algorithm 3, these are *cell_publish* and *neighb_subscribe* functions. It is essential to mention that the cell function parameter is just a reference to the cell object and not the cell object itself.

3.1.5. Message processing function

The message processing function plays the role of an asynchronous CA transition function and is used for both versions. As shown in Algorithm 4, this function takes for argument a cell reference and a pubsub publication callback, and then returns a function that processes a message, starts or updates a cell given its current neighbourhood state according to Conway's Game of Life rules: any alive cell surrounded by two alive neighbours lives on the next state, any cell surrounded by three

Algorithm 2 Cell Object : Neighbour subscribing version

```

class Cell_NS(coord,state,length)
  coord ← coord
  state ← state
  alive_neighb ← 0
  current_neighb ← 0
  subs ← neighb_subscribe(this,length,cell_publish)
end class

```

Algorithm 3 Publish and Subscribe functions

```

procedure cell_publish(cell)
  pubsub.pub(cell_name(cell.coord),cell.state)
end procedure
procedure neighb_publish(cell)
  for i ∈ cell.neighb do
    pubsub.pub(cell_name(cell.neighb[i]),cell.state)
  end for
end procedure
function cell_subscribe(cell,publish)
  sub ← pubsub.sub(cell_name(cell.coord))
  sub.on("message",process(cell,publish))
  return sub
end function
function neighb_subscribe(cell, length, publish)
  neighb ← neighb_list(cell.coord,length)
  sub ← Array(8)
  for i ∈ neighb do
    name ← cell_name(neighb[i].coord)
    cur_neighb ← pubsub.sub(name)
    cur_neighb.on("message",process(cell,publish))
    sub[i] ← cur_neighb
  end for
  return sub
end function

```

alive neighbours lives on the next state, and in any other case, the cell dies on the next state. The publication callback is called either when a start request is made ($msg = 2$), and in that case, it starts its neighbourhood by broadcasting its current state, or the cell has reached a new round and has updated its state, and in that case, it publishes its new state.

Algorithm 4 Message processing function

```

function process(cell,publish_callback)
  return
  function message_processor(msg)
    if  $msg = 2$  then                                     ▷ CA start request
      return  $publish\_callback(cell)$ 
    else
       $alive\_neighb \leftarrow alive\_neighb + msg$ 
       $current\_neighb \leftarrow current\_neighb + 1$ 
      if  $current\_neighb = 8$  then
        if ( $alive\_neighb = 2$  and  $cell.state = 1$ ) or  $alive\_neighb = 3$  then
           $cell.state \leftarrow 1$ 
        else
           $cell.state \leftarrow 0$ 
        end if
        ( $alive\_neighb, current\_neighb$ )  $\leftarrow (0, 0)$ 
        return  $publish\_callback(cell)$ 
      end if
    end if
  end function
end function

```

3.1.6. Miscellaneous functions

Two functions remain necessary to complete the CA, they are used in both versions and are described in Algorithm 5. Function `cell_name` returns the topic name of a given cell, and `neighb_list` function returns an array of neighbours coordinates for a given cell of coordinates `coord` and a given `length = N` on a $N \times N$, $N \geq 3$, circular grid.

Algorithm 5 Miscellaneous functions

```

function cell_name(coord)
  return " $cell -$ " .  $coord[0]$  . " - " .  $coord[1]$ 
end function
function neighb_list(coord, len)
   $neighb \leftarrow Array(0)$ 
  ( $x, y$ )  $\leftarrow coord$ 
   $\Delta \leftarrow [-1, 0, 1]$ 
  for ( $\alpha, \beta$ )  $\in \Delta \times \Delta$  do
    if  $\Delta[\alpha] \neq 0$  or  $\Delta[\beta] \neq 0$  then
       $\mu \leftarrow (len + x + \Delta[\alpha]) \bmod len$ 
       $\nu \leftarrow (len + y + \Delta[\beta]) \bmod len$ 
       $neighb.push((\mu, \nu))$ 
    end if
  end for
  return  $neighb$ 
end function

```

3.1.7. CA initialization and known limitations

The cellular automata can be initialized by instantiating the $N \times N$ cells and then send a message $\text{msg}=2$ to any topic "cell- x - y " where $(x, y) \in N \times N$. The triggered cell then propagates its current state and effectively starts the whole CA in an asynchronous domino effect. This straightforward implementation has known limitations, as a cell does not keep track of its current round and thus does not know if received messages are related to its current or next round. The following CA implementation will tackle this problem.

3.2. Decentralized swarms of distributed CA

3.2.1. Description and conventions

For large scale purposes, the asynchronous distributed CA will be divided into swarms of autonomous smaller CA with inner-swarm and inter-swarm communication capabilities. Communication is necessary inside the swarm and between the cells located at the edges of two distinct swarms. This will be achieved through local Cell-to-Cell and global Swarm-to-Swarm communication channels as shown in figure 2. Each cell has to keep track of its current round, and it will process its neighbourhood messages according to its round number. At round 0, each cell is unaware of its neighbours' swarm. Each cell identifies itself to its neighbourhood during this discovery phase through direct messaging as in the NP DCA version previously described. From round 1 begins the processing phase, where each cell is aware of the corresponding swarms of its neighbourhood, and communication is provided through inner-swarm and inter-swarm broadcast.

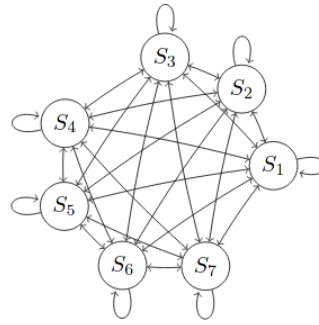


Figure 2. Swarms S_i of distributed CA communicating

3.2.2. Cell and Swarm objects

Algorithm 6 describes a single cell class, instantiating a cell object member of a swarm. The initialization parameter coord is in the form $(x, y) \in N \times N$, state parameter is either 0 or 1, swarm parameter is a reference to the related swarm of the cell and length parameter is N . Variable neighb is an array of 8 neighbouring cells where each item has its swarm id, cell coordinates, and its two last state rounds stored. Variable subscribe keeps the pubsub mechanism at a local level, processing each event received, updating the cell and broadcasting cell state updates to its neighbourhood. Algorithm 7 describes a single swarm class, instantiating a swarm object member of the CA. The initialization parameter swarm_id represents the id of the swarm, and cell_array parameter assigned to cells variable is the array of cell objects from the swarm. Variable subscribe keeps the pubsub mechanism at both local and global levels, processing and dispatching each event received depending on its origin and destination.

3.2.3. Publish and Subscribe functions

Publish-Subscribe functions enable local and inter-swarm communication. On a local level, each cell uses the publish function to send a message to its swarm, and the subscribe function listens and processes messages from its topic. In Algorithm 8, these are the functions cell_publish and cell_subscribe . On a global level, each swarm uses the publish function to dispatch a received message, either to its cells or to another swarm, and the subscribe function listens and processes messages from its topic. In Algorithm 8, these are the functions swarm_publish and swarm_subscribe . Functions

Algorithm 6 Cell Object : Swarm version

```

class Cell(coord,state,swarm,length)
  coord ← coord
  state ← state
  swarm_id ← swarm.id
  round ← 0
  neighb ← neighb_list(coord,length)
  subscribe ← cell_subscribe(this,swarm)
end class

```

Algorithm 7 Swarm Object

```

class Swarm(swarm_id,cell_array)
  id ← swarm_id
  cells ← cell_array
  subscribe ← swarm_subscribe(this)
end class

```

pubsub.sub and pubsub.sub are generic enough to implement either FloodSub, GossipSub or EpiSub Publish-Subscribe models.

Algorithm 8 Publish and Subscribe functions

```

function swarm_subscribe(swarm)
  sub ← pubsub.sub(swarm_name(swarm.id))
  sub.on("message",swarm_process(swarm))
  return sub
end function
procedure swarm_publish(swarm,msg)
  if swarm.id = msg[0] then
    (swarm_process(swarm))(msg)
  else
    pubsub.pub(swarm_name(msg[0]),msg)
  end if
end procedure
function cell_subscribe(cell,swarm)
  sub ← pubsub.sub(cell_name(cell.coord))
  sub.on("message",cell_process(cell,swarm))
  return sub
end function
procedure cell_publish(msg)
  pubsub.pub(cell_name(msg[1]),msg)
end procedure

```

3.2.4. Cell message broadcast

The cell message broadcast procedure described in Algorithm 9 sends a message to all neighbours of a given cell through swarm_publish procedure.

3.2.5. Swarm processing function

The swarm processing function plays the role of an asynchronous message router. As shown in Algorithm 10, this function takes for argument the current swarm reference. It returns a function that processes each message as follows: if the message is heading to an inner cell, the target cell processes it directly. If the message is heading to a known swarm, the pubsub swarm_publish function will

Algorithm 9 Cell message broadcast

```

procedure cell_msg_broadcast(cell,swarm)
  for  $i \in cell.neighb$  do
     $cur\_nehb \leftarrow cell.neighb[i]$ 
     $msg \leftarrow (cur\_nehb.swarm\_id, cur\_nehb.coord, cell.coord, cell.round, cell.state)$ 
     $swarm\_publish(swarm, msg)$ 
  end for
end procedure

```

send it. In the initialization phase, the swarm destination is unknown. Thus the message is directly dispatched to the corresponding cell via pubsub `cell_publish` function.

Algorithm 10 Swarm processing function

```

function swarm_process(swarm)
  return
function message_router(msg)
   $coord \leftarrow msg[1]$ 
   $cell\_id \leftarrow find\_cell\_id(swarm.cells, coord)$ 
  if  $cell\_id \geq 0$  then
    return  $(cell\_process(swarm.cells[cell\_id], swarm))(msg)$ 
    ▷ Incoming message
  else if  $msg[0] \geq 0$  then
    return  $swarm\_publish(swarm, msg)$ 
    ▷ Known swarm publish
  else
    return  $cell\_publish(msg)$ 
    ▷ Unknown swarm publish
  end if
end function
end function

```

3.2.6. Cell processing function

The cell processing function plays the role of an asynchronous transition function. As shown in Algorithm 11, this function takes for argument a cell reference and a pubsub publication callback and returns a function that processes a message, starts or updates a cell depending on the message received. If the received swarm id is -1, the cell identifies itself on the network, broadcasting its swarm id, coordinates and current state to its neighbourhood. In other cases, the received message is processed in the neighbour array. Suppose the cell received enough neighbours for its round (processed via `alive_n` function). In that case, Conway's Game of Life rules are applied, the neighbour state array is reset for this round (via `reset_neighb` procedure), and the cell message broadcast procedure sends the next state to its neighbours.

4. PRELIMINARY RESULTS - PERFORMANCE COMPARISONS

4.1. Implementation and experimentation testbed

DCA and DSDCA algorithms were implemented in the NodeJS framework and interfaced with a tweaked version of IPFS, compatible with Linux x64 and OS X platforms. Source code is freely available on GitHub[26]. The following performance measures were carried out using Grid'5000 testbed[27], a large-scale testbed for experiment-driven research supported by a scientific interest group (GIS) hosted by Inria, including CNRS, RENATER and several French Universities as well as other organizations. Only IPFS bandwidth was measured during preliminary experiments, but inconsistencies between IPFS and network bandwidth, including TCP/IP overhead, were observed. It was then decided that total network bandwidth usage, including TCP/IP overhead, was the most relevant resource to monitor.

Algorithm 11 Cell processing function

```

function cell_process(cell,swarm)
  return
  function message_processor(msg)
    (swarm_id,coord,target,round,state)  $\leftarrow$  msg
    if swarm_id = -1 then
      cell_msg_broadcast(cell, swarm)
      return true
    end if
    neighb_id  $\leftarrow$  coord_to_id(cell.neighb,target)
    cell.neighb[neighb_id].swarm_id  $\leftarrow$  swarm_id
    cell.neighb[neighb_id].s[round mod 2]  $\leftarrow$  state
    (alive,total)  $\leftarrow$  alive_n(cell.neighb,cell.round)
    if total = 8 then
      if (alive = 2 and cell.state = 1) or alive = 3 then
        cell.state  $\leftarrow$  1
      else
        cell.state  $\leftarrow$  0
      end if
      cell.round  $\leftarrow$  cell.round + 1
      reset_neighb(cell)
      cell_msg_broadcast(cell, swarm)
      return true
    end if
    return false
  end function
end function

```

▷ Start the CA

▷ All neighbours received

4.2. DCA NP/NS Grid'5000 Experimentation Protocol

Experimentations with DCA NP and NS versions were carried out on various clusters from 5 different sites (Lyon, Lille, Nancy, Nantes, Grenoble). Each experiment was launched for 100 rounds on five machines in parallel, and network bandwidth was measured with vnstat. Various settings were used, such as the number of nodes (from 10 to 226), FloodSub and GossipSub router protocols, and NP / NS DCA versions. At the end of each experiment, the total network bandwidth used for IPFS nodes bootstrapping and CA processing was measured, and the average network bandwidth per node was processed. Each DCA ran on individual machines, and network bandwidth was measured through vnstat software in interactive mode, monitoring each machine's loopback (lo).

4.3. Distributed Cellular Automata Performance

DCA performance was measured during 250 experiments with several nodes ranging from 10 to 226. The focus is kept on results for experiments with more than 50 nodes. As shown in figures 3a and 3b, there is no meaningful difference in network bandwidth between both PubSub protocols when comparing NP or NS implementations, 0.1% for NP version and 0.3% for NS version on average.

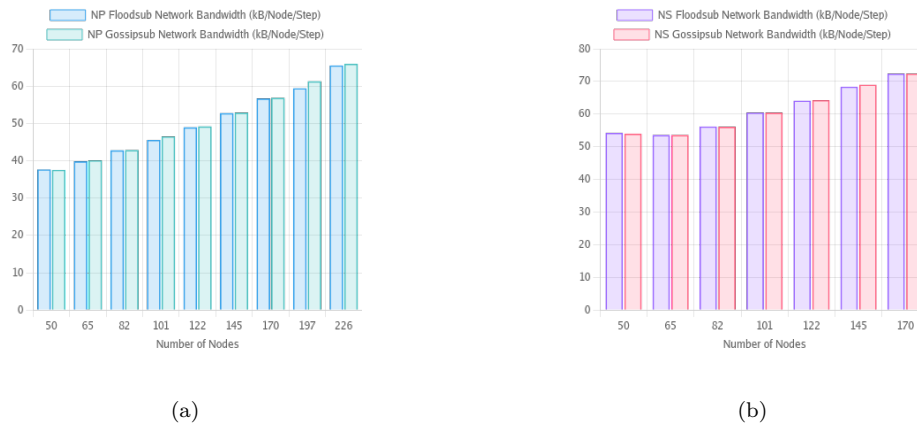


Figure 3. DCA: Floodsub vs Gossipsub Network Bandwidth,
(a) NP version in kilobit per node per step, (b) NS version in kilobit per node per step

Contrarily, it is much relevant to compare NP and NS implementations running with the same PubSub protocol, as shown in figures 4a and 4b.

On average, there is a 24.7% bandwidth difference for Floodsub protocol and a 24.3% bandwidth difference for GossipSub protocol. In any case, and contrarily to what was supposed theoretically, the NP version seems to be the most bandwidth-efficient algorithm, even independently from FloodSub or GossipSub protocol.

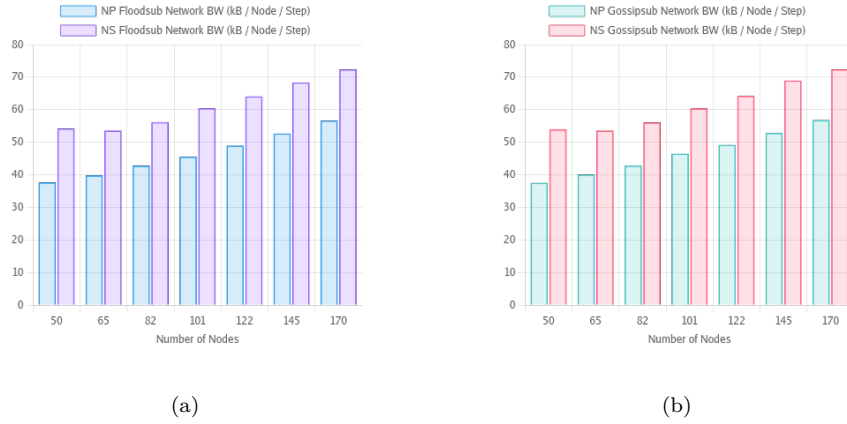


Figure 4. DCA: Floodsub and Gossipsub Network Bandwidth,
(a) NP Floodsub vs NS Floodsub in kb/node/step, (b) NP Gossipsub vs NS Gossipsub in kb/node/step

4.4. DSDCA Grid'5000 Experimentation Protocol

Experimentations for Decentralized Swarms of DCA were carried out across five sites with machines from various clusters with 10GBps connectivity (Nova-Lyon, Chiclet-Lille, Gros-Nancy, Ecotype-Nantes, Dahu-Grenoble). Each Swarm was running on a specific site, and network bandwidth was measured with vnstat software. Nova was used as the IPFS private network and Swarm bootstrapping node, Chiclet, Gros, Ecotype and Dahu as Swarms of DCA nodes. Various settings were tried, such as the number of swarms (3 and 4), nodes (from 13 to 404), and FloodSub and GossipSub router protocols. Each experiment was run five times for 1000 rounds. At the end of each experiment, the total network bandwidth used between sites was measured, and the average network bandwidth per node was processed.

4.5. Decentralized Swarms of DCA Performance

During 250 different experiments, DSDCA implementation performance was measured with several nodes ranging from 13 to 404. Network bandwidth measures for the Bootstrap phase (IPFS Private network and Node bootstrapping) was dissociated from the DSDCA Processing phase. Concerning the Bootstrap phase in three swarms and four swarms experiments, there is a substantial network bandwidth difference between FloodSub and GossipSub versions, as shown in figures 5a and 5b. On average, there is a 31.3% bandwidth difference for FloodSub vs GossipSub protocol on three swarms bootstrap phases, and 11.7% bandwidth difference for FloodSub vs GossipSub protocol on four swarms bootstrap phases.




Figure 5. DSDCA: Floodsub vs Gossipsub Bootstrap Bandwidth, (a) 3 Swarms DSDCA in Mb/node, (b) 4 Swarms DSDCA in Mb/node

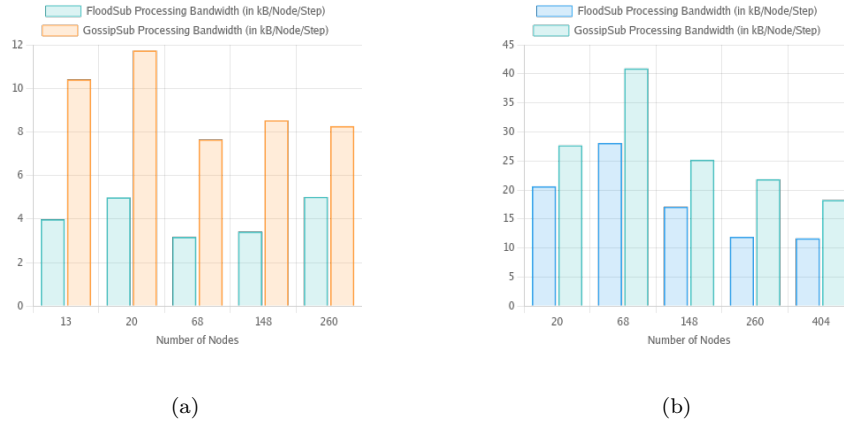


Figure 6. DSDCA: Floodsub vs Gossipsub Processing Bandwidth, (a) 3 Swarms DSDCA in kb/node/step, (b) 4 Swarms DSDCA in kb/node/step

Regarding the DSDCA processing phase in three swarms and four swarms experiments, network bandwidth variations between FloodSub and GossipSub versions are pretty significant and informative, as shown in figures 6a and 6b. On average, there is a 55,45% bandwidth difference for FloodSub vs GossipSub protocol on three swarms bootstrap phases, and 34.16% bandwidth difference for FloodSub vs GossipSub protocol on four swarms bootstrap phases. In any case, as a preliminary result, and contrarily to what was initially supposed, FloodSub is observed to be the most bandwidth-efficient protocol for the DSDCA algorithm, independently from the number of nodes and of swarms.

4.6. A few statistics about Grid'5000 usage

This experiment used 4,701 CPU core hours on 47 machines across Lille, Grenoble, Lyon, Nancy and Nantes datacenters. The total network bandwidth used across Grid'5000 datacenters is estimated to be 137.475 TB for NP/NS DCA versions experiments and 148.396 TB for Decentralized Swarms of DCA experiments. About 500 different experiments were run during two weeks, about 43,000 IPFS nodes were fired up, and 30,000,000 rounds of CA were computed.

5. CONCLUSION AND FURTHER WORKS

This article introduced an experimental implementation of simple asynchronous distributed cellular automata and decentralized swarms of asynchronous distributed cellular automata driven by IPFS and Libp2p PubSub. Publish-Subscribe functions genericity provides the possibility to switch between FloodSub, GossipSub and EpiSub protocols. Experiments were carried out on Grid'5000 testbed for either NP/NS DCA and 3/4 swarms DSDCA. Preliminary results showed the sustainability of NP version over NS DCA and FloodSub protocol over GossipSub for DCA and DSDCA algorithms. Further experiments will be conducted with a larger number of Swarms and Nodes, and a comparison of protocol performance will be made with the EpiSub protocol. EpiSub implementation over Libp2p/IPFS is currently under active development, as an automated DSDCA experiment launcher and bandwidth monitoring over Grid'5000, and a real-time asynchronous distributed cellular automata visualization system. This work is prior to implementing a large-scale decentralized epidemic propagation modelling and prediction system based upon asynchronous distributed cellular automata applied to the current SARS-CoV-2 (COVID-19) epidemic. This system will use a hybrid adapted version of SIR models driven by DSDCA and fed by worldwide data, estimating population dynamics on various factors such as mobility, infection rates, vaccine rates, recoveries and deaths. In this DSDCA system, individual nodes will simulate population dynamics in cities, swarms for regions and upper-level swarms for countries. Although cellular automata can be computed faster in a centralized manner, decentralization will enable to free from memory and processing power limits induced by centralization in the context of such a very large scale simulation application.

ACKNOWLEDGEMENT

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>). I want to express my gratitude and thanks to Pierre Neyron from CNRS, which granted me Grid'5000 open access to run DCA and DSDCA experiments across 5 French datacenters. I want to thank especially George Polyzos from AUEB, Nuno Santos from INESC-ID and Jorge Soares from Protocol Labs, who reviewed this paper at a very early stage and whose feedbacks motivated me to elaborate and dig further. I would also like to thank the French Government for awarding social grants, providing me with the financial means to complete this self-funded, non-profit and independent research.

REFERENCES

- [1] J. Santos, N. Santos and D. Dias, "Censorship-resistant web annotations based on ethereum and IPFS," In Proceedings of the 35th Annual ACM Symposium on Applied Computing (SAC '20). Association for Computing Machinery, New York, NY, USA, doi: 10.1145/3341105.3374049, pp. 2211–2213, 2020.
- [2] A. J. Collados-Lara, E. Pardo-Igúzquiza and D. Pulido-Velazquez, "A distributed cellular automata model to simulate potential future impacts of climate change on snow cover area," *Advances in Water Resources*, vol.124, ISSN 0309-1708, pp. 106–119, 2019.
- [3] J. Benet, "IPFS - Content Addressed, Versioned, P2P File System," *arXiv Networking and Internet Architecture*, arXiv:1407.3561, 2014.
- [4] S. Ulam, "On some Mathematical Properties Connected with Patterns of Growth of Figures," *Proceedings of Symposia on Applied Mathematics*, vol.14, pp. 215–224, 1962.
- [5] J. Neumann and A. Burks, "Theory Of Self Reproducing Automata," University of Illinois Press, 1st ed., ISBN 0252727339, 1966.
- [6] V. Christianto, V. Krasnoholovets, and F. Smarandache, "Cellular Automata Representation of Submicroscopic Physics," *Prespacetime Journal, QuantumDream*, vol. 10, no. 8, pp. 1024–1036, December 2019.
- [7] P.M.A. Sloot, J.A. Kaandorp, A.G. Hoekstra and B. Overeinder, "Distributed cellular automata: Large scale simulation of natural phenomena," *Solutions to Parallel and Distributed Computing Problems: Lessons from Biological Sciences*, Computer Centre University of Tromsø, January 2001.

- [8] H. Huang, J. Lin, B. Zheng, Z. Zheng and J. Bian, "When Blockchain Meets Distributed File Systems: An Overview, Challenges, and Open Issues," IEEE Access, vol. 8, doi: 10.1109/ACCESS.2020.2979881, pp. 50574–50586, 2020.
- [9] D. Dias, J. Benet, "Distributed web applications with IPFS," 16th International Conference on Web Engineering (ICWE), Lugano, Switzerland, 2016.
- [10] S. S. Hasan, N. H. Sultan and F. A. Barbhuiya, "Cloud Data Provenance using IPFS and Blockchain Technology," In Proceedings of the Seventh International Workshop on Security in Cloud Computing (SCC '19). Association for Computing Machinery, New York, NY, USA, doi: 10.1145/3327962.3331457, pp. 5–12, 2019.
- [11] CryptoLab, "Filecoin: A decentralized storage network," Protocol Labs, July 2017.[Online]. Available: <https://research.protocol.ai/publications/filecoin-a-decentralized-storage-network>. [Accessed: 12-Sep-2020].
- [12] S. He, Y. Lu, Q. Tang, G. Wang, C. Qishi Wu, "Peer-to-Peer Content Delivery via Blockchain," arXiv Cryptography and Security, arXiv:2102.04685, 2016.
- [13] J. P. de Araujo, "A communication-efficient causal broadcast publish/subscribe system," PhD Thesis, Sorbonne Université, pp. 19-21, 2019.
- [14] V. Santos, "Js-libp2p-floodsub", Libp2p, December 2020. [Online]. Available: <https://github.com/libp2p/js-libp2p-floodsub>. [Accessed: 14-Feb-2021].
- [15] D. Dias, "PubSub at Scale - Open Problems," Protocol Labs, September 2020. [Online]. Available: https://github.com/protocol/ResNetLab/blob/master/OPEN_PROBLEMS/PUBSUB_AT_SCALE.md. [Accessed: 20-Feb-2021].
- [16] R. Baldoni, R. Beraldi, V. Quema, L. Querzoni and S. Tucci-Piergiovanni, "TERA: topic-based event routing for peer-to-peer architectures," DEBS '07 : Proceedings of the First ACM International Conference on Distributed Event-Based Systems, doi:10.1145/1266894.1266898, pp. 2–13, January 2007.
- [17] D. Vyzovitis, Y. Nopora, D. McCormick, D. Dias and Y. Psaras, "GossipSub: Attack-Resilient Message Propagation in the Filecoin and ETH2.0 Networks," In Proceedings of Protocol Labs TechRep (PL-TechRep-2020-002), Protocol Labs, arXiv:2007.02754, 2020.
- [18] "Publish / Subscribe - Libp2p Concepts," Libp2p Documentation [Online]. Available: <https://docs.libp2p.io/concepts/publish-subscribe>. [Accessed: 19-Feb-2021].
- [19] D. Vyzovitis, "Gossipsub v1.1: Security extensions to improve on attack resilience and bootstrapping," Libp2p, December 2020. [Online]. Available: <https://github.com/libp2p/specs/blob/master/pubsub/gossipsub/gossipsub-v1.1.md>. [Accessed: 2-Feb-2021].
- [20] D. Vyzovitis, "Episub: Proximity Aware Epidemic PubSub for libp2p," Libp2p, June 2019. [Online]. Available: <https://github.com/libp2p/specs/blob/master/pubsub/gossipsub/episub.md>. [Accessed: 14-Feb-2021].
- [21] J. Leita, J. Pereira and L. Rodrigues, "Epidemic Broadcast Trees," 2007 26th IEEE International Symposium on Reliable Distributed Systems (SRDS 2007), Beijing, China, doi: 10.1109/SRDS.2007.27, pp. 301–310, 2007.
- [22] J. Leita, J. Pereira and L. Rodrigues, "HyParView: A Membership Protocol for Reliable Gossip-Based Broadcast," 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07), Edinburgh, UK, doi: 10.1109/DSN.2007.56, pp. 419-429, 2007.
- [23] C. Tang, R. N. Chang and C. Ward, "GoCast: gossip-enhanced overlay multicast for fast and dependable group communication," 2005 International Conference on Dependable Systems and Networks (DSN'05), Yokohama, Japan, doi: 10.1109/DSN.2005.52, pp. 140-149, 2005.
- [24] S. Wolfram, "A New Kind of Science," Wolfram Media, ISBN 9781579550080, pp. 231–249, 2002.
- [25] T. Toffoli and N. Margolus, "Cellular Automata Machines: A New Environment for Modeling," MIT Press, Cambridge, MA, USA, ISBN 0262200600, ch. 7, 1987.
- [26] V. Manuceau, "About Decentralized Swarms of Asynchronous Distributed Cellular Automata using IPFS," GitHub, July 2021.[Online]. Available: https://github.com/vincent-manuceau/Decentralized_Cellular_Automata_over_IPFS. [Accessed: 30-Aug-2021].
- [27] F. Desprez and L. Nussbaum, "The data-centers facet of SILECS (A.K.A. Grid'5000)" SILECS/Datacenters - Grid'5000, April 2019.[Online]. Available:

<https://www.grid5000.fr/mediawiki/images/Grid5000.pdf>. [Accessed: 30-Aug-2021].

BIOGRAPHY OF AUTHOR



Vincent Manuceau is an independent Researcher with main interests in Cellular Automata, IPFS (Inter-Planetary File System) and RINA (Recursive Internetwork Architecture).

Currently Head of Creole Shop company, the leading Caribbean food and beverages on-line grocery store with 24h/48h shipping worldwide (<https://www.creole-shop.fr/en>), he develops innovative technologies with direct applications to the company, such as full automation of the logistics chain.

He does fundamental research on his spare time, and writes Computer Science and Mathematics courses, he also is reviewer for International Journal of Informatics and Communication Technology (IAES IJ-ICT) since 2021. Vincent is always open to knowledge sharing and teaching.

Further info on his homepage: <http://vincent.manuceau.net>