Information Processing Letters 115 (2015) 738–743

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

Gossip Latin square and the meet-all gossipers problem

Nethanel Gelernter *, Amir Herzberg

Department of Computer Science, Bar Ilan University, Ramat Gan, 52900, Israel

A R T I C L E I N F O

Article history:
Received 2 August 2013
Received in revised form 29 May 2014
Accepted 3 March 2015
Available online 5 March 2015
Communicated by Gregory Chockler

Keywords:
Latin square
Gossip
Combinatorial problems

A B S T R A C T

Given a network of $n = 2^k$ gossipers, we want to schedule a cyclic calendar of meetings between all of them, such that: **(1)** each gossiper meets (gossips) only once a day, with one other gossiper, **(2)** in every $(n - 1)$ consecutive days, each gossiper meets all other gossipers, and **(3)** every gossip, initiated by any gossiper, will reach all gossipers within $k = \log(n)$ days.

In this paper we study the above stated *meet-all gossipers problem*, by defining and constructing the *Gossip Latin Square* (GLS), a combinatorial structure which solves the problem. We then present an efficient construction of GLS, based on maximal Fibonacci LFSR.

1. Introduction

The gossip problem [3,4] is one of the most studied problems in distributed computing. The original problem discusses how can n *gossipers*, each knowing some *rumor*, spread all rumors between them, where in each *round* (step), every gossiper can exchange rumors with only one other gossiper. We refer to such an exchange of rumors between two gossipers as a *meeting*. Note that such "meeting" only refers to communication and does not require physical proximity.

The problem has a lot of versions, which differ by topologies, model (offline/online, sync/async, with or without considering faults) or by restrictions on the communication (e.g., instead of pairwise communication as in phone calls, group communication as in conference calls [6]). There are also many solutions, both randomized and deterministic.

The two main efficiency measures for solutions to the gossiping problem are: (1) message complexity; namely, how many meetings are needed for spreading all the gossips? (2) Time/round complexity; how many rounds are needed?

In the abstract above, we defined the *meet-all gossipers problem* which corresponds to the second question, as the minimal time to broadcast one gossip between $n = 2^k$ gossipers is k [2]. The difference between the gossipers problem and the *meet-all gossipers problem*, is in the second requirement: every gossiper must meet all the other $n - 1$ gossipers in every $n - 1$ consecutive days. In the rest of this paper we present a deterministic solution for the problem. We focus on the offline-scheduling, synchronous and fault-free model.

Beyond the theoretical interest, the *meet-all gossipers problem* has properties that may come handy in the design of communication protocols related to guaranteed delivery (QoS) and anonymity. Relying on other gossipers, we ensure the minimal delay, but even without them (on load, or when some of the gossipers are not entirely reliable) the maximal delay is $n - 1$, as every gossiper meets all the others every consecutive $n - 1$ days. This might be relevant also for communication mechanisms where there are messages that can be sent through other parties and secret messages that cannot. The deterministic routing that involves different paths and intermediate relays, can be used

* Corresponding author.
E-mail address: nethanel.gelernter@gmail.com (N. Gelernter).

http://dx.doi.org/10.1016/j.ipl.2015.03.001
0020-0190/

as a constant rate communication protocol with minimal delay for anonymous communication goals [8].

In the second section we formalize the *meet-all gossipers problem* and define *gossip Latin square (GLS)*, a combinatorial structure that solves the problem. In the third section we present a construction for GLS, and prove its correctness. In the last section we conclude and briefly discuss generalizations of the problem.

2. Definitions

2.1. The meet-all gossipers problem

To formally define the *meet-all gossipers problem*, we first define a *meeting schedule function*, which defines which pair of gossipers meet at any given day.

Definition 1. Consider a finite set H (of gossipers). A *meeting schedule* is a function $Meet : \mathbb{N} \times H \to H$. For every $d \in \mathbb{N}$, we say that h *meets* $Meet_d(h)$ in day d.

We next formalize requirements (1) and (2) in the abstract. A meeting schedule function *Meet* that satisfies these requirements, is said to be *perfectly fair*.

Definition 2. Meeting schedule function $Meet : \mathbb{N} \times H \to H$ is *perfectly fair* if it satisfies the following two requirements:

Pairwise daily meetings For every $d \in \mathbb{N}$ and $h \in H$ holds (1) $Meet_d(h) \neq h$ and (2) $h = Meet_d(Meet_d(h))$.

Meet all For every $d \in \mathbb{N}$ and $h \in H$, $\{Meet_{d+i}(h)\}_{0 \leq i < n-1} = H \backslash \{h\}$.

To define the third requirement, we have to define the set of *recipients*, denoted $R^{Meet}_{d,m}(h)$, of a rumor initiated by $h \in H$ on day d and propagated m days. We first define the *recipient* relation between a pair of gossipers, h and h'. Informally, gossiper h' is a (d, m)-*recipient* from h, if there is a sequence of meetings between gossipers, beginning from day d, such that a rumor initiated on day d by h, will reach h' within m days, via the meetings sequence.

Definition 3. Consider set of n gossipers H, two gossipers $h, h' \in H$, and $d, m \in \mathbb{N}$. We say that h' is a (d, m)-*recipient* from h according to the meeting schedule $Meet : \mathbb{N} \times H \to H$, *if and only if* there is a sequence of $y < m$ pairs $\{(p_t, q_t)| p_t \in H,\ 0 \leq q_t < q_{t+1} < m\}_{t=0}^{y}$ such that

1. $p_0 = Meet_{d+q_0}(h)$.
2. For every $0 < t \leq y$, $p_t = Meet_{d+q_t}(p_{t-1})$.
3. $h' = p_y$.

Let $R^{Meet}_{d,m}(h)$ denote the set of all gossipers h', s.t. h' is a (d, m)-*recipient* from h according to *Meet*. We call $R^{Meet}_{d,m}(h)$ the (d, m)-*recipients set* of h according to *Meet*.

Informally, $R^{Meet}_{d,m}(h)$ is the set of gossipers, that will receive a gossip, initiated by h on day d within m days, according to the meeting schedule function *Meet*.

We now formally define a *meet-all gossipers schedule*, i.e., a schedule that satisfies all three requirements in the abstract; such a meeting schedule is said to be *perfectly fair* and *round-optimal*.

Definition 4. Meeting schedule function $Meet : \mathbb{N} \times H \to H$ ($|H| = 2^k$) is a *meet-all gossipers schedule*, if it is perfectly fair (Definition 2) and also satisfies the following (*round-optimal*) requirement: For every $d \in \mathbb{N}$ and $h \in H$, $R^{Meet}_{d,k}(h) = H \backslash \{h\}$.

If a meeting schedule function is *meet-all gossipers schedule*, then for broadcasting gossips, it is enough that every *gossip* stays in the network during k days; this can be done by attaching to each gossip a creation timestamp. During these days, every gossiper that knows the gossip, shares it with every other gossiper he meets. Namely, during a meeting between two gossipers, they exchange all the gossips they know, that were originated less than k days before the meeting.

2.2. Schedule matrix

We next present a *schedule matrix*, a square matrix whose bottom row is the set of gossipers H, and whose content defines the meetings schedule.

Definition 5. Let M be an $n \times n$ matrix. We say that M is a *schedule matrix* if

1. The last row (index $n-1$), called also the *headline row*, contains n distinct elements denoted $H = \{h_i\}_{i=0}^{n-1}$ (gossipers).
2. For $0 \leq i < n$, $0 \leq j < n$, $M_{i,j}$ appears in the headline row.

We now define a mapping from a given schedule matrix M, to a meeting schedule function.

Definition 6. Let M be an $n \times n$ schedule matrix, and let $H = \{h_i\}_{i=0}^{n-1}$ be its headline row. The *meeting schedule function* of M is denoted $Meet^M : \mathbb{N} \times H \to H$, and defined as: $Meet^M_d(h_j) = M_{d \bmod (n-1), j}$.

We denote the (d, m)-recipients set of $Meet^M$ by $R^M_{d,m}$.

2.3. Gossip Latin Square (GLS)

A *Latin square (LS)* [7] is a $n \times n$ matrix whose rows and columns are permutations over n distinct elements. We now show that every Latin square is also a *schedule matrix*.

Lemma 7. *Every $n \times n$ Latin square matrix, is also a schedule matrix.*

Proof. The *headline row* contains n distinct elements, and every other row is a permutation over the elements of the *headline row*. $\square$

We now define GLS, a Latin square that maps to meeting schedule function that solves the *meet-all gossipers problem*.