R.A.I.R.O. Theoretical Informatics, to appear

THE COMPLEXITY OF SYSTOLIC DISSEMINATION OF INFORMATION IN INTERCONNECTION NETWORKS *

Juraj Hromkovič[†], Ralf Klasing, Walter Unger, Hubert Wagener[‡]
Department of Mathematics and Computer Science
University of Paderborn
33095 Paderborn, Germany

Dana Pardubská[§]
Faculty of Mathematics and Physics
Comenius University
84215 Bratislava, Slovakia

Keywords: communication algorithms, parallel computations.

Abstract

A concept of systolic dissemination of information in interconnection networks is presented, and the complexity of systolic gossip and broadcast in one-way (telegraph) and two-way (telephone) communication mode is investigated. The following main results are established:

(i) a general relation between systolic broadcast and systolic gossip,

(ii) optimal systolic gossip algorithms on paths in both communication modes, and

(iii) optimal systolic gossip algorithms for complete k-ary trees in both communication modes.

*This work was partially supported by grants Mo 285/4-1, Mo 285/9-1 and Me 872/6-1 (Leibniz Award) of the German Research Association (DFG), and by the ESPRIT Basic Research Action No. 7141 (ALCOM II).

[†]This author was partially supported by SAV Grant No. 88 and by EC Cooperative Action IC 1000 Algorithms for Future Technologies (ALTEC).

[‡]This author was supported by the Ministerium für Wissenschaft und Forschung des Landes Nordrhein-Westfalen.

[§]This author was partially supported by MSV SR and by EC Cooperative Action IC 1000 Algorithms for Future Technologies (ALTEC). Most of the work of this author was done when she was visiting the University of Paderborn.

1 Introduction

One of the most intensively investigated areas of computation theory is the study and comparison of the computational power of distinct interconnection networks as candidates for the use as parallel architectures for existing parallel computers. There are several approaches enabling to compare the efficiency and the "suitability" of different parallel architectures from distinct point of views. One extensively used approach deals with the possibility to simulate one network by another without any essential increase of computational complexity (parallel time, number of processors). Such an effective simulation of a network A by a network B surely exists if the network A can be embedded into B (more details and an overview about this research direction can be found in [MS90]).

Another approach to measure the computational power of interconnection networks is to investigate which class of computing problems can be computed by a given class of networks. Obviously, this question is reasonable only by additional restrictions on the networks because each class of networks of unbounded number of processors (like paths, grids, complete binary trees, hypercubes, etc.) can recognize all recursive sets. These additional restrictions mostly restrict the time of computations (for example, to $\log_2 n$ by complete binary trees or to real time by paths) and/or the kind of computation assuring a regular flow of data in the given network. A nice concept for the study of the power of networks from this point of view has been introduced by Culik II et al. [CGS84], and investigated in [IK84, CC84, IPK85, CSW84, CGS83, IKM85]. This concept considers classes of languages recognized only by systolic computations on the given parallel architecture (network) in the shortest possible time for a given network. The notion "systolic computation" has been introduced by Kung [Ku79], and it means that the computation consists only of the repetition of simple computation and communication steps in a periodic way. The reason to prefer systolic computations is based on the fact that each processor of a network executing a systolic algorithm works very regularly repeating only a short sequence of simple instructions during the whole computation. Thus, the hardware and/or software realization of systolic algorithms is essentially cheaper than the realization of parallel algorithms containing many irregularities in the data flow or in the behaviour of the processors.

The last of the approaches mentioned here helping to search for the best (most effective) structures of interconnection networks is the study of the complexity of information dissemination in networks (for an overview see [HHL88, HKMP93]). This approach is based on the observation that the realization of the communication (data flow between the processes) of several parallel algorithms on networks requires at least as much (or sometimes even more) time as the computation time of the processors. This means that the time spent with communication is an important parameter of the quality of interconnection networks. To get a comparison of networks from the communication point of view, the complexity of the realization of some basic communication tasks like broadcast (one processor wants to tell something to all other processors), accumulation (one processor wants to get some pieces of information from all other processors) or gossip (each processor wants to tell something to each other) is investigated for different networks.

The aim of this paper is to combine the ideas of the last two approaches mentioned above to get a concept of systolic communication algorithms enabling to study the communicational effectivity of networks when a very regular behaviour of each processor of the network is required.

The first step in this direction was made by Liestman and Richards [LR93b] who introduced a very regular form of communication based on graph coloring. (This kind of communication has later been called "periodic gossiping" by Labahn et al. [LHL93].) This "periodic" communication was introduced in order to solve a special gossip problem introduced by Liestman and Richards [LR93a] and called "perpetual gossiping", where each processor may get a new piece of information at any time from the outside of the network and the never halting communication algorithm has to broadcast it to all other processors as soon as possible.

The concept of Liestman and Richards [LR93b] includes some restrictions which bound the possibility of systolic communication in a non-necessary way (more about this in the next section). Another drawback is that the complexity considered in [LR93b, LHL93] is the number of systolic periods and not the number of communication rounds, i.e. only rough approximations on the precise number of rounds executed are achieved in [LR93b, LHL93]. Here, we introduce a more general concept of systolic (periodic) dissemination of information in order to evaluate the quality of interconnection networks from this point of view. Our main aim of the investigation of systolic communication is not only to establish the complexity of systolic realisation of basic communication tasks in distinct networks, but also to learn how much must be paid for the change from arbitrary "irregular" communication to nice, regular systolic one.

This paper is organized as follows. Section 2 contains the formal description of the concept of systolic communication and some basic observations comparing general communication algorithms with systolic ones and systolic gossip with systolic broadcast. Section 3 is devoted to the complexity of systolic gossip in paths. Optimal systolic gossip algorithms are presented for the two-way mode of communication. The complexity for the one-way mode is determined up to a small constant independent of the length of the path and the length of the period. Section 4 is devoted to systolic gossip in trees. Optimal systolic gossip algorithms for complete k-ary trees are described in the one-way and the two-way communication mode. In the Conclusion, the results achieved are discussed and some open problems are formulated. More information about the contents of Sections 3 and 4 and exact statements of the achieved results of this paper are given in Section 2 after the introduction of the concept of systolic communication.

2 The Concept of Systolic Communication

The aim of this section is to define the concept of systolic communication algorithms and to give some fundamental observations about systolic broadcast and gossip. Before doing this, we give a more precise description of the communication problems investigated and of the communication modes considered.

An (interconnection) network is viewed as a connected undirected graph $G = (V, E)$ where the nodes of V correspond to the processors and the edges correspond to the communication links of the network. An infinite sequence $\{G_i\}_{i=1}^{\infty}$ with $G_i = (V_i, E_i)$, $|V_i| > |V_j|$ for $i > j$, is called a class of interconnection networks. Examples of classes of interconnection networks are paths - $\{P_n\}_{n=1}^{\infty}$ (P_n is the path of n nodes), cycles - $\{C_n\}_{n=1}^{\infty}$ (C_n is the cycle of n nodes), complete balanced k-ary trees - $\{T_k^h\}_{h=1}^{\infty}$ (T_k^h is the complete balanced k-ary tree of depth h), cube-connected cycles - $\{CCC_k\}_{k=1}^{\infty}$ (CCC_k is the cube-connected cycles network of dimension k), and two-dimensional square grids - $\{Gr_m^2\}_{m=1}^{\infty}$ (Gr_m^2 is the $m \times m$ grid).

In what follows, we shall investigate the following three fundamental communication problems in networks:

1. Broadcast problem for a network G and a node v of G.
   Let $G = (V, E)$ be a network and let $v \in V$ be a node of G. Let v know a piece of information $I(v)$ which is unknown to all nodes in $V - \{v\}$. The problem is to find a communication strategy such that all nodes in G learn the piece of information $I(v)$.

2. Accumulation problem for a network G and a node v of G.
   Let $G = (V, E)$ be a network, and let $v \in V$ be a node of G. Let each node $u \in V$ know a piece of information $I(u)$ which is independent of all other pieces of information distributed in other nodes (i.e. $I(u)$ cannot be derived from $\bigcup_{v \in V - \{u\}}\{I(v)\}$). The set $I(G) = \{I(w) | w \in V\}$ is called the cumulative message of G. The problem is to find a communication strategy such that the node v learns the cumulative message of G.

3. Gossip problem for a network G.
   Let $G = (V, E)$ be a network, and let, for all $v \in V$, $I(v)$ be a piece of information residing in v. The problem is to find a communication strategy such that each node from V learns $I(G)$.

Now, it remains to explain what the notion "communication strategy" means. The communication strategy is meant to be a communication algorithm from an allowed set of synchronized communication algorithms. Each communication algorithm is a sequence of simple communication steps called communication rounds (or simply rounds). To specify the set of allowed communication algorithms one defines a so-called communication mode which precisely defines what may happen in one communication step (round). Here, we consider the following two basic communication modes:

a, one-way mode (also called telegraph mode)
   In this mode, in a single round, each node may be active only via one of its adjacent edges either as a sender or as a receiver. This means that if one edge (u, v) is active as a communication link, then the information is flowing only in one direction. Formally, let $G = (V, E)$ be a network, $\vec{E} = \{(v \to u), (u \to v) | (u, v) \in E\}$. A one-way communication algorithm for G is a sequence of rounds $A_1, A_2, \ldots, A_k$, where $A_i \subseteq \vec{E}$ for every $i \in \{1, \ldots, k\}$, and if $(x_1 \to y_1), (x_2 \to y_2) \in A_i$ and $(x_1, y_1) \neq (x_2, y_2)$ for some $i \in \{1, \ldots, k\}$, then $x_1 \neq x_2 \wedge x_1 \neq y_2 \wedge y_1 \neq x_2 \wedge y_1 \neq y_2$ (i.e. each A_i is a matching in the directed graph $(V, \vec{E})$). If $(u \to v) \in A_i$ for some $i \in \{1, \ldots, k\}$, then it is assumed that the whole current knowledge of the node u is known to the node v after the execution of the i-th round A_i.

b, two-way mode (also called telephone mode)
   In two-way mode, in a single round, each node may be active only via one of its adjacent edges and if it is active then it simultaneously sends a message and receives a message through the given, active edge. Formally, let G be a network. A

two-way communication algorithm for G is a sequence of rounds $B_1, B_2, \ldots, B_r$, where each round $B_j \subseteq E$, and for each $i \in \{1,\ldots,r\}, \forall (x_1,y_1),(x_2,y_2) \in B_i : (x_1,y_1) \neq (x_2,y_2)$ implies $x_1 \neq x_2 \wedge x_1 \neq y_1 \wedge y_1 \neq y_2 \wedge x_2 \neq y_1$ (i.e. B_i is a matching in G). If $(u,v) \in B_i$ for some i, then it is assumed that the whole current knowledge of u is submitted to v, and the whole current knowledge of v is submitted to u in the i-th round.

Another possibility to describe a communication algorithm for $G = (V,E)$ is to say what happens on which edge of E in which time unit. For every two-way communication algorithm $A = A_1, A_2, \ldots, A_m$ for G, edge (A) $= \{(e, S_e^A) \mid e \in E, S_e^A = \{j | e \in A_j\}$ for any $e \in E\}$.

Analogously, for every one-way communication algorithm $B = B_1, \ldots, B_m$ for G, edge (B) $= \{(u \to v, S^B_{u\to v}) \mid (u,v) \in E$, and $S^B_{u\to v} = \{j|(u \to v) \in B_j\}$ for any $u \to v$ such that $(u,v) \in E\}$. We define edge (C) $\leq$ edge (D) for two two-way (one-way) communication algorithms C, D for $G = (V,E)$, if for every edge $e \in E$, $S_e^C \subseteq S_e^D$ (if for every $(u \to v)$ such that $(u,v) \in E$, $S^C_{u\to v} \subseteq S^D_{u\to v}$). Obviously, if C is a broadcast (accumulation, gossip) algorithm for some G, and edge$(C) \leq$ edge (D) for some communication algorithm D for G, then D is a broadcast (accumulation, gossip) algorithm for G, too.

Now, we are prepared to introduce systolic communication algorithms.

Definition 2.1 *Any one-way (two-way) communication algorithm $A = A_1, A_2, A_3, \ldots, A_m$ for some $m \in I\!N$ is called k-systolic for some positive integer k, if there exist some $r \in \{1,\ldots,m\}$ and some $j \in \{1,\ldots,k\}$ such that*

$$A = (A_1, A_2, \ldots, A_k)^r, A_1, A_2, \ldots, A_j.$$

$P = A_1, \ldots, A_k$ is called the period/cycle of A, k is called the length of P. □

In what follows the complexity of communication algorithms is considered as the number of rounds they consist of.

Definition 2.2 *Let $G = (V,E)$ be a network. Let $A = A_1, A_2, \ldots, A_m$ be a one-way (two-way) communication algorithm on G. The complexity of A is $c(A) = m$ (the number of rounds of A). Let v be a node of V. The one-way complexity of the broadcast problem for G and v is*

$$b_v(G) = min\ \{c(A) | A \text{ is a one-way communication algorithm solving the broadcast problem for } G \text{ and } v\}.$$

The one-way complexity of the accumulation problem for G and v is

$$a_v(G) = min\ \{c(A) | A \text{ is a one-way communication algorithm solving the accumulation problem for } G \text{ and } v\}.$$

We define

$\underline{b(G)} = max\ \{b_v(G)|v \in V\}$ *as the* *broadcast complexity of G,*

$\underline{minb(G)} = min\ \{b_v(G)|v \in V\}$ *as the* *min-broadcast complexity of G,*

$\underline{a(G)} = max\ \{a_v(G)|v \in V\}$ *as the* *accumulation complexity of G, and*

$\underline{mina(G)} = min\ \{a_v(G)|v \in V\}$ *as the* *min-accumulation complexity of G.*

The one-way gossip complexity of G is

$$\underline{r(G)} = min\ \{c(A)|A \text{ is a one-way communication algorithm solving the gossip problem for } G\},$$

and the two-way gossip complexity of G is

$$\underline{r_2(G)} = min\ \{c(A)|A \text{ is a two-way communication algorithm solving the gossip problem for } G\}.$$

□

Note that we do not define the broadcast (accumulation) complexity for two-way mode because each two-way broadcast (accumulation) algorithm for a network G can be transformed into a one-way broadcast (accumulation) algorithm consisting of the same number of rounds. Observe also that $a_v(G) = b_v(G)$ (and consequently $a(G) = b(G)$, $\text{mina}(G) = \text{minb}(G)$) for any network G and any node v of G (cf. [HKMP93]).

Now, we give the notation for the complexity of systolic broadcast, accumulation, and gossip.

Definition 2.3 *Let $G = (V, E)$ be a network, and let k be a positive integer. Let v be a node of V. The k-systolic complexity of the broadcast problem for G and v is*

$$\underline{[k] - sb_v(G)} = min\ \{c(A)|A \text{ is a one-way } k\text{-systolic communication algorithm solving the broadcast problem for } G \text{ and } v\},$$

and the k-systolic complexity of the accumulation problem for G and v is

$$\underline{[k] - sa_v(G)} = min\ \{c(A)|A \text{ is a one-way } k\text{-systolic communication algorithm solving the accumulation problem for } G \text{ and } v\}.$$

We define

$$\underline{[k] - sb(G)} = max\ \{[k] - sb_v(G)|v \in V\}$$

as the k-systolic broadcast complexity of G, and

$$\underline{[k] - minsb(G)} = min\ \{[k] - sb_v(G)|v \in V\}$$

as the k-systolic min-broadcast complexity of G.

The one-way k-systolic gossip complexity of G is

$$\underline{[k]-sr(G)} = min\ \{c(A) | A \textit{ is a one-way } k\textit{-systolic communication algorithm solving the gossip problem for } G\},$$

and the *two-way k-systolic gossip complexity of G* *is*

$$\underline{[k]-sr_2(G)} = min\ \{c(A) | A \textit{ is a two-way } k\textit{-systolic communication algorithm solving the gossip problem for } G\}. \qquad \square$$

Obviously, each communication algorithm is k-systolic for some sufficiently large k. But we want to consider k-systolic communication algorithms for fixed k for some classes of networks. In this approach, k is a constant independent of the sizes of the networks of the class. This means that our k-systolic algorithms are simply realized by the repetition of a cycle of k simple instructions by any processor of the network.

The "periodic" gossip introduced in [LR93b] is a special case of systolic gossip introduced above. The periodic communication is based on the coloring of edges in G, which means that each edge can be used at most once in one period (cycle). We are giving no restriction on the number of occurrences of an edge in the rounds of a period, and Section 4 shows that this can be helpful for designing quick communication algorithms. Moreover, the periodic communication based on coloring works for two-way communication mode only, i.e. the one-way mode was not considered in [LR93b, LHL93]. Finally, the complexity of periodic gossip in [LR93b, LHL93] is measured as the number of executed periods which gives only a rough estimation on the number of rounds sufficient and necessary to solve the given communication problem. In our systolic concept we prefer to precisely measure the complexity of communication tasks as the number of rounds executed. This approach enables also a precise comparison of the systolic gossip and the general gossip. For some networks G we can even prove that $r(G) = [k]$–$sr(G)$ for some suitable constant k, i.e. that some optimal gossip algorithm can be systolized.

We observe that each k-systolic algorithm uses (activates) at most k adjacent edges of every node of the network during the whole work of the algorithm. Thus, there is no reason to consider classes of networks like hypercubes and complete graphs, because a k-systolic algorithm can use only a subgraph of these graphs with the degree bounded by k. For this reason, we shall investigate systolic complexity of broadcast and gossip for constant-degree bounded classes of networks only. Our aim is not only to get some lower and upper bounds on the systolic broadcast and gossip complexity of some concrete networks, but also to compare the general complexities of unrestricted communication algorithms with the systolic ones. In this way, we can learn which is the price for our systolization, i.e. how many additional rounds are needed to go from an optimal broadcast (gossip) algorithm to an optimal systolic one.

Our first result shows that, in some sense, the broadcast complexity is the same as the systolic broadcast complexity for any network.

Lemma 2.4 *For every class of networks* $\{G_i\}_{i=1}^{\infty}$ *and every positive integer* d *such that* G_i *has the degree at most* d *for any* $i \in \mathbb{N}$,

$$[d] - sb_v(G_i) = b_v(G_i)$$

for any $i \in \mathbb{N}$ *and any node* v *of* G_i.

Proof. Let $B = B_1, B_2, \ldots, B_m$ be a broadcast algorithm from v in $G_i = (V, E)$ for some positive integer i. W.l.o.g. we may assume $B_i \bigcap B_j = \emptyset$ for $i \neq j$. One can reconstruct B to get another broadcast algorithm $A = A_1, A_2, \ldots, A_m$ for v and G_i with the following three properties:

(i) $\bigcup_{i=1}^{m} A_i \subseteq \bigcup_{j=1}^{m} B_j$, $c(B) = c(A)$,

(ii) $T = (V, \bigcup_{i=1}^{m} A_i)$ is a directed tree with the root v (all edges are directed from the root to the leaves), and

(iii) for every node $w \in V$, if w gets $I(v)$ in time t and the degree (indegree plus outdegree) of w in T is k, then w submits $I(v)$ in the rounds $t+1, t+2, \ldots, t+k-1$ to all its $k-1$ descendants in T.

The construction of broadcast algorithm A' from B with the properties (i) and (ii) is simple and can be found in Observation 1.2.5 of [HKMP93]. How to get the property (iii) (i.e. to construct A from A') is obvious.

To define a d-systolic broadcast algorithm $C = (C_1, C_2, \ldots, C_d)^r, C_1, C_2, \ldots, C_j$ for some $r \in I\!N$, $j \in \{0, \ldots, d-1\}$, $m = r \cdot d + j$, it is sufficient to specify C_i for $i = 1, \ldots, d$. For $1 \leq s \leq d$, let $In(s) = \{n | 1 \leq n \leq m, n \bmod d = s - 1\}$. Then,

$$C_s = \bigcup_{\ell \in In(s)} A_\ell \text{ for } s \in \{1, \ldots, d\}.$$

Since edge$(A) \leq$ edge(C), it is obvious that C is a broadcast algorithm for G and v. □

Clearly, the same consideration as in the proof of Lemma 2.4 leads to the following result.

Lemma 2.5 *Let d be a positive integer. Let $\{G_i\}_{i=1}^{\infty}$ be a class of networks, where for every $i \in I\!N$, G_i has degree bounded by d. Then, for every $i \in I\!N$ and every node v of G_i*

$$[d] - sa_v(G_i) = a_v(G_i).$$ □

The next important question is which relation holds between gossip complexity and systolic gossip complexity. The following sections show that, as opposed to broadcast (accumulation), there are already essential differences between the complexities of general gossip and systolic gossip. Here we shall still deal with the relation between gossip and broadcast. It is well-known (see, for instance [HKMP93]) that $r(G) \leq$ mina(G)+ minb$(G) = 2\cdot$minb(G) and $r_2(G) \leq 2\cdot$minb$(G) - 1$ for any graph G. Note that for trees [BHMS90] and some cyclic graphs with "weak connectivity" [HJM93] the equalities $r(G) = 2\cdot$minb(G) and $r_2(G) = 2\cdot$minb$(G) - 1$ hold. The idea of the proof of $r(G) \leq$ mina(G)+ minb(G) is very simple: One node of G first accumulates $I(G)$, and then it broadcasts $I(G)$ to all other nodes. Unfortunately, we cannot use this scheme to get systolic gossip from systolic broadcast and systolic accumulation, because we have to use every edge of an optimal broadcast (accumulation) scheme in both directions in each repetition of the cycle of a systolic gossip algorithm which already increases the time for the broadcast phase twice. Thus, using this straightforward idea we only obtain the following.

Theorem 2.6 *Let G be a communication network of degree bounded by some positive integer k. Then*

$$[2k]\text{–}sr(G) \leq 4 \cdot [k]\text{–}minsb(G) + 2k.$$

Proof. Let $B = B_1, B_2, \ldots, B_m$, $m = \text{minb}(G) = [k]\text{–minsb}(G)$, be an optimal broadcast algorithm for G and some node v of G with the properties (i), (ii), (iii) as in Lemma 2.4. Let $T_B = (V, \bigcup_{i=1}^{m} B_i)$. Obviously, $A = A_1, A_2, \ldots, A_m$, where $A_i = \vec{B}_{m-i+1} = \{(x \to y) | (y \to x) \in B_{m-i+1}\}$ is an optimal accumulation algorithm for G and v. Moreover, $T_A = (V, \bigcup_{i=1}^{m} A_i)$ is the same scheme as T_B, only the edges are directed in opposite directions. Building

$$C_s = \bigcup_{j \in In(s)} B_j \text{ for } s \in \{1, \ldots, k\} \text{ (where } In(s) \text{ as defined in Lemma 2.4),}$$

$$D_s = \bigcup_{j \in In(s)} A_j \text{ for } s \in \{1, \ldots, k\},$$

one obtains optimal k-systolic broadcast and accumulation algorithms $C = (C_1, C_2, \ldots, C_k)^r$, $C_1, C_2, \ldots, C_j$ and $D = (D_1, \ldots, D_k)^r, D_1, D_2, \ldots, D_j$ resp. for some positive integer r, and $j \in \{0, \ldots, k-1\}$. Now, we consider the $2k$-systolic communication algorithm

$$F = (D_1, \ldots, D_k, C_1, \ldots, C_k)^{2r+1}.$$

The initial part of F, $(D_1, \ldots, D_k, C_1, \ldots, C_k)^r D_1, \ldots, D_k$, is a one-way $2k$-systolic accumulation algorithm for G and the node v. The rest, $C_1, \ldots, C_k(D_1, \ldots, D_k, C_1, \ldots, C_k)^r$, is a one-way $2k$-systolic broadcast algorithm for G and the node v. Thus, F is a one-way $2k$-systolic gossip algorithm with $c(F) = 2k \cdot (2r+1) = 4r \cdot k + 2k \leq 4 \cdot [k]\text{–minsb}(G) + 2k$. $\square$

The next sections deal with the systolic gossip problem in concrete networks. The next Section 3 is devoted to gossiping in paths P_n. For systolic algorithms in two-way (telephone) communication mode, the optimal gossip algorithm for paths is in fact a 2-systolic communication algorithm. For the one-way (telegraph) communication mode, upper and lower bounds on $[k]\text{–}sr(P_n)$ are proved which differ only in a small constant independent of n and the length k of the period. More precisely, we show that for any $n \geq 2$, $k \geq 4$,

$$[k]\text{–}sr(P_n) = \frac{k}{k-2} \cdot (n-2) + c_{n,k} \quad \text{for some constant } 0 \leq c_{n,k} \leq 3.$$

As a consequence, we obtain that for the one-way communication mode one can systolically gossip faster in P_n with a longer period k. More precisely, for growing period k, the function $r(P_n)$ of n can be approached more and more but never achieved (namely $[k]\text{–}sr(P_n) \approx (1 + \varepsilon_k) \cdot r(P_n)$ and $\lim_{k \to \infty} \varepsilon_k = 0$).

Section 4 is devoted to gossiping in complete k-ary trees T_k^h for $k \geq 2$. Surprisingly we show for sufficiently large periods d (d independent of the depth h of the tree and depending only on the degree k) that $[d]\text{–}sr_2(T_k^h) = r_2(T_k^h) = 2\cdot\text{minb}(T_k^h)-1$ and $[d]\text{–}sr(T_k^h) = r(T_k^h) = 2\cdot\text{minb}(T_k^h)$ for any $h \in I\!N$, i.e. we can systolically gossip in complete trees in optimal gossip time in both modes. We also show for the minimal possible period length $d = k+1$ of any two-way systolic communication algorithm for k-ary trees that $[k+1]\text{–}sr_2(T_k^h) \leq r_2(T_k^h) + 1$.

3 Systolic Gossiping in Paths

In this section we consider systolic gossiping in the path P_n of n nodes.

For the two-way mode, we can find a systolic gossip algorithm for P_n with an optimal period length that works as efficiently as the algorithm in the general gossip mode.

Theorem 3.1

(i) $[2]\text{–}sr_2(P_n) \;=\; n-1 \;=\; r_2(P_n)$ *for even* $n \geq 2$,

(ii) $[2]\text{–}sr_2(P_n) \;=\; n \;=\; r_2(P_n)$ *for odd* $n \geq 3$.

Proof. The lower bounds for the general gossip mode are presented in [HKMP93]. The upper bounds are variations of the algorithms described in [HKMP93].

Algorithm A for P_n (where $V(P_n) = \{x_1, \ldots, x_n\}$, $E(P_n) = \{(x_1, x_2), \ldots, (x_{n-1}, x_n)\}$) has the following systolic period:

$$A_1 = \{(x_1, x_2), (x_3, x_4), (x_5, x_6), \ldots\},$$

$$A_2 = \{(x_2, x_3), (x_4, x_5), (x_6, x_7), \ldots\}.$$

A simple analysis shows that Algorithm A takes $n-1$ rounds if n is even and n rounds if n is odd. $\square$

Note that for an optimal gossip in P_n, the period lengths in Theorem 3.1 are the best possible. Any systolic algorithm for P_n must have at least period length 2.

Let us turn to the one-way mode of communication now. For the complexity of systolic gossiping in the path P_n of n nodes, we obtain upper and lower bounds which are tight up to a constant. An important observation contrasting to the two-way case is that there is no constant d such that $[d]\text{–}sr(P_n) = r(P_n)$ for every $n \in I\!N$. Instead, the next theorem shows that one can essentially gossip faster in P_n with a longer period k. For growing k, $r(P_n)$ can be approached more and more but never achieved.

Theorem 3.2 *For any* $n \geq 2$, $k \geq 4$*:*

(i) $\frac{k}{k-2}\cdot(n-2) \;\leq\; [k]\text{–}sr(P_n) \;\leq\; \frac{k}{k-2}\cdot n - 1$ *for* k *even,*

(ii) $\frac{k}{k-2}\cdot(n-2) \;\leq\; [k]\text{–}sr(P_n) \;\leq\; \frac{k}{k-2}\cdot(n-1) + 1$ *for* k *odd.*

Proof. Let us first describe the upper bounds.

(i), upper bound for even k:

Let us first assume that n is a multiple of $k-2$. Then the path P_n is divided into subpaths $B_1, B_2, \ldots, B_{n/(k-2)}$ of $k-2$ nodes as follows:

$$B_i \;:=\; \{(k-2)\cdot(i-1)+1, (k-2)\cdot(i-1)+2, \ldots, (k-2)\cdot i\} \quad \text{for } 1 \le i \le n/(k-2).$$

In each period, the systolic one-way gossip algorithm A does the following:

1. Gossip in B_i for all $1 \le i \le n/(k-2)$.
2. Exchange the information between the endnodes of adjacent blocks.

As $k-2$ is even, Step 1. takes $k-2$ rounds by using the well-known gossip algorithm described e.g. in [HKMP93] in each block B_i, $1 \le i \le n/(k-2)$. Step 2. can be achieved in 2 rounds. Thus, the period length of the systolic algorithm is k.

For a complete gossip, it is enough to ensure that the message I_1 from the left end of the path reaches the right end of the path, and that the message I_2 from the right end of the path reaches the left end. With each period (of length k), I_1 moves one block to the right, and I_2 moves one block to the left. Hence, after $n/(k-2)-1$ periods, i.e. after $k\cdot(n/(k-2)-1)$ rounds, I_1 has moved to block $B_{\lceil n/(k-2)\rceil}$, and I_2 has moved to block B_1. Now, the gossip in the first $k-2$ rounds of the next period suffices to get I_1 and I_2 to the endpoints. Hence, the overall time is at most

$$k\cdot\left(\frac{n}{k-2}\right)-2.$$

If n is not a multiple of $k-2$, we consider the gossip scheme A for a path $P_{n'}$ of $n' := (k-2)\cdot\lceil\frac{n}{k-2}\rceil$ nodes (where $V(P_{n'}) = \{x_1, x_2, \ldots, x_{n'}\}$). Consider the subpath $P_n = \{x_{\ell+1}, \ldots, x_{\ell+n}\}$ of n nodes where

$$\ell \;:=\; \left\lfloor\frac{(k-2)\cdot\lceil\frac{n}{k-2}\rceil-n}{2}\right\rfloor \;\ge\; \frac{k-2}{2}\,.$$

Then the scheme A restricted from $P_{n'}$ to P_n achieves gossiping on P_n in

$$\begin{aligned}
&\left(k\cdot\left\lceil\frac{n}{k-2}\right\rceil-2\right)-2\cdot\left(\left\lfloor\frac{(k-2)\cdot\left\lceil\frac{n}{k-2}\right\rceil-n}{2}\right\rfloor+1\right)\\
\le\;&\left(k\cdot\left\lceil\frac{n}{k-2}\right\rceil-2\right)-\left((k-2)\cdot\left\lceil\frac{n}{k-2}\right\rceil-n\right)-1\\
=\;&2\cdot\left\lceil\frac{n}{k-2}\right\rceil+n-3\;\le\;\frac{k}{k-2}\cdot n-1
\end{aligned}$$

rounds (starting from round ℓ). To make A start with round 1 from the node x_ℓ, the old systolic period $A_1, \ldots, A_k$ has to be rotated by $\ell-1$ positions to $A_\ell, A_{\ell+1}, \ldots, A_k, A_1, \ldots, A_{\ell-1}$.

(ii), upper bound for odd k:

Let us first assume that $n-1$ is a multiple of $k-2$. Then the path P_n is divided into subpaths $B_1, B_2, \ldots, B_{(n-1)/(k-2)}$ of $k-1$ nodes as follows:

$$\begin{aligned} B_i \; &:= \; \{(k-2)\cdot(i-1)+1, (k-2)\cdot(i-1)+2, \ldots, (k-2)\cdot i+1\} \\ & \quad \text{for } 1 \leq i \leq (n-1)/(k-2). \end{aligned}$$

Note that two adjacent blocks overlap by one node. In each period, the systolic one-way gossip algorithm performs a complete gossip in each block B_i for all $1 \leq i \leq (n-1)/(k-2)$. For doing this, the well-known gossip algorithm described e.g. in [HKMP93] is used in each block. As the number of nodes, $k-1$, is even in each block, the gossip takes $k-1$ rounds. The only problem is how communication conflicts between two adjacent blocks can be avoided. For this purpose, note that the gossip algorithm from [HKMP93] for gossiping in a path $P_m = \{x_1, \ldots, x_m\}$ of m nodes only uses the edges (x_1, x_2) and (x_{m-1}, x_m) in the first and the last round. Hence, if we start the gossip in blocks B_i, i odd, in the first round of each systolic period, and in blocks B_i, i even, in the second round of the period, blocks B_i, i odd, only use their leftmost and rightmost edge in rounds 1 and $k-1$ of the period, and blocks B_i, i even, only use their leftmost and rightmost edge in rounds 2 and k of each period. This way, no conflict will occur.

For the analysis of the complexity of the algorithm, let us consider again the messages I_1 and I_2 from the left and the right end of the path. With each period (of length k), I_1 moves one block to the right, and I_2 moves one block to the left. Hence, after $(n-1)/(k-2)$ periods, i.e. after

$$k \cdot \left(\frac{n-1}{k-2}\right)$$

rounds, I_1 has reached the right end of the path and I_2 the left end.

If $n-1$ is not a multiple of $k-2$, we consider the gossip scheme A for a path $P_{n'}$ of $n' := (k-2) \cdot \lceil \frac{n-1}{k-2} \rceil + 1$ nodes (where $V(P_{n'}) = \{x_1, x_2, \ldots, x_{n'}\}$). Consider the subpath $P_n = \{x_{\ell+1}, \ldots, x_{\ell+n}\}$ of n nodes where

$$\ell \; := \; \left\lfloor \frac{(k-2)\cdot\lceil\frac{n-1}{k-2}\rceil - (n-1)}{2} \right\rfloor \; \geq \; \left\lceil \frac{k-2}{2} \right\rceil \; = \; \frac{k-1}{2} \, .$$

Then the scheme A restricted from $P_{n'}$ to P_n achieves gossiping on P_n in

$$\begin{aligned} & k \cdot \left\lceil \frac{n-1}{k-2} \right\rceil - 2 \cdot \left(\left\lfloor \frac{(k-2)\cdot\left\lceil\frac{n-1}{k-2}\right\rceil - (n-1)}{2} \right\rfloor + 1 \right) \\ \leq \; & \frac{k}{k-2} \cdot (n-1) + 1 \end{aligned}$$

rounds (starting from round ℓ). To make A start with round 1 from the node x_ℓ, the old systolic period $A_1, \ldots, A_k$ has to be rotated by $\ell-1$ positions to $A_\ell, A_{\ell+1}, \ldots, A_k, A_1, \ldots, A_{\ell-1}$.

This completes the proof of the upper bounds of Theorem 3.2. To derive the lower bounds, we start by introducing the concept of viewing a gossip algorithm as a set of time-paths. This concept has been successfully used in [HJM93] to get an optimal gossip algorithm for cycles.

Definition 3.3 *Let $G = (V, E)$ be a graph, and let $X = x_1, \ldots, x_m$ be a simple path (i.e. $x_i \neq x_j$ for $i \neq j$) in G. Let $A = A_1, \ldots, A_s$ be a communication algorithm in one–way mode. Let $T = t_1, \ldots, t_{m-1}$ be an increasing sequence of positive integers such that $(x_i \to x_{i+1}) \in A_{t_i}$ for $i = 1, \ldots, m-1$. We say that $X[t_1, \ldots, t_{m-1}]$ is a time–path of A because it provides the information flow from x_1 to x_m in A. If $t_{i+1} - t_i - 1 = k_i \geq 0$ for some $i \in \{1, \ldots, m-2\}$ then we say that $X[t_1, \ldots, t_{m-1}]$ has a k_i-delay at the node x_{i+1}. The global delay of $X[t_1, \ldots, t_{m-1}]$ is $d(X[t_1, \ldots, t_{m-1}]) = t_1 - 1 + \sum_{i=1}^{m-2} k_i$. The global time of $X[t_1, \ldots, t_{m-1}]$ is $m - 1 + d(X[t_1, \ldots, t_{m-1}])$.*

Obviously, the necessary and sufficient condition for a communication algorithm to be a gossip algorithm in a graph G is the existence of time–paths between all pairs of nodes in G.

So, one can view the gossip algorithm for a graph G as a set of time–paths between any ordered pair of nodes.

The complexity (the number of rounds) of a communication algorithm can be measured as

$$\max\{\text{global time of } X[T] \mid X[T] \text{ is a time-path of } A\}.$$

To see a gossip algorithm as a set of time–paths is mainly helpful for proving lower bounds. A conflict of two time–paths (the meeting of two time–paths going in "opposite directions" at the same node and at the same time of the systolic period) causes some delays in these time–paths (because of the restriction given by the communication modes). Too many unavoidable conflicts mean too many delays, and so one can get much better lower bounds for gossiping in some graphs G than the trivial diameter lower bound. A combinatorial analysis providing lower bounds by analyzing the number of conflicts and delays requires a precise definition and use of these two notions. Thus, we define these notions for the one–way communication mode and the path P_n as follows. Note that the next definition essentially differs from the definition of conflicts in [HJM93] because that definition allows at most one conflict between two time-paths going in opposite directions on the same physical path of the network. The essential point here is that the time-paths from one end-point to another are realized in a systolic manner and that this systolic realisation causes conflicts and delays in nodes where the crucial information pieces flowing between the two end-points do not meet in a physical time.

Definition 3.4 *Let $P_n = (\{x_1, x_2, \ldots, x_n\}, \{(x_i, x_{i+1}) | i = 1, \ldots, n\})$ be a path of n nodes. Let $A = (A_1, A_2, \ldots, A_k)^r A_1, A_2, \ldots, A_s$ be a k-systolic one-way gossip algorithm for P_n for some $k, r \in I\!N$, $0 \leq s < k$. Let $X = x_1, x_2, \ldots, x_n$ and $Y = x_n, x_{n-1}, \ldots, x_1$. An X-direction of A is $X[A] = (S_1, S_2, \ldots, S_{n-1})$, where $S_i = \{1 \leq j \leq k | (x_i \to x_{i+1}) \in A_j\}$ for $i = 1, \ldots, n-1$. A Y-direction of A is $Y[A] = (Q_1, Q_2, \ldots, Q_{n-1})$, where $Q_i = \{1 \leq j \leq k | (x_{i+1} \to x_i) \in A_j\}$ for $i = 1, \ldots, n-1$.*

For any $p \in \{1, \ldots, n-1\}$ and $q \in \{1, \ldots, k\}$, let

$$\begin{aligned} \text{next}_q(x_p \to x_{p+1}) &= \min\{a, b+k \mid a, b \in S_p \text{ and } q < a \leq k,\ 1 \leq b \leq q\}, \text{ and} \\ \text{next}_q(x_{p+1} \to x_p) &= \min\{a, b+k \mid a, b \in Q_p \text{ and } q < a \leq k,\ 1 \leq b \leq q\}. \end{aligned}$$

Let i be a positive integer, $1 \leq i \leq k$. We say that the directions of A $X[A]$ and $Y[A]$ have an i-collision in a node x_m for some $m \in \{2, \ldots, n-1\}$ if one of the following four conditions holds:

(i) $i \in S_{m-1}$, $\text{next}_i(x_{m+1} \to x_m) \;<\; \text{next}_i(x_m \to x_{m+1}) \;<\; \text{next}_i(x_m \to x_{m-1})$,

(ii) $i \in Q_m$, $\text{next}_i(x_{m-1} \to x_m) \;<\; \text{next}_i(x_m \to x_{m-1}) \;<\; \text{next}_i(x_m \to x_{m+1})$,

(iii) $i \in S_{m-1}$, $\text{next}_i(x_{m+1} \to x_m) \;<\; \text{next}_i(x_m \to x_{m-1}) \;<\; \text{next}_i(x_m \to x_{m+1})$,

(iv) $i \in Q_m$, $\text{next}_i(x_{m-1} \to x_m) \;<\; \text{next}_i(x_m \to x_{m+1}) \;<\; \text{next}_i(x_m \to x_{m-1})$.

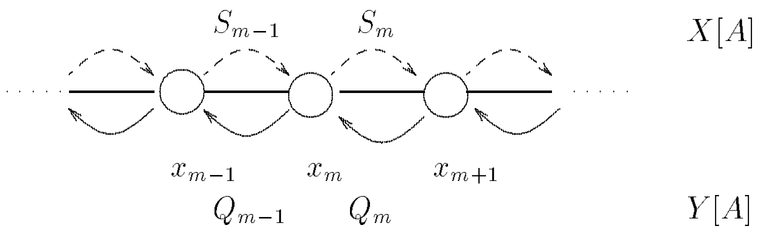

Fig. 1: An i-collision in a node x_m

If one of the cases (i) or (ii) happens, then we say that the i-collision in x_m causes at least a 1–delay on $X[A]$ and at least a 1–delay on $Y[A]$. If (iii) happens, then we say that the i-collision in x_m causes at least a 2–delay on $X[A]$. If (iv) happens, then we say that the i-collision causes at least a 2–delay on $Y[A]$. For every $i \in \{1, \ldots, k\}$, every i-collision in x_m is called a collision in x_m.

Let $X[T]$ and $Y[T']$ be two time-paths of A for some $T = t_1, t_2, \ldots, t_{n-1}$ and $T' = t'_1, t'_2, \ldots, t'_{n-1}$ $= r_{n-1}, \ldots, r_1$. Let $i \in \{1, \ldots, k\}$ and $m \in \{2, \ldots, n-1\}$. We say that the time-path $X[T]$ meets the X-direction $X[A]$ in the node x_m in the i-th periodical moment if

$$(t_{m-1} = kj + d, \text{ where } 1 \le d \le i) \text{ and}$$
$$((t_m = kj + d_1, \text{ where } i < d_1 \le k) \text{ or } (t_m = k(j+1) + d_2, \text{ where } 1 \le d_2 < i)) \ .$$

We say that $Y[T']$ meets $Y[A]$ in the node x_m in the i-th periodical moment if

$$(r_m = kj + d, \text{ where } 1 \le d \le i) \text{ and}$$
$$((r_{m-1} = kj + d_1, \text{ where } i < d_1 \le k) \text{ or } (r_{m-1} = k(j+1) + d_2, \text{ where } 1 \le d_2 < i)) \ .$$

Finally, we say that there is a (systolic) conflict between $X[T]$ and $Y[T']$ in a node x_m if there exists $i \in \{1, \ldots, k\}$ such that $X[A]$ and $Y[A]$ have an i-collision in the node x_m and

(i) $X[T]$ *meets* $X[A]$ *in the node* x_m *in the* $((i+1) \bmod k+1)$*-th periodical moment, and*

(ii) $Y[T']$ *meets* $Y[A]$ *in the node* x_m *in the* $((i+1) \bmod k+1)$*-th periodical moment.* □

We see that the conflict between the time-paths $X[T]$ and $Y[T']$ in a node x_m means that there exist $i \in \{1, \ldots, k\}$, $p, q \in I\!N$ such that

(i) the piece of information $I(x_1)$ flowing to x_n according to $X[T]$ is visiting x_m (still not in x_{m+1}) after the execution of $pk + i + 1$ rounds,

(ii) the piece of information $I(x_n)$ flowing to x_1 according to $Y[T']$ is visiting x_m after the execution of $qk+i+1$ rounds,

(iii) there is an i-collision between $X[A]$ and $Y[A]$ in the node x_m.

Despite of the fact that $I(x_1)$ and $I(x_n)$ possibly do not meet in x_m in the same time (p may differ from q), the systolic realization defined by the directions $X[A]$ and $Y[A]$ causes that the delay on $X[A]$ (Y[A]) caused by the the i-collision in x_m is also a delay on the time-path $X[T]$ ($Y[T']$).

Overall, it can be stated that the sum of the delays on $X[T]$ and $Y[T']$ caused by any conflict is at least 2.

Now, we are prepared to prove the lower bounds of Theorem 3.2. Let $V(P_n) = \{x_1, \ldots, x_n\}$, $E(P_n) = \{(x_1, x_2), \ldots, (x_{n-1}, x_n)\}$, and consider the paths $X = x_1, \ldots, x_n$ and $Y = x_n, \ldots, x_1$. Each gossip algorithm A for P_n must contain two time-paths $X[T]$ and $Y[T']$ for some $T = t_1, \ldots, t_{n-1}$ and $T' = t'_1, \ldots, t'_{n-1}$.

The aim will be to bound the number of conflicts between $X[T]$ and $Y[T']$ from below. The fact that each conflict causes a delay of at least 2 on $X[T]$ and $Y[T']$ will then give rise to a lower bound on the number of rounds of A.

To make the lower bound proof more transparent, we first show a lower bound which is weaker than the one stated in Theorem 3.2, but which is less technical to prove.

Lemma 3.5 *For any $n \geq 2$:*

(i) $[k]$*–*$sr(P_n) \;\geq\; \frac{k+2}{k} \cdot (n-1) - 1$ *for k even,*

(ii) $[k]$*–*$sr(P_n) \;\geq\; \frac{k+1}{k-1} \cdot (n-1) - 1$ *for k odd.*

The proof of Lemma 3.5. The core of the proof is to show that there are at least $\left\lceil \frac{n-1}{\lfloor k/2 \rfloor} \right\rceil - 1$ conflicts in the inner nodes of P_n. This is done by first establishing that the distance between two neighbouring conflicts is not too "large". (Two conflicts c_1 in node x_i and c_2 in x_j, $i < j$, are called *neighbouring* if there is no conflict in $x_{i+1}, \ldots, x_{j-1}$ between $X[T]$ and $Y[T']$.)

Claim 1 The distance between two neighbouring conflicts c_1 in x_i and c_2 in x_j, $i < j$, is at most $\lfloor k/2 \rfloor$.

The proof of Claim 1. Consider the time-paths $X[T] = X[t_1, t_2, \ldots, t_{n-1}]$ and $Y[T'] = Y[t'_1, t'_2, \ldots, t'_{n-1}]$. For any $\ell \in \{1, \ldots, n-1\}$, let *diff*$(\ell) := (t_\ell - t'_\ell) \bmod k$ measure the time-difference in the systolic period between the communication $x_\ell \to x_{\ell+1}$ and $x_{\ell+1} \to x_\ell$, and let

$$\mathrm{reldiff}(i, \ell) := (t_\ell - t'_\ell) - (t_i - t'_i - (t_i - t'_i) \bmod k)$$

denote the time-difference relative to the time-difference in (x_i, x_{i+1}). As there is a conflict c_1 in x_i, we have

$$\mathrm{reldiff}(i, i) \geq 1.$$

Following the definition of the time-paths $X[T]$ and $Y[T']$, we have $t_1 < t_2 < \ldots < t_{n-1}$ and $t'_1 > t'_2 > \ldots > t'_{n-1}$. Hence,

$$\begin{aligned}
&\text{reldiff}(i, i+1) \geq 3,\\
&\text{reldiff}(i, i+2) \geq 5,\\
&\vdots\\
&\text{reldiff}(i, i+s) \geq 2s+1 \quad \text{for } s \geq 1.
\end{aligned}$$

As soon as $\mathit{reldiff}(i, i+s) \geq k$ (i.e. $\mathit{reldiff}(i, i+s-1) < k$ and $\mathit{reldiff}(i, i+s) \geq k$), there is a conflict in x_{i+s}. It follows that there is a conflict in at least one of the nodes $x_{i+1}, x_{i+2}, \ldots, x_{i+s}$ if $2s+1 \geq k$ or $s \geq \lfloor k/2 \rfloor$ respectively. This completes the proof of Claim 1. □

The proof of Claim 1 shows that the largest distance between two neighbouring conflicts c_1 in x_i and c_2 in x_j, $i < j$, can only be achieved if $\mathit{reldiff}(i, i+s)$ is as small as possible for any s, i.e.

$$\begin{aligned}
&\text{reldiff}(i, i) = 1,\\
&\text{reldiff}(i, i+1) = 3,\\
&\vdots\\
&\text{reldiff}(i, i+s) = 2s+1 \quad \text{for } s \geq 1.
\end{aligned}$$

For the corresponding time-paths $X[T]$ and $Y[T']$, this means that

$$\begin{aligned}
&(t_i - t'_i) \bmod k = 1,\\
&(t_{i+1} - t'_{i+1}) \bmod k = 3,\\
&\vdots\\
&(t_{i+s} - t'_{i+s}) \bmod k = 2s+1 \quad \text{for } s \geq 1,
\end{aligned}$$

and for the directions $X[A] = (S_1, S_2, \ldots, S_{n-1})$ and $Y[A] = (Q_1, Q_2, \ldots, Q_{n-1})$ it follows that for $r = \text{Mod}(t_i)$,

$$\begin{array}{ll}
\text{Mod}(r-1) \in Q_i & r \in S_i \\
\text{Mod}(r-2) \in Q_{i+1} & \text{Mod}(r+1) \in S_{i+1} \\
\text{Mod}(r-3) \in Q_{i+2} & \text{Mod}(r+2) \in S_{i+2} \\
\vdots & \vdots
\end{array}$$

(where $\text{Mod}(m) := (m-1) \bmod k + 1$ for any $m \in I\!N$). This optimal pattern of length $\lfloor k/2 \rfloor$ on the X- and Y-direction between the two neighbouring conflicts c_1 and c_2 will be referred to as pattern P^1_{opt}.

The proof of Lemma 3.5 continued. According to Claim 1, there is a conflict in each $\lfloor k/2 \rfloor$ steps. Hence, the number ℓ of conflicts in the inner nodes of P_n is at least $\left\lceil \frac{n-1}{\lfloor k/2 \rfloor} \right\rceil - 1$. As each conflict causes an overall delay of at least 2 on $X[T]$ and $Y[T']$, one of the time-paths incurs a delay of at least ℓ. Therefore, the message on this time-path needs at least $(n-1)+\ell$ rounds. Applying $\ell \geq \left\lceil \frac{n-1}{\lfloor k/2 \rfloor} \right\rceil - 1$ leads to a lower bound on the number of rounds of at least

$$(n-1) + \left(\frac{n-1}{\lfloor k/2 \rfloor} - 1 \right) = \begin{cases} \frac{k+2}{k} \cdot (n-1) - 1 & \text{if } k \text{ is even,} \\ \frac{k+1}{k-1} \cdot (n-1) - 1 & \text{if } k \text{ is odd.} \end{cases}$$

This completes the proof of Lemma 3.5. $\square$

The proof of Theorem 3.2 continued. A technically more involved consideration than the one of Lemma 3.5 provides the precise lower bound of Theorem 3.2.

Lemma 3.6 *For any $n, k \geq 3$:*

$$[k]\text{-}sr(P_n) \;\geq\; \frac{k}{k-2} \cdot (n-2)\;.$$

The proof of Lemma 3.6. The core of the proof is to show an improved lower bound on the number of conflicts in the inner nodes of P_n. To obtain this improved bound, it is not enough to bound the distance between two neighbouring conflicts from above. Instead, we will argue about the distance between s successive conflicts. The improvement in the argument derives from the fact that the average distance between two neighbouring conflicts is less than the maximum distance. Technically, we prove the following fact.

Claim 2 Let $c_1, c_2, \ldots, c_s$ be s successive conflicts. Then the distance between c_1 and c_s is at most $(s-1) \cdot (k/2 - 1) + 1$.

If Claim 2 is true, one can easily complete the proof of Lemma 3.6 in the following way. Using Claim 2, we see that the number s of conflicts in the inner nodes of P_n must satisfy

$$((s+2)-1) \cdot (k/2 - 1) + 1 \;\geq\; n-1$$

(if this inequality is not true, then the inner nodes contain at least $s+1$ conflicts). This implies

$$s \geq 2 \cdot \frac{n-2}{k-2} - 1.$$

As each conflict causes an overall delay of at least 2 on $X[T]$ and $Y[T']$, one of these time-paths has a delay of at least s. Therefore, the number of executed rounds of any k-systolic gossip algorithm on P_n is at least

$$(n-1) + s \;\geq\; \frac{k}{k-2} \cdot (n-2).$$

Thus, to complete the proofs of Theorem 3.2 and Lemma 3.6 it is sufficient to prove Claim 2.

Claim 2 will be proved separately for even k and odd k. The proof itself will be an induction on the number s of conflicts. For the inductive step, an additional property about the structure of the conflicts is needed. Hence, Claim 2 is reformulated in an appropriate way.

For this purpose, let us first specify some further notation. For two conflicts c_1 in x_i and c_2 in x_j, $i < j$, let $dist(c_1, c_s)$ denote the distance between c_1 and c_2 on the path, i.e. the number of edges between x_i and x_j. Consider the time-paths $X[T] = X[t_1, t_2, \ldots, t_{n-1}]$ and $Y[T'] = Y[t'_1, t'_2, \ldots, t'_{n-1}]$. Let c be a conflict in some inner node x_i of P_n. Then $rdiff(c) := (t_i - t'_i) \bmod k$ and $ldiff(c) := (t'_{i-1} - t_{i-1}) \bmod k$. $rdiff(c)$ $[ldiff(c)]$ measures the time-difference in the systolic period between the communication $x_i \to x_{i+1}$ and $x_{i+1} \to x_i$ $[x_{i-1} \to x_i$ and $x_i \to x_{i-1}]$. For the conflict situations described in (i)–(iv) of Definition 3.4, we have the following time-differences:

(i) $\mathrm{ldiff}(c) \geq 3$, $\mathrm{rdiff}(c) \geq 1$,

(ii) $\mathrm{ldiff}(c) \geq 1$, $\mathrm{rdiff}(c) \geq 3$,

(iii) $\mathrm{ldiff}(c) \geq 2$, $\mathrm{rdiff}(c) \geq 2$,

(iv) $\mathrm{ldiff}(c) \geq 2$, $\mathrm{rdiff}(c) \geq 2$.

Note that the pattern P^1_{opt} achieving optimal length $\lfloor k/2 \rfloor$ between two neighbouring conflicts c_1 and c_2 fulfills $rdiff(c_1) = 1$ and

$$\mathrm{ldiff}(c_2) \;=\; \begin{cases} 1 & \text{if } k \text{ is even,} \\ 2 & \text{if } k \text{ is odd.} \end{cases}$$

Now, we are able to reformulate Claim 2 as follows.

Claim 2a Let $k \in I\!N$, k even. Let $c_1, c_2, \ldots, c_s$ be s successive conflicts, $s \geq 2$. Then the following statements hold:

(1) $\mathrm{dist}(c_1, c_s) \;\leq\; (s-1) \cdot (k/2 - 1) + 1$,

(2) $\mathrm{dist}(c_1, c_s) \;=\; (s-1) \cdot (k/2 - 1) + 1$
$\Longrightarrow\; (\mathrm{rdiff}(c_1) = 1 \;\wedge\; \mathrm{ldiff}(c_s) = 1)$.

Claim 2b Let $k \in I\!N$, k odd. Let $c_1, c_2, \ldots, c_s$ be s successive conflicts, $s \geq 2$. Then the following statements hold:

(1) If s is odd:

(1a) $\mathrm{dist}(c_1, c_s) \;\leq\; \left(\frac{s-1}{2}\right) \cdot (k-2) + 1$,

(1b) $\mathrm{dist}(c_1, c_s) \;=\; \left(\frac{s-1}{2}\right) \cdot (k-2) + 1$
$\Longrightarrow\; (\mathrm{rdiff}(c_1) = 1 \;\wedge\; \mathrm{ldiff}(c_s) = 1)$.

(2) If s is even:

$$\mathrm{dist}(c_1, c_s) \;\leq\; \left(\frac{s-2}{2}\right) \cdot (k-2) + \lfloor k/2 \rfloor \;.$$

Clearly, Claim 2 follows immediately from Claims 2a and 2b. So, the only thing left to show for the proof of Lemma 3.6 is the validity of Claims 2a and 2b.

The proof of Claim 2a. (Induction on s)
The induction base for $s = 2$ is clear from the remark about P^1_{opt} above. Assume that (1) and (2) hold for $s \in I\!N$, $s \geq 2$. We try to extend $X[A]$ and $Y[A]$ from c_s to c_{s+1}.

(i) If $\mathrm{dist}(c_1, c_s) \leq (s-1) \cdot (k/2 - 1)$, then

$$\mathrm{dist}(c_1, c_{s+1}) \;\leq\; \mathrm{dist}(c_1, c_s) + k/2 \;\leq\; ((s+1) - 1) \cdot (k/2 - 1) + 1$$

according to Claim 1 in the proof of Lemma 3.5. Hence, (1) holds for $s+1$.

It remains to show that (2) holds for $s+1$. To obtain $\mathrm{dist}(c_1, c_{s+1}) = ((s+1)-1)\cdot(k/2-1)+1$, $\mathrm{dist}(c_1, c_s) = (s-1)\cdot(k/2-1)$ must hold, and the extension from c_s to c_{s+1} must be of length $k/2$. The only pattern for achieving this is P^1_{opt} from the proof of Claim 1 in Lemma 3.5. For this pattern, $\mathrm{rdiff}(c_s) = \mathrm{ldiff}(c_{s+1}) = 1$ holds. Hence, $\mathrm{ldiff}(c_s) \geq 3$. Consider the conflicts $c_1, c_2, \ldots, c_s$. As $\mathrm{ldiff}(c_s) \geq 3$, the distance between c_1 and c_s can be increased by one by moving c_s one node to the right on the path (and by extending $X[A]$ and $Y[A]$). Then, the distance between c_1 and c_s is $(s-1)\cdot(k/2-1)+1$, and the induction hypothesis yields $\mathrm{rdiff}(c_1) = 1$. Hence, (2) holds for $s+1$.

(ii) If $\mathrm{dist}(c_1, c_s) = (s-1)\cdot(k/2-1)+1$, then $\mathrm{rdiff}(c_1) = \mathrm{ldiff}(c_s) = 1$ according to the induction hypothesis. Hence, the extension from c_s to c_{s+1} must start with $\mathrm{rdiff}(c_s) = 3$. The same argumentation as in the proof of Claim 1 in Lemma 3.5 shows that $\mathrm{dist}(c_s, c_{s+1}) \leq k/2-1$ must hold, and if $\mathrm{dist}(c_s, c_{s+1}) = k/2-1$ then $\mathrm{ldiff}(c_{s+1}) = 1$. Hence, (1) and (2) hold for $s+1$.

This completes the proof of Claim 2a. $\square$

The proof of Claim 2b (1). (Induction on s)
Let $s = 3$. Then

$$\mathrm{dist}(c_1, c_s) \;=\; \mathrm{dist}(c_1, c_3) \;\leq\; 2\cdot\lfloor k/2\rfloor \;=\; k-1 \;=\; \left(\frac{s-1}{2}\right)\cdot(k-2)+1$$

follows from Claim 1 in the proof of Lemma 3.5. The only way to achieve $\mathrm{dist}(c_1, c_3) = k-1$ is to construct $X[A]$ and $Y[A]$ by using the pattern P^1_{opt} from the proof of Claim 1 in Lemma 3.5 between c_1 and c_2, which leads to $\mathrm{rdiff}(c_1) = 1$, $\mathrm{ldiff}(c_2) = 2$. Hence, an optimal extension from c_2 to c_3 has to start with $\mathrm{rdiff}(c_2) = 2$, has length $\lfloor k/2\rfloor$ and leads to $\mathrm{ldiff}(c_3) = 1$ (by using the same arguments as in the proof of Claim 1 of Lemma 3.5). This whole optimal pattern between c_1 and c_3 is referred to as pattern P^2_{opt}.

Assume that (1a) and (1b) hold for $s \in I\!N$, $s \geq 3$ odd. We try to extend $X[A]$ and $Y[A]$ from c_s to c_{s+2}.

(i) If $\mathrm{dist}(c_1, c_s) \leq \left(\frac{s-1}{2}\right)\cdot(k-2)$, then

$$\mathrm{dist}(c_1, c_{s+2}) \;\leq\; \mathrm{dist}(c_1, c_s) + (k-1) \;\leq\; \left(\frac{(s+2)-1}{2}\right)\cdot(k-2)+1\;.$$

Hence, (1a) holds for $s+2$.

It remains to show that (1b) holds for $s+2$. To obtain $\mathrm{dist}(c_1, c_{s+2}) = \left(\frac{(s+2)-1}{2}\right)\cdot(k-2)+1$, $\mathrm{dist}(c_1, c_s) = \left(\frac{s-1}{2}\right)\cdot(k-2)$ must hold, and the extension from c_s to c_{s+2} must be of length $k-1$. The only pattern for achieving this is P^2_{opt}. For this pattern, $\mathrm{rdiff}(c_s) = \mathrm{ldiff}(c_{s+2}) = 1$ holds. Hence, $\mathrm{ldiff}(c_s) \geq 3$. Consider the conflicts $c_1, c_2, \ldots, c_s$. As $\mathrm{ldiff}(c_s) \geq 3$, the distance between c_1 and c_s can be increased by one by moving c_s one node to the right on the path (and by extending $X[A]$ and $Y[A]$). Then, the distance between c_1 and c_s is $\left(\frac{s-1}{2}\right)\cdot(k-2)+1$, and the induction hypothesis yields $\mathrm{rdiff}(c_1) = 1$. Hence, (1b) holds for $s+2$.

(ii) If $\mathrm{dist}(c_1, c_s) = \binom{s-1}{2} \cdot (k-2) + 1$, then $\mathrm{rdiff}(c_1) = \mathrm{ldiff}(c_s) = 1$ according to the induction hypothesis. Hence, the extension from c_s to c_{s+2} must start with $\mathrm{rdiff}(c_s) = 3$. The same argumentation as for P^2_{opt} shows that $\mathrm{dist}(c_s, c_{s+2}) \leq k-2$ must hold, and if $\mathrm{dist}(c_s, c_{s+2}) = k-2$ then $\mathrm{ldiff}(c_{s+2}) = 1$. Hence, (1a) and (1b) hold for $s+2$. □

The proof of Claim 2b (2). According to Claim 2b (1), $\mathrm{dist}(c_1, c_{s-1}) \leq \binom{s-2}{2} \cdot (k-2) + 1$. If $\mathrm{dist}(c_1, c_{s-1}) \leq \binom{s-2}{2} \cdot (k-2)$, then

$$\mathrm{dist}(c_1, c_s) \;\leq\; \mathrm{dist}(c_1, c_{s-1}) + \mathrm{dist}(c_{s-1}, c_s) \;\leq\; \binom{s-2}{2} \cdot (k-2) + \lfloor k/2 \rfloor$$

according to Claim 1 of Lemma 3.5. If $\mathrm{dist}(c_1, c_{s-1}) = \binom{s-2}{2} \cdot (k-2) + 1$, then $\mathrm{rdiff}(c_1) = \mathrm{ldiff}(c_{s-1}) = 1$ according to Claim 2b (1). Hence, $\mathrm{rdiff}(c_s) \geq 3$, $\mathrm{dist}(c_{s-1}, c_s) \leq \lfloor k/2 \rfloor - 1$, and it follows that

$$\mathrm{dist}(c_1, c_s) \;\leq\; \mathrm{dist}(c_1, c_{s-1}) + \mathrm{dist}(c_{s-1}, c_s) \;\leq\; \binom{s-2}{2} \cdot (k-2) + \lfloor k/2 \rfloor$$

This completes the proof of Claim 2b. □

This completes the proof of Lemma 3.6 and Theorem 3.2, too. □

Thus, Theorem 3.2 provides upper and lower bounds on $[k]\text{–}sr(P_n)$ which differ only in a small constant independent of n and k.

Corollary 3.7 *For any $n \geq 2$, $k \geq 4$:*

$$[k]\text{–}sr(P_n) = \frac{k}{k-2} \cdot (n-2) + c_{n,k} \quad \textit{for some constant } 0 \leq c_{n,k} \leq 3.$$

4 Systolic Gossip in k-ary trees

In this section we investigate the systolic gossip complexity of complete, balanced k–ary trees. The main result of this section is that there exist systolic gossip algorithms with constant period whose complexity matches the lower bound for even non–systolic algorithms.

Let us first state the lower bound for gossiping in complete, balanced k–ary trees. It is shown in [BHMS90] that the gossip complexity in two–way mode $r_2(T)$ for any tree T is exactly $2 \cdot minb(T) - 1$, and that for one–way mode $r_1(T) = 2 \cdot minb(T)$ holds. For a complete, balanced k–ary tree T^h_k of height h it is not hard to see, that $minb(T^h_k)$ is given by $k \cdot h$ (for a proof consult [FHMMM92, HKMP93]). This implies the following proposition.

Proposition 4.1 *For a complete, balanced k–ary tree T^h_k of height h and any period p*

i) $[p] - sr(T^h_k) \geq r(T^h_k) = 2 \cdot k \cdot h$

ii) $[p] - sr_2(T^h_k) \geq r(T^h_k) = 2 \cdot k \cdot h - 1$

To describe our algorithms we introduce the following notations. In a systolic algorithm with period p each vertex has to repeat a communication pattern of length p. For the two-way mode of communication we specify such a pattern by a string of length p over the alphabet $C_1, C_2, \ldots, C_k, P, N$. The semantics of this specification is that any vertex v performs a communication with its i-th child (parent, resp.) in round j, iff the pattern of v contains C_i (P, resp.) at position j mod p. The letter N indicates that no communication is performed. In one-way mode we use the alphabet $C_1^{\uparrow}, C_1^{\downarrow}, \ldots, C_k^{\uparrow}, C_k^{\downarrow}, P^{\uparrow}, P^{\downarrow}, N$, where $\uparrow$ ($\downarrow$,resp.) indicates that the flow of information is directed towards the root (towards the leaves, resp.). A gossip algorithm can now be given by specifying a communication pattern for each vertex. Note that the patterns of incident vertices have to be compatible in the sense that whenever the pattern of some vertex v being the i-th child of its parent $p(v)$ indicates a parent communication ($P, P^{\uparrow}$ or $P^{\downarrow}$), the pattern of $p(v)$ has to contain the matching communication ($C_i, C_i^{\uparrow}$, or $C_i^{\downarrow}$) at the corresponding position.

Another point of view emphasizing on this compatibility constraint is to specify a round of communication by a (directed) matching in the tree, where vertices communicate in the given round, iff an edge from the matching connects the vertices. Thus a sequence of p matchings can be used alternatively to specify a systolic algorithms with period p.

Note that there exists no systolic algorithm of period $\leq k$, if $h > 1$, because in this case there are vertices of degree $k+1$. Any algorithm with period $\leq k$ would ignore some edge and no information between the components of the tree connected by this edge can be exchanged.

Now we are able to state our first result, namely a nearly optimal gossiping scheme with minimal period in two mode of communication.

Theorem 4.2 *For $k \geq 2$ and $h \geq 0$*

$$[k+1] - sr_2(T_k^h) \leq 2 \cdot k \cdot h.$$

Proof. We give a gossiping scheme of period $k+1$ by specifying the communication pattern of every vertex. All occurring patterns are cyclic shifts of $S = (P, C_1, C_2, \ldots, C_k)$, provided we substitute the parent communication for the root and the child communications for all leaves by N. In the following we will assume that these obvious substitutions are applied where appropriate, without explicit mention. Let $S_i = (C_i, C_{i+1}, \ldots, P, \ldots C_{i-1})$ be the pattern obtained by cyclically shifting the string S i positions to the left. Thus $S = S_0 = S_{k+1}$ holds. The actual patterns for the gossiping scheme are now obtained recursively as follows:

i) the root uses pattern $S_{h \bmod (k+1)}$,

ii) if v is the i–th child of $p(v)$ and $p(v)$ uses S_j, then v uses $S_{(j-i) \bmod (k+1)}$.

Four simple observations are in order:

1. The patterns are chosen in such a way that the parent communication of each vertex, being the i–th child of its parent, aligns with letter C_i in the pattern of its parent. Thus the given patterns obey the compatibility constraints.

2. The subtree of the first child of the root performs the given gossiping scheme for T_k^{h-1}. And all vertices of the i-th subtree perform the pattern of the corresponding vertex of the first tree shifted $(i-1)$ positions to the right.

3. In round $k(h-1)+i \equiv i+1-h \equiv h+i(\bmod (k+1))$ the root performs a communication with its i–th child according to $S_{h \bmod (k+1)}$.

4. The leftmost leaf has pattern $S_0 = S$, and therefore starts with a parent communication.

We now show by induction on h, that this communication scheme performs simultaneously a fast accumulation and a fast broadcasting with respect to the root. From these results we then can conclude our claim. First, we show that after $k \cdot h$ rounds the cumulative message of T_k^h is known to the root.

For $h = 0$ this statement is true. Assume that it holds for all trees T_k^{h-1}, $h > 0$. In T_k^h we consider now the i–th child r_i of the root. By induction hypothesis and observation 2.) we can conclude that the cumulative message of r_i's subtree is known to r_i after round $k \cdot (h-1)+i-1$, for $1 \leq i \leq k$. Observation 3.) now states that in the next round $k \cdot (h-1)+i$ the cumulative message of the i–th subtree is given to the root. Thus after round $k \cdot (h-1)+k = k \cdot h$ all messages have arrived in the root. At this point it is also worthwhile to mention that not only the root holds the cumulative message after $k \cdot h$ rounds, but also its k–th child. This is a consequence of the two-way mode of communication. Before round $k \cdot h$ the root knows at least all messages not contained in the k-th subtree and its k-th child knows the complementary information. Since the information in two-way mode is exchanged, both vertices learn the cumulative message in this last round.

Next we consider the broadcast capabilities of our scheme. By induction on h it follows that any information known to the root of T_k^h before round t is broadcasted to all vertices after round $t + kh - 1$, if in round t a communication with its first child is performed, and after round $t + kh$, otherwise. For the induction step we observe that all children of the root obtain the broadcast information before round $t+k$, if in round t the root communicates with its first child, and before round $t + k + 1$, otherwise. Since for all vertices each parent communication is directly followed by a communication with the first child, we can inductively assume that the broadcast in the subtrees is finished after round $(t + k) + k(h-1) - 1 = t + kh - 1$, or round $(t + k + 1) + k(h-1) - 1 = t + kh$, respectively.

Concerning the gossip complexity of the communication scheme we now can argue as follows. After kh rounds the cumulative message is known to the root and its k–th child. According to the communication pattern of the root in round $kh + i + 1$ the i-th child is informed. The broadcasting of the cumulative message in the i-th subtree is therefore finished after round $kh + i + 2 + k(h-1) - 1$, for $1 \leq i \leq k-1$, and in the k–th subtree of the root after round $kh + 1 + k(h-1) - 1$. Thus the time critical subtree is the $(k-1)$-st subtree. The broadcast in this tree and the entire gossip is finished after round $2kh$. □

The above algorithm is not time-optimal. When the root has received the cumulative message for the first time — after kh rounds — this message is delayed by one round because of the N in the communication pattern of the root. To overcome this delay the root should perform a pattern like $(C_1, C_2, \ldots, C_k, C_1, C_2, \ldots, C_{k-1}, \ldots)$. But such a pattern does not fit within $k+1$ rounds, thus we have to increase the period to proof the following Theorem.

This new time-optimal algorithm will consist of two parts. Most of the node of T_k^h will perform exactly the same pattern as specified in Theorem 4.2. To be precise, the period is $2 \cdot (k+1)$ and the pattern is $C_1, C_2, \dots, C_k, P, C_1, C_2, \dots, C_k, P$. The nodes of the three top levels of the tree will follow some special patterns. Note that by using a at period of $2 \cdot (k+1)$ the optimal algorithm for a T_k^1 is already systolic. We use now the optimal algorithm for the communication within the top two levels. Level three produces the correct interaction between the top part and level four, where the algorithm for subtrees from Theorem 4.2 is implemented. Before presenting the new algorithm in Theorem 4.3 we take a closer look at the algorithm from Theorem 4.2. The subtree rooted at the first son of the root is the only one which has to start the communication in the first round. As a consequence all other subtrees may start its communication one round earlier. The subtree rooted at the $(k-1)$-th son of the root is the only one which has to communicate in round $2 \cdot k \cdot h$. Thus we have to shift this subtree by one round. Due to our first observation is this possible iff the $(k-1)$-th son of the root is different to the first son of the root. Thus the next Theorem 4.3 deals with the case $k \geq 3$ and the case $k = 2$ is solved in Theorem 4.4

Theorem 4.3 *For $k \geq 3$ and $h \geq 0$*

$$[2 \cdot (k+1)] - sr_2(T_k^h) = 2 \cdot k \cdot h - 1 = r_2(T_k^h) \textit{ holds.}$$

Proof.

Within this new algorithm are several patterns:

$$\begin{aligned}
S^r &= (N, C_1, C_2, \dots, C_k, C_1, C_2, \dots, C_{k-1}, N, N) \\
S^s &= (P, \underbrace{N, \dots, N}_{k-1 \text{ times}}, P, N, C_1, C_2, \dots, C_k) \\
S^{s'} &= (P, \underbrace{N, \dots, N}_{k-1 \text{ times}}, P, C_1, C_2, \dots, C_k, N) \\
S^t &= (N, C_1, C_2, \dots, C_k, P, C_1, C_2, \dots, C_k) \\
S^u &= (P, C_1, C_2, \dots, C_k, P, C_1, C_2, \dots, C_k)
\end{aligned}$$

Let S_i^r $(S_i^s, S_i^{s'}, S_i^t, S_i^u)$ be the pattern obtained from S^r $(S^s, S^{s'}, S^t, S^u)$ by cyclically shifting the string i positions to the left. The gossiping scheme for T_k^h is defined in the following recursive way:

i) the root uses pattern $S^r_{h \bmod (2 \cdot (k+1))}$

ii) the i-th child of the root $(1 \leq i \leq k-2)$ uses pattern $S^s_{(h-i) \bmod (2 \cdot (k+1))}$

iii) the $k-1$-th child of the root uses pattern $S^{s'}_{(h-k+1) \bmod (2 \cdot (k+1))}$

iv) the k-th child of the root uses pattern $S^t_{(h+1) \bmod (2 \cdot (k+1))}$

v) if v is the i–th child of $p(v)$ and $p(v)$ uses S_j^s, then v uses $S^t_{(j-i) \bmod 2 \cdot (k+1))}$.

vi) if v is the i–th child of $p(v)$ and $p(v)$ uses $S_j^{s'}$, then v uses $S^t_{(j-i+1) \bmod 2 \cdot (k+1))}$.

vii) if v is the i–th child of $p(v)$ and $p(v)$ uses S^t_j or S^u_j, then v uses $S^u_{(j-i)\mathrm{mod}2\cdot(k+1))}$.

Several simple observations are in order:

1. The pattern are chosen such that the parent communication of each vertex, being the i–th child of its parent, aligns with letter C_i in the pattern of its parent. Thus the given patterns obey the compatibility constraints. This is true because of the following:
   - If the root uses pattern S^r_j then the i-th son $(1 \le j \le k-1)$ uses pattern S^s_{j-i} or $S^{s'}_{j-i}$ and the last son uses pattern S^t_{j+1}. Note that S^s and $S^{s'}$ have the P communication at the same position.
   - If a node uses pattern S^s_j then the i-th son $(1 \le j \le k)$ uses pattern S^t_{j-i}.
   - If a node uses pattern $S^{s'}_j$ then the i-th son $(1 \le j \le k)$ uses pattern S^t_{j-i+1}. Note that the communications with the children in $S^{s'}$ are shifted one position to the left compared with S^s.
   - If a node uses pattern S^t_j or S^u_j then the i-th son $(1 \le j \le k)$ uses pattern S^u_{j-i}.

   It is easy to see the the compatibility constraints are valid in all cases.
2. If a node v uses pattern S^t then the subtree rooted at v performs the communication pattern from Theorem 4.2. This is true because all descendants will use the pattern S^u.
3. The leftmost leaf has pattern $S^v_0 = S$, and therefore starts with a parent communication.
4. Let f_i be the i-th son of the root. The subtree of f_k performs the given gossiping scheme for T^{h-1}_k from Theorem 4.2. Note that f_k uses pattern S^t.
5. Let f_{11} the first son of f_1. The subtree of f_{11} perform the given gossiping scheme for T^{h-2}_k from Theorem 4.2. Note that f_{11} use pattern S^t.
6. The Node f_{11} sends the cumulative message to f_1 without delay. Due to Theorem 4.2 at time $k \cdot (h-2)$ the cumulative message of the subtree rooted at f_{11} has arrived in f_{11}. The node f_{11} uses pattern $S^t_{(h-2)\mathrm{mod}2\cdot(k+1)}$ which has a parent communication at time $k \cdot (h-2) + 1$.
7. Any node v at level two sends the cumulative message to its parent $p(v)$ without delay.
8. Within the top three levels of T^h_k any cumulative message is passed to the parent node without any delay. Thus the root receives the cumulative message at time kh. A more detailed proof from fact 1 will produce also this result.
9. The root sends the cumulative message one step prior to the algorithm from Theorem 4.2.
10. A node f_i $(1 \le i \le k-2)$ delays the cumulative message by one step before sending it to its sons. Thus all nodes within the subtree rooted at f_i receive the cumulative message within $2 \cdot k \cdot h - 1$ steps. Note that within this part of the tree the new algorithm behaves like the algorithm from Theorem 4.2. But in that algorithm there is only one leaf which receives the cumulative message at time $2 \cdot k \cdot h$. This leaf is a descendent of f_{k-1}

11. All nodes within the subtree rooted at f_{k-1} receive the cumulative message within $2 \cdot k \cdot h - 1$ steps. Note that by using pattern $S^{s'}$ the sending down of the cumulative message is not delayed.

12. All nodes within the subtree rooted at f_k receive the cumulative message within $2 \cdot k \cdot h - k$ steps.

From all the above remarks we conclude the validity of this algorithm. $\square$

Note that the above algorithm works not in the case $k = 2$. But using the same technic from Theorem 4.3 we get the following theorem:

Theorem 4.4 *For binary trees of height $h \geq 0$*

$$[9] - sr_2(T_2^h) = 4 \cdot h - 1 = r_2(T_2^h) \textit{ holds.}$$

Proof. We just define the algorithm. Within this algorithm are several patterns:

$$\begin{aligned}
S^r &= (N, C_1, C_2, C_1, N, N, N, N, N)\\
S^s &= (P, N, P, C_1, C_2, N, N, C_1, C_2)\\
S^t &= (N, N, P, N, C_1, C_2, P, N, N, N)\\
S^{t'} &= (N, N, P, C_1, C_2, N, P, N, N, N)\\
S^u &= (N, C_1, C_2, P, C_1, C_2, N, C_1, C_2)\\
S^v &= (P, C_1, C_2, P, C_1, C_2, P, C_1, C_2)
\end{aligned}$$

Let S_i^r $(S_i^s, S_i^t, S_i^{t'}, S_i^u, S_i^v)$ be the pattern obtained from S^r $(S^s, S^t, S^{t'}, S^u, S^v)$ by cyclically shifting the string i positions to the left. The gossiping scheme for T_k^h is defined in the following recursive way:

i) the root uses pattern $S^r_{h \bmod 9}$

ii) the i-th child of the root $(1 \leq i < k)$ uses pattern $S^s_{(h-i) \bmod 9}$

iii) the k-th child of the root uses pattern $S^u_{(h-k) \bmod 9}$

iv) if v is the i–th child of $p(v)$ and $p(v)$ is the first son of the root and $p(v)$ uses S_j^s, then v uses $S^t_{(j-i) \bmod 9}$.

v) if v is the i–th child of $p(v)$ and $p(v)$ is not the first son of the root and $p(v)$ uses S_j^s, then v uses $S^{t'}_{(j-i) \bmod 9}$.

vi) if v is the i–th child of $p(v)$ and $p(v)$ uses S_j^t, then v uses $S^u_{(j-i) \bmod 9}$.

vii) if v is the i–th child of $p(v)$ and $p(v)$ uses $S_j^{t'}$, then v uses $S^u_{(j-i+1) \bmod 9}$.

viii) if v is the i–th child of $p(v)$ and $p(v)$ uses S_j^u, then v uses $S^v_{(j-i) \bmod 9}$.

ix) if v is the i–th child of $p(v)$ and $p(v)$ uses S_j^v, then v uses $S^v_{(j-i) \bmod 9}$.

Using the arguments similar to the ones from Theorem 4.3 we conclude the correctness and time bound of this algorithm. $\square$

Next we turn to the one-way mode of communication. We will derive a systolic gossip algorithm, that requires $2kh+1$ rounds and has a period of $(3+\lceil\frac{3}{k-1}\rceil)(k+1)$. Thus for binary trees a period of length 18, for ternary trees a period of length 20, and for k-ary trees with $k \geq 4$ a period of length $4(k+1)$ is suffices.

For $i \geq 0$, we will recursively derive communication schemes $A^i_\uparrow$ and $A^i_\downarrow$ having the following properties when applied to T^h_k:

These communication schemes perform a parent communication at the root. Such a communication is interpreted as an output from the root to the environment (if $P^\uparrow$ is specified), or as an input to the root of the tree from the environment (if $P^\downarrow$ is specified).

- P1 $A^i_\downarrow$ as well as $A^i_\uparrow$ perform gossiping in $2kh+1+i$ rounds.
- P2 $A^i_\downarrow$ guarantees accumulation of all messages in the root after $kh+i$ rounds, and outputs the cumulative message to the environment in round $kh+i+k+2 = k(h+1)+i+2$. Moreover $A^i_\uparrow$ broadcasts any message received from the environment in round r using kh additional rounds, i.e the message is distributed after round $r+kh$.
- P3 $A^i_\uparrow$ guarantees accumulation of all messages in the root after $kh+i$ rounds also, but outputs the cumulative message to the environment in round $kh+i+1$. Moreover $A^i_\uparrow$ broadcasts any message received from the environment in round r using $k(h+1)+1$ additional rounds, i.e the message is distributed after round $r+k(h+1)+1$.

Surely, given communication schemes $A^0_\uparrow$ and $A^0_\downarrow$ it is trivial to obtain schemes for $A^i_\uparrow$ and $A^i_\downarrow$ just by shifting cyclically all communication patterns of $A^0_\uparrow$, and $A^0_\downarrow$ respectively, i positions to the right. Thus we concentrate on deriving schemes $A^0_\uparrow$ and $A^0_\downarrow$.

Up to cyclic shifts h^*+1 different patterns of length $(h^*+1)(k+1)$ are used, for some appropriate constant h^* to be chosen later. All vertices of depth $d \equiv 0(\mathrm{mod} h^*)$ use (up to cyclic shifts) one of the following $[0 \bmod h^*]$–patterns

$$\text{either } (C^\uparrow_1,\ldots,C^\uparrow_k,P^\downarrow,C^\downarrow_1,\ldots,C^\downarrow_k,P^\uparrow,N,\ldots,N)$$
$$\text{or } (C^\uparrow_1,\ldots,C^\uparrow_k,P^\uparrow,C^\downarrow_1,\ldots,C^\downarrow_k,N,\ldots,N,P^\downarrow)$$

All vertices of depth $d \equiv j(\mathrm{mod} h^*), 1 \leq j \leq h^*-1$ use (up to cyclic shifts) the $[j \bmod h^*]$–pattern

$$(C^\uparrow_1,\ldots,C^\uparrow_k,P^\uparrow,\underbrace{N,\ldots,N}_{(j(k+1)-1)\text{times}},P^\downarrow,C^\downarrow_1,\ldots,C^\downarrow_k,N,\ldots,N)$$

Note that after an appropriate shift these patterns are compatible to each other. Each pattern involves two communications with the same child, say C_ℓ. More precisely, if $C^\uparrow_\ell$ occurs in a $[(j-1) \bmod h^*]$–pattern at positions t, then $C^\downarrow_\ell$ occurs at position $t+j(k+1)$. The same distance, namely $j(k+1)$ occurs in the $[j \bmod h^*]$–pattern for the next level between $P^\uparrow$ and

$P^{\downarrow}$. Note especially that for both [0 mod h^*]–patterns the cyclic distance between $P^{\uparrow}$ and $P^{\downarrow}$ is exactly $h^*(k+1)$.

Consider now $T_k^{h^*}$. By fixing the pattern for the root and using the pattern given above, we indeed fix the entire communication scheme, provided the compatibilty constraints are obeyed. (Note that in the leaves we have to use one of the root patterns, but since all child communications are substituted by N, both patterns become indistinguishable when cyclical shifts are allowed). In general the only choice in designing our communication scheme after fixing the patterns is the choice of the pattern in vertices of depth $d \equiv 0(\mathrm{mod} h^*)$.

For scheme $A^0_{\uparrow}$ we use in the root the pattern $(C_1^{\uparrow}, \ldots, C_k^{\uparrow}, P^{\uparrow}, C_1^{\downarrow}, \ldots, C_k^{\downarrow}, N, \ldots, N, P^{\downarrow})$ shifted such that $C_k^{\uparrow}$ is performed in round kh^*. Similarily, for scheme $A^0_{\downarrow}$ we use the pattern $(C_1^{\uparrow}, \ldots, C_k^{\uparrow}, P^{\downarrow}, C_1^{\downarrow}, \ldots, C_k^{\downarrow}, P^{\uparrow}, N, \ldots, N)$ again shifted such that $(C_k^{\uparrow})$ is performed in round kh^*. Actually, for $T_k^{h^*}$ both choices lead to exactly the same scheme, except for the parent communications of the root. It is instructive to have a closer look at the execution of this scheme. For this purpose, it will be convenient to have the notion of signatures of vertices.

Definition 4.5 *Assume that each edge e in T_k^h connecting some arbitrary inner vertex v with its j–th child is labelled with $\ell(e) = j$, for $1 \leq j \leq k$. Let $e_1, \ldots, e_s$ be the edges lying on the unique path from vertex w to the root. Then the signature of w is defined as*

$$\mathrm{sig}(w) = \sum_{i=1}^{s} \ell(e_i)$$

Assume that in round r the root performs operation $C_k^{\uparrow}$ for some sufficiently large r. Since for any inner vertex v any information received from its j-th child has to wait exactly $k - j$ rounds in v before it is delivered to the next level or until round r is elapsed, in case v is the root, any information given from leaf ℓ to its parent in round $r - k(h^* - 1) + \mathrm{sig}(\ell) + 1$ will arrive at the root before or in round r. Note that the incurred overall delay time, i.e. the sum of delays incured in the vertices on the path from ℓ to the root, for an information starting in leaf ℓ is just $kh^* - \mathrm{sig}(\ell)$, since the delay in depth d and the contribution of the edge between depth d and $d + 1$ on the root path of ℓ to $\mathrm{sig}(\ell)$ add up exactly to k, for $0 \leq d < h$.

An obvious consequence is that accumulation in the root requires only kh^* rounds, if the root pattern is adjusted such that operation $C_k^{\uparrow}$ is performed in round kh^*. This is because leaf ℓ performs $P^{\uparrow}$ in round $\mathrm{sig}(\ell) - h^* \geq 1$, and therefore all informations reach the root before round kh^*.

With respect to the broadcast capabilities of this scheme we observe that at any vertex v, except possibly at the root, an information sent to the j–th child incurs a delay of $j - 1$. For any message obtained from the environment, the delay in the root is either $j - 1$ (when using $A^0_{\downarrow}$) or $(j - 1) + (k + 1)$ (when using $A^0_{\uparrow}$). Thus the total delay incurred by a message received by the root and being forwarded to ℓ is either $\mathrm{sig}(\ell) - h^*$ or $\mathrm{sig}(\ell) - h^* + (k + 1)$. Since the forward path has length h^* the message arrives in ℓ at round $r' + \mathrm{sig}(\ell)$ or $r' + \mathrm{sig}(\ell) + (k + 1)$ respectively, where r' is the round in which the root has received the input message. Since $h^* \leq \mathrm{sig}(\ell) \leq kh^*$ all leaves have obtained the message after kh^* additional rounds, in case $A^0_{\downarrow}$ is used, and after $k(h^* + 1) + 1$ additional rounds, if $A^0_{\uparrow}$ is used.

The gossip in $T_k^{h^*}$ requires $2kh^* + 1$ rounds. Accumulation in the root is finished after kh^*, the next round is declared in the root as parent communication and then the broadcast is started.

Note that the broadcast requires kh^* rounds by either communication scheme, since in both schemes the root causes a delay of $j-1$ for messages sent to the j-th child, if round kh^*+1 is considered as first round of the broadcast. Indeed, we can make a slightly more accurate statement. If each leaf ℓ delivers its message for the first time in round $\mathrm{sig}(\ell)-(h^*-1)$, then ℓ receives the cumulative message in round $\mathrm{sig}(\ell)-(h^*-1)+h^*(k+1)$.

Now it is easy to check that these schemes restricted to the first c levels of T_k^h achieve the accumulation-, broadcast- and gossip capabilities postulated in P1, P2, and P3 for T_k^c with $c \leq h^*$.

The extension of these schemes for T_k^h of arbitrary height $h > h^*$ is quite easy. Assume that we have constructed appropriate schemes for trees of height $h-h^*$. We cut off the first h^*+1 levels of T_k^h. This toptree of height h^* is handled exactly as $T_k^{h^*}$, where the root pattern is adjusted such that $C_k^{\uparrow}$ is performed in round kh.

Now consider the a subtree T_ℓ of T_k^h rooted at an arbitrary leaf ℓ of the toptree. If $\mathrm{sig}(\ell) \leq h^*+k$ we use scheme $A_{\uparrow}^{\mathrm{sig}(\ell)-(h-1)}$ for this subtree, otherwise we use $A_{\downarrow}^{\mathrm{sig}(\ell)-(h-1)-(k+1)}$.

Note that according to P2 and P3 ℓ, the root of T_ℓ, delivers the cumulative message of T_ℓ to its parent in round $kh-k(h^*-1)+\mathrm{sig}(\ell)+1$, independent of the pattern used for T_ℓ. This guarantees that accumulation in the root is finished after round kh, as well as the compatibility of the pattern used in ℓ with the pattern in its parent. Moreover the cumulative message of T_k^h is given as output in round $kh+1$ for $A_{\uparrow}^0$, and in round $k(h+1)+2$ for $A_{\downarrow}^0$, as required.

To analyse the broadcast properties of the scheme, we consider scheme $A_{\downarrow}^0$ and note that any input message received by the root in round r arrives in leaf ℓ of the toptree in round $r+\mathrm{sig}(\ell)$ for $A_{\downarrow}^0$. Inductively we may assume that the broadcast in T_ℓ requires additional $k(h-h^*)$ rounds, if $\mathrm{sig}(\ell) > h^*+k$. In this case all vertices of T_ℓ have received the message after round $r+k(h-h^*)+\mathrm{sig}(\ell) \leq kh$, since $\mathrm{sig}(\ell) \leq kh^*$ holds for any leaf of the toptree. If $\mathrm{sig}(\ell) \leq h^*+k$, then scheme $A_{\uparrow}^{\mathrm{sig}(\ell)-h^*}$ is used for T_ℓ, requiring $k(h-h^*)+k+1$ additional rounds to broadcast in T_ℓ. In this case the broadcast in T_ℓ is finished after round $r+kh-(k-1)h^*+2k+2$. Note that $r+kh-(k-1)h^*+2k+2 < r+kh+1$ holds, iff $h^* \geq 2+\lceil\frac{3}{k-1}\rceil$ holds. In case scheme $A_{\downarrow}^0$ is used, obviously $k+1$ additional rounds, due to delays in the root, are required for the broadcast in T_ℓ. Thus the required broadcast capabilities are achieved only in case $h^* \geq 2+\lceil\frac{3}{k-1}\rceil$ holds. Recall that the required period is h^*+1.

To see that the given scheme performs a gossip in $2kh+1$ rounds, we note that accumulation in the root is achieved in kh rounds. Leaf ℓ of the toptree thus receives the cumulative message in round $kh+1+\mathrm{sig}(\ell)$. Performing the same analysis as in the case of broadcast (i.e. substituting r by $kh+1$) yields the claimed result.

Summarizing the above discussion we get.

Theorem 4.6 *For $k \geq 2$ and $h \geq 0$*

$$\left[\left(3+\left\lceil\frac{3}{k-1}\right\rceil\right)(k+1)\right] - sr(T_k^h) \leq 2\cdot k\cdot h+1.$$

As in case of two–way mode, the gossip scheme given achieves a runtime that requires just one round of communication more than stated in the lower bound. In the rest of this section we will speed up the scheme by one round at the cost of increasing slightly the length of the period using a idea very similar to the one applied in two–way mode.Thus we obtain systolic gossip schemes with an optimal number of communication rounds.

The improvement is based on the observation, that only a few vertices receive the cumulative message in the last round. We first consider the one–way mode of communication.

Recall the gossip scheme applied to T_k^h. Let ℓ be any leaf of the toptree of height $h*$. The cumulative message was broadcast successfully in T_ℓ after round $kh + 1 + k(h - h^*) + \mathrm{sig}(\ell)$, if $\mathrm{sig}(\ell) > h^* + k$. Thus, among all these subtrees, only the rightmost one requires $2kh + 1$ round to broadcast the cumulative message, because for all leaves ℓ of the toptree, except the rightmost one, $\mathrm{sig}(\ell) \leq kh^* - 1$ holds. If $\mathrm{sig}(\ell) \leq h^* + k$, then the broadcast of the cumulative message is finished in T_ℓ after round $kh+1+k(h-h^*)+\mathrm{sig}(\ell)+k+1 \leq 2kh+2k+2-(k-1)h^*$. If we now choose $h^* \geq 2 + \lceil\frac{4}{k-1}\rceil$, all these subtrees finish their broadcast in or before round $2kh$. Thus inductively we can conclude, that whenever we apply the previous scheme with a period of $(3 + \lceil\frac{4}{k-1}\rceil)(k + 1)$ only one vertex, namely the rightmost leaf has not received the cumulative message after round $2kh$. We will now modify the scheme in such a way, that all vertices receive the cumulative message either in the same round as before or one round earlier. Especially all vertices in the rightmost subtree of the root will receive the cumulative message one round earlier, which guarantees a gossip complexity of $2kh$ rounds.

The modification is as follows:

1. The pattern of the root is changed from
   $(C_1^\uparrow, \ldots, C_k^\uparrow, P^\downarrow, C_1^\downarrow, \ldots, C_k^\downarrow, P^\uparrow, N, \ldots, N)$ to
   $(C_1^\uparrow, \ldots, C_k^\uparrow, C_1^\downarrow, \ldots, C_k^\downarrow, N, \ldots, N)$
2. The pattern of the j–th child r_j of the root for $1 \leq j \leq k-1$ is changed from
   $(C_1^\uparrow, \ldots, C_k^\uparrow, P^\uparrow, \underbrace{N, \ldots, N}_{k \text{ times}}, P^\downarrow, C_1^\downarrow, \ldots, C_k^\downarrow, N, \ldots, N)$ to
   $(C_1^\uparrow, \ldots, C_k^\uparrow, P^\uparrow, \underbrace{N, \ldots, N}_{(k-1) \text{ times}}, P^\downarrow, N, C_1^\downarrow, \ldots, C_k^\downarrow, N, \ldots, N)$
3. The pattern of the k–th child r_k of the root is changed from
   $(C_1^\uparrow, \ldots, C_k^\uparrow, P^\uparrow, \underbrace{N, \ldots, N}_{k \text{ times}}, P^\downarrow, C_1^\downarrow, \ldots, C_k^\downarrow, N, \ldots, N)$ to
   $(C_1^\uparrow, \ldots, C_k^\uparrow, N, P^\uparrow, \underbrace{N, \ldots, N}_{(k-1) \text{ times}}, P^\downarrow C_1^\downarrow, \ldots, C_k^\downarrow, N, \ldots, N)$
4. All vertices in the subtree rooted at r_k, except r_k, obtain the pattern of the corresponding vertex in the subtree rooted at r_{k-1}.

The alignment of these patterns in time is such that the root pattern, performs operation $C_k^\uparrow$ in round kh. All other patterns are then fixed by the compatibility constraints. This modification has the effect that the subtrees rooted at r_k and at r_{k-1} now work absolutely synchronously, except for the parent communication in r_k and at r_{k-1}. Moreover, all vertices with depth $d > 1$ not in the subtree rooted at r_k perform exactly the same communication as in the unmodified scheme. Especially, they deliver the cumulative message of their subtrees in the same round as before, and also expect the overall cumulative message in the same rounds as before. It is now easy to check that $r_1, \ldots, r_{k-1}$ indeed perform their child communications exactly as before, especially that r_j, for $j \leq k-1$, holds the cumulative message of its subtree after round $k(h-1)+j$ and passes the overall cumulative message in round $kh+j+i+1$ to its i–th child. It follows that the broadcast of the cumulative message in the subtree at r_j for $j \leq k-1$ is finished after round $2kh$ as before. Since the broadcast in the subtree at r_k is performed synchronously with the broadcast in the subtree of r_{k-1}, we can conclude that gossip is performed in $2kh$ rounds.

This yields

Theorem 4.7 *For $k \geq 2$ and $h \geq 0$*

$$\left[\left(3 + \left\lceil \frac{4}{k-1} \right\rceil\right)(k+1)\right] - sr(T_k^h) \leq 2 \cdot k \cdot h = r(T_k^h)$$

Note that the lengths of the periods in the time-optimal and nearly time-optimal gossip schemes from Theorem 4.6 and Theorem 4.7 differ only for $k = 2$ and 4.

5 Conclusion

Here we discuss the results achieved and formulate some of the main resulting open problems.

In this paper we have introduced the concept of systolic communication. In Section 2 we have shown that the complexity of systolic gossip is at most four times the complexity of systolic min-broadcast. This contrasts to the general relation $r(G) \leq 2 \cdot \text{minb}(G)$ for any G [BHMS90].

Open problem 1. Can the multiplicative constant 4 in the result $[2k]$–$sr(G) \leq 4 \cdot [k]$–$minsb(G) + 2k$ of Theorem 2.6 improved? Note that 2 does not suffice because due to Theorem 3.2, $[k]$–$sr(P_n) \geq d_k \cdot 2 \cdot \text{minb}(P_n)$, where $d_k > 1$ for any $k \in I\!N$, holds. On the opposite, trees are the hardest graphs for the relation between general gossip and min-broadcast ($r(T) = 2 \cdot \text{minb}(T)$ for any tree T), and we can prove $[d]$–$sr(T_k^h) = 2 \cdot \text{minb}(T_k^h)$ for some suitable constant d. This gives us hope for a much better relation between systolic gossip and broadcast than the relation given in Theorem 2.6.

Section 4 shows that we can systolically gossip in T_k^h in the optimal gossip time $r(T_k^h)$ $[r_2(T_k^h)]$. We only have to pay for this with a systolic period longer than the minimal possible period length $k+1$ $[2k+2]$ for one-way [two-way] systolic communication algorithms for T_k^h.

Open problem 2. What is the minimal period length for a time-optimal gossip? Which time can be achieved by a $[2k+2]$-systolic one-way gossip algorithm? (An upper bound of $3kh - (k-1)$ can easily be obtained for $[2k+2]$-systolic gossiping.)

References

[BHMS90] A. Bagchi, S. L. Hakimi, J. Mitchem, E. Schmeichel: *Parallel algorithms for gossiping by mail.* Information Process. Letters 34 (1990), 197-202.

[CC84] C. Choffrut, K. Culik II: *On real-time cellular automata and trellis automata.* Acta Informatica 21 (1984), 393-407.

[CGS83] K. Culik II, J. Gruska, A. Salomaa: *Systolic automata for VLSI on balanced trees.* Acta Informatica 18 (1983), 335-344.

[CGS84] K. Culik II, I. Gruska, A. Salomaa: *Systolic trellis automata. Part I.* Intern. J. Comput. Math. 15 (1984), 195-212.

[CSW84] K. Culik II, A. Salomaa, D. Wood: *Systolic tree acceptors.* R.A.I.R.O. Theoretical Informatics 18 (1984), 53-69.

[FHMMM92] R. Feldmann, J. Hromkovič, S. Madhavapeddy,B. Monien, P. Mysliwietz: *Optimal algorithms for dissemination of information in generalized communication modes.* Proc. PARLE'92, Lecture Notes in Computer Science 605, Springer Verlag 1992, pp. 115-130.

[HHL88] S. M. Hedetniemi, S. T. Hedetniemi, A. L. Liestmann: *A survey of gossiping and broadcasting in communication networks.* Networks 18 (1988), 319-349.

[HJM93] J. Hromkovič, C.-D. Jeschke, B. Monien: *Note an optimal gossiping in some weak-connected graphs.* Theoretical Computer Science, to appear.

[HKMP93] J. Hromkovič, R. Klasing, B. Monien, R. Peine: *Dissemination of information in interconnection networks (broadcasting & gossiping).* In: Combinatorial Network Theory (Frank Hsu, Ding-Zhu Du, Eds.), Science Press, AMS 1993, to appear.

[IK84] O. H. Ibarra, S. M. Kim: *Characterizations and computational complexity of systolic trellis automata.* Theoretical Computer Science 29 (1984), 123-153.

[IKM85] O. H. Ibarra, S. M. Kim, S. Moran: *Sequential machine characterizations of trellis and cellular automata and applications.* SIAM J. Comput. 14 (1985), 426-447.

[IPK85] O. H. Ibarra, M. A. Palis, S. M. Kim: *Fast parallel language recognition by cellular automata.* Theoretical Computer Science 41 (1985), 231-246.

[Ku79] H.T. Kung: *Let's design algorithms for VLSI systems.* In: Proc. of the Caltech Conference of VLSI (CL.L. Seifz Ed.), Pasadena, California 1979, pp. 65-90.

[LHL93] R. Labahn, S.T. Hedetniemi, R. Laskar, *Periodic gossiping on trees.* to appear in: Discrete Applied Mathematics.

[LR93a] A.L. Liestman, D. Richards: *Perpetual gossiping.* Technical Report, Simon Fraser University, Vancouver, Canada, 1993.

[LR93b] A.L. Liestman, D. Richards: *Network communication in edge-colored graphs: Gossiping.* IEEE Trans. Par. Distr. Syst. 4 (1993), 438-445.

[MS90] B. Monien, I. H. Sudborough: *Embedding one interconnection network in another.* Computing Suppl. 7 (1990), 257-282.