Proc. 10th Int. Conference on Fundamentals of Computation Theory (FCT '95)

Effective Systolic Algorithms for Gossiping in Cycles and Two-Dimensional Grids ⋆

(Extended Abstract)

Juraj Hromkovič [1,⋆⋆], Ralf Klasing [2], Dana Pardubská [3,⋆⋆⋆]
Walter Unger [2], Juraj Waczulik [3], Hubert Wagener [2,†]

[1] Institut für Informatik und Praktische Mathematik
Universität zu Kiel, 24098 Kiel, Germany

[2] Department of Mathematics and Computer Science
University of Paderborn, 33095 Paderborn, Germany

[3] Faculty of Mathematics and Physics
Comenius University, 84215 Bratislava, Slovakia

Abstract. The complexity of systolic dissemination of information in one-way (telegraph) and two-way (telephone) communication mode is investigated. The following main results are established:

(i) tight lower and upper bounds on the complexity of one-way systolic gossip in cycles for any length of the systolic period,

(ii) optimal one-way and two-way systolic gossip algorithms in 2-dimensional grids whose complexity (the number of rounds) meets the trivial lower bound (the sum of the sizes of its dimensions).

The second result (ii) shows that we can systolically gossip in grids with the general gossip complexity, i.e., we do not need to pay for the systolization in grids. Opposed to this, the first result (i) shows that the complexity of systolic gossip in cycles is greater than the gossip complexity in cycles for any systolic period length independent of the size of the cycle.

⋆ This work was partially supported by grants 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).

⋆⋆ Supported by SAV Grant 2/1138/94 of the Computer Science Institute of the Slovak Academy of Sciences.

⋆⋆⋆ 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.

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

1 Introduction

One of the main quantitative measures characterizing the effectivity (computational power) of interconnection networks is the complexity of the solution of fundamental communication tasks like broadcast and gossip. The complexity of these communication tasks was intensively studied in the last two decades (for a survey see [HHL88, FL94, HKMP95]). This paper continues in the study of the complexity of very regular communication algorithms called "periodic" algorithms in [LR93a, LR93b, LHL94] and "systolic" (or "traffic-light") algorithms in [HKPUW94, KP94]. The main motivation behind this concept corresponds to the idea of Kung [Ku79] who has introduced so-called "systolic computations" as parallel computations with cheap realization due to a very regular, synchronized periodic behaviour of all processors of the interconnection network during the whole execution of the computation (see also some of the first papers about systolic computations in trees and grids [CC84, CGS83, CGS84, CSW84]). Liestman and Richards [LR93a, LR93b] were the first who considered a very regular form of communication algorithms for broadcast and gossip. This form, later called "periodic" communication algorithms [LHL94, LR94], was based on edge coloring of a given interconnection network and on periodic (cyclic) execution of communications via edges with the same colour in one communication step. This periodic gossip based on graph colouring was implicitly used also in [XL92, XL94] in order to solve a completely practical problem – termination detection in synchronous interconnection networks. A little more general concept of regular communication was given by Hromkovič et. al. [HKPUW94] who have introduced k-systolic communication algorithms as a repetition of a given sequence of k communication steps (for some $k \in \mathbb{N}$, k independent of the number of processors of the network). This kind of communication was also considered in [KP94] for the problem of route scheduling under the two-way (telephone) model of communication. Since there is no difference in the complexity of general broadcast algorithms and systolic broadcast algorithms (every broadcast algorithm can be systolized without any increase in the number of communication steps [HKPUW94]), the main research problems formulated in [HKPUW94] are the following two:

1. Which is the k-systolic gossip complexity of fundamental interconnection networks? Which is the most appropriate k (if it exists) for a given network?
2. How much must be paid for the systolization of communication algorithms in concrete interconnection networks (i.e., which is the difference between the complexity of gossip and the complexity of systolic gossip)?

These problems were investigated for one-way (telegraph) and two-way (telephone) communication modes in [HKPUW94], where optimal gossip algorithms for paths and complete d-ary trees are presented. While the optimal two-way gossip algorithm for the paths of odd length is 2-systolic, the complexity of k-systolic one-way gossip in paths is always greater than the complexity of $(k+1)$-systolic one-way gossip in paths for any $k \in \mathbb{N}$. A little surprising result of [HKPUW94] is that we do not need to pay for the systolization in k-ary complete trees. More precisely, for every complete k-ary tree there exists a constant $c_k \in \mathbb{N}$ (depending on k and on the communication mode but independent of the depth of the tree) such that the complexity of c_k-systolic two-way (one-way) gossip is the same as the complexity of two-way (one-way) gossip.

An optimal algorithm for non-systolic gossiping in the two-dimensional grid in the two-way mode of communication was presented in [FP80] matching the diameter lower bound. In [LR93a, KP94], two-way systolic algorithms were described taking only $O(1)$ additional rounds.

Here, we continue in the investigation of the problems 1 and 2 for cycles and for 2-dimensional grids. The organization and the contributions of this paper are as follows. The next section provides the fundamental definitions and notations. The complexity of one-way gossip in the cycle C_n of n nodes is studied in Section 3. It is proved that with growing k, the complexity of k-systolic one-way gossip in C_n decreases and tends to $n/2$. Optimal systolic one-way and two-way gossip algorithms in two-dimensional grids are presented in Section 4. The complexity of these optimal algorithms meet the trivial lower bound given by the diameter (sum of the sizes of the dimensions) i.e., we have a little surprising result (especially for one-way communication mode) that we do not need to pay for the systolization in grids. This improves the result of [LR93a, KP94], where a weaker result with $O(1)$ additional rounds in the more powerful two-way mode is established.

2 Definitions

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 $G_i = (V_i, E_i), |V_i| \geq |V_j|$ for $i > j$ is called a class of interconnection networks. Here, we consider the class of cycles $\{C_n\}_{n=1}^{\infty}$ (C_n is the cycle of n nodes) and d-dimensional grids $M(n_1, n_2, \ldots, n_d)$ of the size $n_1 \times n_2 \times \cdots n_d$ for $d, n_1, n_2, \ldots, n_d \in \mathbb{N}$.

The following three fundamental communication problems in networks are defined as follows.

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.

Now, we define the systolic communication as introduced in [HKPUW94].

Definition 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. 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

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

$$\mathrm{minb}(G) = \min \{b_v(G) | v \in V\} \text{ as the min-broadcast complexity of } G.$$

The one-way gossip complexity of G is

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

and the two-way gossip complexity of G is

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

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 $a_v(G) = b_v(G)$ for any work G and any node v of G (cf. [HKMP95]).

Now, we give the notation for the complexity of systolic gossip. Note that we do not need to define the complexity of k-systolic broadcast and accumulation because it does not differ from the complexity of broadcast and accumulation [HKPUW94].

Definition 3. Let $G = (V, E)$ be a network, and let k be a positive integer. The one-way k-systolic gossip complexity of G is

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

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

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

□

3 Systolic Gossip in Cycles

The aim of this section is to investigate the complexity of systolic gossip in cycles in one-way and two-way modes. Similarly as in the non-systolic case [KCV92, HJM94], it is much easier to find optimal systolic two-way gossip algorithms than to get very closed lower and upper bounds on the complexity of one-way systolic gossip in cycles.

We start with the simple two-way case. The next assertion is based on the only fact that the optimal two-way gossip algorithm in cycles can be seen as a systolic one for any even period length.

Theorem 4. *For any even $n \geq 4$, and any even $k \geq 2$:*

$$[k]\text{-}sr_2(C_n) = n/2.$$

If n is odd, the situation is a little more complicated. But also in this case, already period length 4 is enough to get an optimal systolic two-way gossip algorithm for C_n with odd n. We omit the proof of the following assertion in this extended abstract. For the statement of the next theorem, let P_n denote the path of n nodes.

Theorem 5. *For any odd $n \geq 3$:*

$$\begin{aligned}
[2]\text{-}sr_2(C_n) &= n = r_2(P_n),\\
[3]\text{-}sr_2(C_n) &\leq \lceil (2n)/3 \rceil + 2,\\
[4]\text{-}sr_2(C_n) &= \lceil n/2 \rceil + 1 = r_2(C_n).
\end{aligned}$$

We call attention to the fact that odd period lengths are not appropriate for systolic gossip in cycles. We can prove for odd k that $[k]\text{-}sr_2(C_n)$ does not match the trivial lower bound $\lceil n/2 \rceil$ given by the diameter as in the case of even k. But we omit the formulation of these lower bounds in this extended abstract.

Now, we give the results for the one-way case which presents the main contribution of this section. The proof is too technical to be completely presented in this extended abstract.

Theorem 6. *For any $k \geq 6$ and any $n \geq 2$:*

$$\frac{n}{2} \cdot \frac{k+2}{k} - 2 \ \leq [k]\text{-}sr(C_n) \ \leq \ \frac{n}{2} \cdot \frac{k+2}{k} + 2k - 2 \quad \textit{for } k \textit{ even},$$
$$\frac{n}{2} \cdot \frac{k+2}{k} - 2 \ \leq [k]\text{-}sr(C_n) \ \leq \ \frac{n}{2} \cdot \frac{k+1}{k-1} + 2k - 2 \quad \textit{for } k \textit{ odd}.$$

Idea of the Proof. The upper bounds are based on some analysis showing that the best strategy for one-way systolic gossiping in cycles is to totally prefer the information flow in one direction over the flow in the opposite direction. The lower bound is based on a careful analysis of the number of "collisions" between messages flowing in opposite directions in the cycle. □

Finally, we observe that there are some similarities in the results for paths [HKPUW94] and in the results stated above. For both path and cycle we do not need to pay for the systolization in the two-way case, but we have to pay with $const \cdot n$ additional rounds for any period length k in the one-way case.

4 Systolic Gossiping in Two-Dimensional Grids

The main result of this section is an optimal systolic algorithm for the gossip problem in two-dimensional grids. In [KP94], an algorithm was presented for the two–way mode of communication which completes gossiping on an $(m_1 \times m_2)$–grid M (i.e. a grid consisting of $m_1 + 1$ columns and $m_2 + 1$ rows) in $d + O(1)$ rounds, where $d = m_1 + m_2$ is the diameter of the grid. We derive an algorithm for the one–way mode of communication that performs gossiping on M in d rounds, provided m_1 and m_2 are large enough. The period–length of the presented algorithm is 16. Since the diameter is the trivial lower bound on the number of required rounds for any gossip algorithm (even for a non–systolic one in two–way mode of communication), the optimality of our algorithm is obvious. The fact itself that we do not need to pay for the systolization in two–dimensional grids is of special interest, because grids are typical parallel architectures. Note that any solution in the one–way mode of communication is also a solution for the two–way mode. Thus, the possible advantages of using the two–way model solely lie in shorter periods and a weaker restriction on the size of the grids. We have some results of this kind for the two–way mode, but we omit their explicit formulation in this extended abstract.

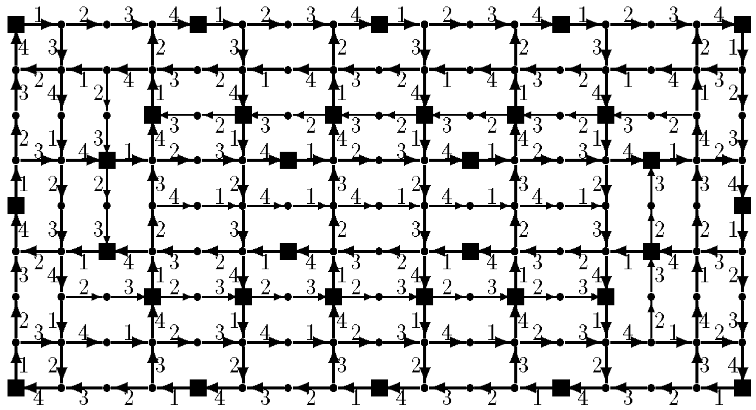

Fig. 1. One–way gossiping in a (16×8)–grid (near–optimal scheme)

Theorem 7. *For any $(m_1 \times m_2)$–grid M with m_i even, and $m_i \geq 28$,*

$$[16] - sr(M) \;=\; m_1 + m_2 \;.$$

Idea of the Proof. The proof is realized as follows. First, one designs a near–optimal algorithm S_1 that requires $d + 16$ rounds for gossiping in any 2–dimensional grid. Then, one extends this solution resulting in an optimal algorithm S_2.

To design the near–optimal algorithm we employ the following strategy. We say that a subset K of the grid–nodes forms a *knowledge set* after round t if any information of the system is known to at least one node of K. We specify such a subset K and show that according to our communication scheme K forms a knowledge set after 8 rounds. After additional d rounds any individual node of K knows the cumulative message, i.e. in these d rounds a gossip restricted to K is performed. Thereafter, additional 8 rounds suffice to broadcast the cumulative message to any node of the grid. We can observe that for special values of m_1 and m_2 every fourth node of the border including all corners can be taken into K (see Fig. 1 for a near–optimal gossiping in a (16×8)–grid).

To obtain an optimal gossip algorithm we split the grid into a center and a seam as shown in Figure 2.

In the center we use the communication pattern of the near–optimal algorithm. Assume that K_{GC} is the subset of center–nodes among which gossiping in d_{GC} rounds is possible, where d_{GC} denotes the diameter of the center. If we are able to make K_{GC} a knowledge set for the entire grid in $(d - d_{GC})/2$ rounds, and if broadcasting from K_{GC} to the entire grid can be performed within the same number of rounds, we obtain an optimal algorithm using only d rounds of communication. The communication pattern for the seam consequently is

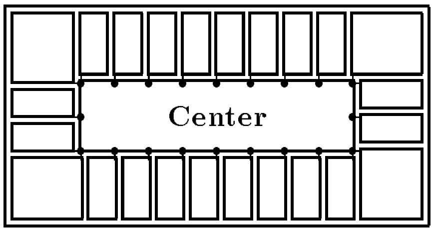

Fig. 2. Decomposition of the grid

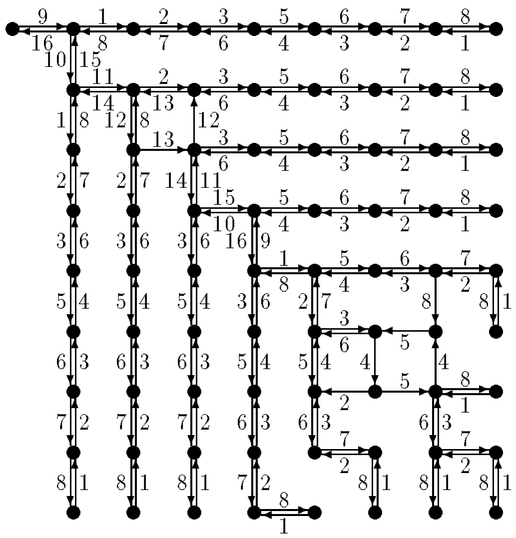

Fig. 3. Accumulation and broadcast in a corner of size 8×9 (optimal scheme)

specialized in the accumulation and broadcast task. Fig. 3 gives an example of an optimal scheme for accumulation and broadcast in a corner of the seam. □

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.

[FL94] P. Fraigniaud, E. Lazard, *Methods and problems of communication in usual networks.* Discrete Applied Mathematics, special issue on Broadcasting and Gossiping (DAM-BG), Vol. 53 Nos. 1-3, September 1994, pp. 79-134.

[FP80] A.M. Farley, A. Proskurowski, *Gossiping in grid graphs.* J. Combin. Inform. System Sci. 5 (1980), pp. 161-172.

[HJM94] J. Hromkovič, C.-D. Jeschke, B. Monien: *Note on Optimal Gossiping in some Weak-Connected Graphs.* Theoretical Computer Science 127 (1994), No. 2, 395-402.

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

[HKMP95] 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, 1994, to appear.

[HKPUW94] J. Hromkovič, R. Klasing, D. Pardubská, W. Unger, H. Wagener: *The complexity of systolic dissemination of information in interconnection networks.* In: Proc. of the 1st Canada-France Conference on Parallel Computing (CFCP '94), Springer LNCS 805, pp. 235-249. R.A.I.R.O. Theoretical Informatics and Applications, Vol. 28 Nos. 3-4, 1994, pp. 303-342.

[KCV92] D.W. Krumme, G. Cybenko, K.N. Venkataraman: *Gossiping in minimal time.* SIAM J. Comput. 21 (1992), pp. 111-139.

[KP94] G. Kortsarz, D. Peleg: *Traffic-light scheduling on the grid.* Discrete Applied Mathematics, special issue on Broadcasting and Gossiping (DAM-BG), Vol. 53 Nos. 1-3, September 1994, pp. 211-234.

[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.

[LHL94] R. Labahn, S.T. Hedetniemi, R. Laskar, *Periodic gossiping on trees.* Discrete Applied Mathematics, special issue on Broadcasting and Gossiping (DAM-BG), Vol. 53 Nos. 1-3, September 1994, pp. 235-246.

[LR93a] A.L. Liestman, D. Richards: *Network communication in edge-colored graphs: Gossiping.* IEEE Trans. Par. Distr. Syst. 4 (1993), 438-445.

[LR93b] A.L. Liestman, D. Richards: *Perpetual gossiping.* Parallel Processing Letters 3 (1993), No. 4, 347-355.

[LR94] R. Labahn, A. Raspaud: *Periodic gossiping in back-to-back trees.* Technical Report, Universität Rostock, Germany, 1994.

[XL92] Cheng-Zhong Xu, Francis C.M. Lau: *Distributed termination detection of loosely synchronized computations.* In: Proc. 4th IEEE Symp. on Parallel and Distr. Processing, Texas 1992, pp. 196-203.

[XL94] Cheng-Zhong Xu, Francis C.M. Lau: *Efficient distributed termination detection for synchronous computations in multicomputers.* Technical Report TR-94-04, Department of Computer Science, The University of Hong Kong, March 1994.

This article was processed using the LaTeX macro package with LLNCS style