A linear algorithm for the distance in Cayley graph $H_{\ell,p}$

Caroline Patrão ⓘ   Luis Kowada ⓘ
Diane Castonguay   André Ribeiro
Celina de Figueiredo ⓘ

Abstract

The family of graphs $H_{\ell,p}$ has been defined in the context of edge partitions. The established properties such as vertex transitivity and low diameter suggest this family as a good topology for the design of interconnection networks. The vertices of the graph $H_{\ell,p}$ are the ℓ-tuples with values between 0 and $p-1$, such that the sum of the ℓ values is a multiple of p, and there is an edge between two vertices, if the two corresponding tuples have two pairs of entries whose values differ by one unit or $p-1$. As the diameter of the graph $H_{\ell,p}$ is $\Theta(\ell \cdot p)$, then any algorithm to construct a diametral path between two vertices needs $\Omega(\ell \cdot p)$ steps, however in this work we show an algorithm that finds the distance between a pair of vertices of the graph $H_{\ell,p}$ in $\Theta(\ell \cdot \log p)$, which is the input size in bits, and therefore our algorithm has optimal asymptotic complexity.

1   Introduction

In this work, we are motivated by the design and analysis of static networks. Static networks can be modeled using tools from Graph Theory. A graph represents an interconnection network, where the processors are the vertices and the communication links between processors are the edges connecting the vertices. There are several parameters of interest to specify a network: low degree, low diameter, and the distribution of the node disjoint paths between a pair of vertices in the graph. The degree relates to the port capacity of the processors and hence to the hardware cost of the network. The maximum communication delay between a pair of processors in a network is measured by the diameter of the graph. Thus, the diameter is a measure of the running cost. The number of parallel paths between a pair of nodes is limited by the degree of the underlying graph, the knowledge of this distribution is helpful in the evaluation of the fault-tolerance of the network [4, 9]. Cayley graphs are connected vertex-transitive graphs that provide two challenging and extensively studied problems: obtain a Hamiltonian cycle and determine a graph diameter [8]. Cayley graphs are regular, in some cases have logarithmic diameter, and can be used to design interconnection networks [9, 8, 5]. In [6], Ribeiro, Figueiredo and Kowada showed that not only the graphs $H_{\ell,p}$ are Cayley graphs but also are Hamiltonian. Some families of Cayley graphs are: The multidimensional torus graph T_k^n, the hypercube graph H_n, and the graphs $H'_{\ell,p}$ [7] obtained by the removal of some elements of the generating set of the graphs $H_{\ell,p}$. The hypercube graph is considered a particular case of the torus, that is, $H_n = T_n^2$ and the hypercube graph is isomorphic to the graph $H'_{n+1,2}$.

This paper is organized as follows: In Section 2, the definition of Cayley graphs is introduced to explain the concept of abstract groups, which are described by a generating set and we define the graph $H_{\ell,p}$. In Section 3, our goal is to propose an algorithm with linear complexity of the input size in bits $\Theta(\ell \cdot \log p)$ that finds the distance between a pair of vertices of

the graph $H_{\ell,p}$, although the diameter of $H_{\ell,p}$ is $\Theta(\ell \cdot p)$.

2 The Cayley Graph $H_{\ell,p}$

In this section we define some terms that are used in Section 3. Let $(G,+)$ be a group and $C \subseteq G$ be a *generating set* of G such that any element of G can be obtained from elements of C by a finite number of applications of the operation $+$. A directed graph $\Gamma = (V, E)$ is a *Cayley graph* for a group $(G,+)$ with a generating set C, if there is a bijection mapping each $x \in V$ to an element $g_x \in G$, such that xy is a directed edge of E if and only if there exists $c \in C$ such that $g_y = c + g_x$.

If the identity element $\iota \notin C$, then there are no loops in Γ, and Γ satisfies the *identity free* property. If $g \in G$ then there is a unique $g' \in G$ such that $g+g' = \iota$, denote g' by $-g$ and in this case we define $g_x - g_y = g_x + (-g_y)$. If $c \in C$ implies $-c \in C$, then for every edge from g to $g + c$, there is also an edge from $g + c$ to $(g + c) + (-c) = g$, and Γ satisfies the *symmetry condition*. A Cayley graph that satisfies both the identity free property and the symmetry condition is an undirected graph.

For each $\ell \geq 3$ and $p \geq 3$, Holyer [3] defines a graph $H_{\ell,p} = (V_{\ell,p}, E_{\ell,p})$ where

$V_{\ell,p} = \{x = (x_1, ..., x_\ell),$ with $x_k \in \mathbb{Z}_p$ and $\sum\limits_{k=1}^{\ell} x_k \equiv_p 0\}$,

$E_{\ell,p} = \{xy :$ there are distinct i,j such that $y_k \equiv_p x_k$ for $k \neq i,j$ and $y_i \equiv_p x_i + 1,\ y_j \equiv_p x_j - 1\}$.

The vertices of $H_{\ell,p}$ are the elements of a finite group $(V_{\ell,p}, +)$, where the operation $+$ is such that $x + y = (x_1, \ldots, x_\ell) + (y_1, \ldots, y_\ell) = (x_1 + y_1, \ldots, x_\ell + y_\ell)$, where $x_k + y_k$ is the operation in $(\mathbb{Z}_p, +)$, for $x, y \in V_{\ell,p}$ and the element $(0, 0, \ldots, 0, 0) \in V_{\ell,p}$ is the *identity vertex* ι. The graphs $H_{\ell,p}$ considered in this paper are undirected Cayley graphs associated to the group $(V_{\ell,p}, +)$ with generating set $C_{\ell,p}$ [6], defined as the set of ℓ-tuples $e_{i,j} = (c_1, \ldots, c_\ell)$ such that, for $i,j \in \{1, \ldots \ell\}$ with $i \neq j$, we have that $c_k = 1$ when $k = i$, $c_k = p - 1$ when $k = j$, and otherwise $c_k = 0$.

Figure 1: A drawing of $H_{3,4}$, where the 16 vertices are highlighted. This graph is not a planar graph because it is regular of degree 6, but the drawing presented is a representation on the torus without crossings of edges.

3   Distance between a pair of vertices in the graph $H_{\ell,p}$

The *distance* $d(u, v)$ between the vertices u and v in a graph is the number of edges in a shortest path connecting them. For sake of simplicity, the distance between the vertices x and ι is denoted by $d(x)$.

Cayley graphs are vertex-transitive [4], that is, given two vertices z, w in a Cayley graph G, there is some automorphism $\pi : V(G) \to V(G)$ such that $\pi(z) = w$. Therefore, the problem of finding the distance between two vertices $u, v \in V_{\ell,p}$ can be reduced to finding the distance between a vertex $x \in V_{\ell,p}$ and the identity vertex ι, thus $d(u, v) = d(x)$. In fact, given two vertices $z, v \in V_{\ell,p}$, there is an automorphism $\pi : V_{\ell,p} \to V_{\ell,p}$ such that $\pi(z) = z - v$, for each $z \in V_{\ell,p}$. In particular, $\pi(u) = u - v$ and $\pi(v) = \iota$. Thus, we have that $d(u, v) = d(x)$, where $x = u - v$.

In Lemma 3.1, denote by $S_k(x) = \sum\limits_{u=1}^{k} x_u$ for $1 \leq k \leq \ell$, and by $s = \sum\limits_{u=1}^{\ell} x_u$. Denote by x' the permutation of the components of x such that

the components in x' are in ascending order, i.e., $x'_k \leqslant x'_{k+1}$.

Lemma 3.1. *[7] The distance from a vertex $x \in V_{\ell,p}$ to the identity vertex ι is $d(x) = S_b(x')$, where $b = \ell - \frac{s}{p}$.*

Note that to compute the distance $d(x)$ in Lemma 3.1 we can sort the components of a vertex x. Using Radix Sort algorithm it is possible to sort the components of a vertex x with $\Theta(\ell \cdot \log^2 p)$ operations, because to sort each of the $\Theta(\log p)$ digits it is necessary to compute the $\ell \cdot \log p$ bits from the list.

We want to compute the distance between any two vertices with a better complexity, without sorting the ℓ components. For this, we can apply the Select algorithm in Section 9.3 of book [2], which is used to find the b^{th} smallest element in an input arrangement with distinct elements in linear order of the input size.

By Lemma 3.1 we can calculate the distance between any two vertices $u, v \in V_{\ell,p}$ by mapping u into $x = u - v$ and then calculating the distance $d(x)$, as shown in Algorithm 3.1.

Algorithm 3.1: Distance $d(u, v)$

Input : $\ell, p, u \in V_{\ell,p}, v \in V_{\ell,p}$

Output: $d(u, v)$

1 **begin**

2   **for** $k = 1$ *to* ℓ **do**

3          $x_k \leftarrow (u_k - v_k) \mod p$

4      $s \leftarrow \sum\limits_{k=1}^{\ell} x_k$

5      $b \leftarrow \ell - s/p$

6      $x'_b \leftarrow Select(x, b)$

7      $c \leftarrow 0, d \leftarrow 0$

8   **for** $k = 1$ *to* ℓ **do**

9   **if** $x_k < x'_b$ **then**

10              $d \leftarrow d + x_k$ and $c \leftarrow c + 1$

11      $d \leftarrow d + (b - c) \cdot x'_b$

12   **return** (d)

Theorem 3.2. *Algorithm* 3.1 *calculates the distance* $d(u, v)$ *of any two vertices* $u, v \in V_{\ell,p}$ *in time* $\Theta(\ell \cdot \log p)$.

Proof. The algorithm Distance $d(u, v)$ is described as follows: at first, in lines 2 and 3, we map a vertex u into a suitable vertex x to calculate the distance from x to ι, as this distance is the distance from u to v by the vertex-transitivity of the graph $H_{\ell,p}$. The map can be computed in $\Theta(\ell \cdot \log p)$ bits. Now, we sum the ℓ components of the vertex x (line 4), we find $b = \ell - \frac{s}{p}$ (line 5) and we calculate x'_b by the Select algorithm (line 6) , which also runs in $\Theta(\ell \cdot \log p)$ bits. Although our input may contain repeated elements, we can use a modification in the Select algorithm considering as pivot of the partition the first element equal to the median of the medians, thus the Select algorithm in Algorithm 3.1 determines the element that is in position b, denoted by x'_b, where all its predecessors are less or equal to x'_b and all its successors are greater than x'_b. In the next step, we sum the

components that are smaller than x'_b denoted by d and count the number of sums c. Finally, we ensure that b components are entered by adding the value x'_b until we have b components and we return the distance d, which is also done in $\Theta(\ell \cdot \log p)$. Note that, we have that $d = \sum_{k=1}^{b} x_k$, so by Lemma 3.1, d is the distance $d(x)$ and by the vertex-transitivity of the graph $H_{\ell,p}$ we have that $d(x) = d(u,v)$, so $d = d(u,v)$ and the algorithm Distance runs in linear order on the input size, i.e., $\Theta(\ell \cdot \log p)$, therefore an optimal asymptotic complexity. $\qquad\square$

It is important to note that finding the distance between two vertices in $\Theta(\ell \cdot \log p)$ is tight, especially when analyzing the solution of the diameter problem for the graph $H_{\ell,p}$.

As the graph $H_{\ell,p}$ has $p^{\ell-1}$ vertices, Castonguay, Ribeiro, Figueiredo and Kowada[1] highlighted some special vertices to study the diameter of the graph $H_{\ell,p}$ and obtained an upper bound for this diameter to be $\left\lfloor \frac{\ell \cdot p}{4} \right\rfloor$. The diameter can be found by the greatest value of the distance from each special vertex to the identity vertex ι. For the sake of completeness, we rewrite the results presented in [1, Lemma 11] as $D(H_{\ell,p}) = \max_t\{(\ell - t) \cdot \left\lfloor \frac{t \cdot p}{\ell} \right\rfloor\}$, for $0 \le t < \ell$.

Note that a lower bound for the diameter of $H_{\ell,p}$ is found when $t = \left\lfloor \frac{\ell}{2} \right\rfloor$, thus the solution of the diameter problem for the graph $H_{\ell,p}$ is $\Theta(\ell \cdot p)$. So, any algorithm that finds a path between two vertices needs $\Omega(\ell \cdot p)$ steps, which are the vertices in the path that are displayed by the algorithm. However, Algorithm 3.1 finds the distance between any two vertices u and v of the graph $H_{\ell,p}$ in $\Theta(\ell \cdot \log p)$, which is the input size in bits, and therefore has optimal asymptotic complexity.

4 Conclusion

Several authors observed that Cayley graphs provide a useful and unified framework for the design of interconnection networks for parallel computers. For the family of graphs $H_{\ell,p}$, we propose an algorithm that finds the

distance between any two vertices with linear complexity on the input size $\Theta(\ell \cdot \log p)$, whereas the diameter is $\Theta(\ell \cdot p)$. This established property makes the Cayley graph $H_{\ell,p}$ a good scheme of interconnection networks.

References

[1] D. Castonguay and A.C. Ribeiro and C.M.H. Figueiredo and L.A.B. Kowada On the diameter of the Cayley graph $H_{\ell,p}$. *Matematica Contemporanea* (2015), vol. 44, pp. 1–10.

[2] T.H. Cormen and C.E. Leiserson and K.L. Rivest and C. Stein *Introduction to Algorithms*. Second edition. Elsevier,(2002). Section 9.3.

[3] I. Holyer   The NP-completeness of some edge-partition problems. *SIAM Journal Computing* (1981), vol. 10, pp. 713–717.

[4] E. Konstantinova. Some problems on Cayley graphs. *Linear Algebra and its Applications* (2008), vol. 429, pp. 2754–2769.

[5] S. Lakshmivarahan and J.S. Jwo and S.K. Dhall. Symmetry in interconnection networks based on Cayley graphs of permutation groups: A survey. *Parallel Computing* (1993), vol. 19, pp. 361–407.

[6] A.C. Ribeiro and C.M.H. Figueiredo and L.A.B. Kowada. An evidence for Lovász conjecture about Hamiltonian paths and cycles. *Matematica Contemporanea* (2010), vol. 39, pp. 121–128.

[7] A.C. Ribeiro and C.M.H. Figueiredo and L.A.B. Kowada Two families of Cayley graph interconnection networks. *Matematica Contemporanea* (2014), vol. 42, pp. 105–114.

[8] C.S. Reis and L.A.B. Kowada and L.R. Bueno and A.C. Ribeiro and C.M.H. Figueiredo. Hamiltonian cycles in unitary prefix transposition rearrangement graphs. *Discrete Applied Mathematics* (2015), vol. 192, pp. 82–86.

[9] P. Vadapalli and P.K. Srimani. A new family of Cayley graph inter-
connection networks of constant degree Four. *IEEE Transactions on
Parallel and Distributed Systems* (1996), vol. 7, pp. 26–32.

Caroline Patrão   Luis Kowada
COPPE UFRJ, IFG/IFRJ   IC UFF
Rio de Janeiro, Brazil   Rio de Janeiro, Brazil
caroline.patrao@ifrj.edu.br   luis@ic.uff.br

Diane Castonguay   André Ribeiro
INF UFG   IFGOIANO
Goiás, Brazil   Goiás, Brazil
diane@inf.ufg.br   andre.cunha@ifgoiano.edu.br

Celina de Figueiredo
COPPE UFRJ
Rio de Janeiro, Brazil
celina@cos.ufrj.br