Complexity of the Undirected Radius and Diameter Problems for Succinctly Represented Graphs

Till Tantau
Institut für Theoretische Informatik
Universität zu Lübeck
D-23538 Lübeck, Germany

Technical Report
SIIM-TR-A-08-03
Schriftenreihe der Institute für Informatik/Mathematik der
Universität zu Lübeck

Abstract

The diameter of an undirected graph is the minimal number d such that there is a path between any two vertices of the graph of length at most d. The radius of a graph is the minimal number r such that there exists a vertex in the graph from which all other vertices can be reached in at most r steps. In the present paper we study the computational complexity of deciding whether a given graph has some fixed small diameter or radius. For graphs given as adjacency matrices the problem is trivial, but for graph that are represented succinctly using circuits, the complexity is more interesting: We show that for every fixed $d \geq 2$ the problem of deciding whether a succinctly represented undirected graph has diameter at most d is Π_2^p-complete, while for $d = 1$ the problem is Π_1^p-complete; and for every fixed $r \geq 2$ the problem of deciding whether a succinctly represented undirected graph has radius at most r is Σ_3^p-complete, while for $r = 1$ the problem is Σ_2^p-complete. This shows that these problems are fairly natural problems that are complete for higher levels of the polynomial hierarchy.

1 Introduction

In the paper *On the Complexity of Kings* [1] Hemaspaandra et al. showed that the succinct diameter problem for directed graphs (and even for tournaments) is Π_2^p-complete for all fixed $d \geq 2$ and that the succinct radius problem for directed graphs is Σ_3^p-complete for all fixed $r \geq 2$. The purpose of the present paper is to show that these results also hold for undirected graphs.

This is not quite obvious since the main proof that the succinct k-diameter problem for directed graphs is Π_2^p-complete uses tournament graphs, which are inherently directed. The closest undirected analogues to such graphs are cliques and the diameter of a clique is, trivially, always 1. Also, for directed graphs the diameter and the radius of a graph are only loosely related (a directed graph can have infinite diameter and finite radius), while for undirected graphs we always have $r \leq d \leq 2r$. This makes it especially intriguing that the complexities of these problems differ by a whole level of the polynomial hierarchy.

The present paper will focus on just giving the completeness proofs, for a detailed background on these problems please see [1].

2 Basic Definitions and Notations

Throughout this paper $\Sigma = \{0,1\}$. We refer to elements of $\{0,1\}^* = \Sigma^*$ as bitstrings.

An *(undirected) graph* is a pair (V, E) consisting of a nonempty vertex set V together with an edge set $E \subseteq \{\{u,v\} \mid u,v \in V, u \neq v\}$. A *path of length* l in a graph is sequence v_0, v_1, …, v_l of vertices such that $\{v_{i-1}, v_i\} \in E$ holds for all $i \in \{1,\dots,l\}$. We denote the fact that the sequence forms a path by $v_0 \to v_1 \to \cdots \to v_l$, but note that, nevertheless, edges are always undirected in the present paper. The *diameter* of a graph is the smallest number d such that for every pair $u, v \in V$ of vertices there is a path from u to v of length at most d. If a graph has more than one strongly connected component, its diameter is ∞. The *radius* of a graph is the smallest number r such that there exists a vertex c, called *a k-center* of the graph, from which there are paths of length at most r to all other vertices. (A *k-center* is also known as a *k-king* in the literature, but that terminology is less natural for the present paper.) It is possible for a graph to have a radius of ∞. The radius and the diameter of a graph are related by the inequalities $r \leq d \leq 2r$.

A language L is in Π_2^{p} if and only if there exist a polynomial p and a polynomial-time decidable relation $R \subseteq \Sigma^* \times \Sigma^* \times \Sigma^*$ such that for all words $x \in \Sigma^*$ we have

$$x \in L \iff (\forall y \in \Sigma^{p(|x|)})(\exists z \in \Sigma^{p(|x|)})[R(x,y,z)]. \tag{1}$$

Similarly, $L \in \Sigma_2^p$ means that there exists an R with the property

$$x \in L \iff (\exists y \in \Sigma^{p(|x|)})(\forall z \in \Sigma^{p(|x|)})[R(x,y,z)]. \tag{2}$$

A language L is in Σ_3^{p} if and only if there exist a polynomial p and a polynomial-time decidable relation $R \subseteq \Sigma^* \times \Sigma^* \times \Sigma^* \times \Sigma^*$ such that for all words $x \in \Sigma^*$ we have

$$x \in L \iff (\exists w \in \Sigma^{p(|x|)})(\forall y \in \Sigma^{p(|x|)})(\exists z \in \Sigma^{p(|x|)})[R(x,w,y,z)]. \tag{3}$$

By *circuit* we refer to combinatorial circuits containing input-, output-, negation-, and-, and or-gates. The fan-in of each gate is at most 2. Fan-out is not restricted. For a circuit C with n input gates and m output gates, we also use $C\colon \{0,1\}^n \to \{0,1\}^m$ to denote the function computed by the circuit C. Let $\mathrm{code}(C)$ denote some standard binary encoding of the circuit C.

We use circuits to define undirected graphs succinctly as follows. For positive integers n, given an $2n$-input, 1-output circuit C, we say that it *specifies the graph* G whose vertex set is $V = \Sigma^n$ and whose edge set is defined as follows: There is an edge $\{x,y\} \in E$ iff $C(xy) = 1 \lor C(yx) = 1$. We say that C is a *succinct representation of* G.

We next formalize the radius and diameter problems for succinctly specified graphs. Let k be a fixed positive integer.

SUCCINCT-UNDIRECTED-k-RADIUS

$= \{\mathrm{code}(C) \mid \text{the undirected graph specified by } C \text{ has radius at most } k\}$.

SUCCINCT-UNDIRECTED-k-DIAMETER

$= \{\mathrm{code}(C) \mid \text{the undirected graph specified by } C \text{ has diameter at most } k\}$.

3 The Diameter Problem

For $k = 1$, the SUCCINCT-UNDIRECTED-1-DIAMETER problem is the same as asking whether the graph specified by a circuit is a clique. This, in turn, is the same question as asking whether the encoding circuit always outputs 1 either for xy or for yx. This, finally, is clearly a coNP-complete problem.

Theorem 3.1. SUCCINCT-UNDIRECTED-1-DIAMETER *is complete for* $\Pi_1^p = \text{coNP}$.

The more challenging cases are $k \geq 2$.

Theorem 3.2. *For every* $k \geq 2$*, the problem* SUCCINCT-UNDIRECTED-k-DIAMETER *is complete for* Π_2^p.

Proof. Clearly, SUCCINCT-UNDIRECTED-k-DIAMETER $\in \Pi_2^p$ since on input of a circuit C we only have to check whether *for all* $x, y \in \Sigma^n$ *there exist* $l \leq k-1$ intermediate vertices $v_1, \dots, v_l \in \Sigma^n$ such that $x \to v_1 \to \dots \to v_l \to y$.

To prove completeness, let any problem $L \in \Pi_2^p$ be given and let L be characterized by the relation R as in equation (1). Let us write m for $p(|x|)$. We present a reduction that maps any word $x \in \Sigma^*$ to a circuit C such that $x \in L$ iff the undirected graph G specified by the circuit C has diameter at most k. As a matter of fact, if $x \in L$, then the graph will have diameter 2, and if $x \notin L$, then the graph will have diameter ∞. This means that the same reduction will work for all $k \geq 2$.

We start with the core idea of the construction of the graph, the idea will need to be modified slightly later on. The graph is a split graph, which means that it consists of an independent set I and a clique K with some additional edges between these two sets. The independent set I is just Σ^m and contains 2^m many vertices, the clique is the set $\Sigma^m \times \Sigma^m$ and contains $2^m \cdot 2^m$ many vertices. There is an edge between a vertex $y \in I$ and a vertex $(z_1, z_2) \in K$ iff $R(x, y, z_1) \vee R(x, y, z_2)$.

We claim that the resulting graph has diameter 2 for $x \in L$ and diameter ∞ for $x \notin L$. To see this, first assume $x \notin L$. Then there exists a $y \in \Sigma^m$ such that for all $z \in \Sigma^m$ we have $\neg R(x, y, z)$. But, then, the vertex y in the independent set I is isolated and, thus, the diameter is infinite. Now, second, assume $x \in L$. Then every vertex $y \in I$ is connected to at least one vertex in the clique, which means that from any $y \in I$ we can reach all vertices in the clique in at most two steps (one step to reach the clique and another step to reach the desired vertex in the clique). Vice versa, from the clique we can reach all other vertices in the clique in one step and also all $y \in I$ in, possibly, one more step. The most interesting, final, case are two vertices $y_1, y_2 \in I$. For them there exist witnesses $z_1, z_2 \in \Sigma^m$ such that both $R(x, y_1, z_1)$ and $R(x, y_2, z_2)$ hold. Then $y_1 \to (z_1, z_2) \to y_2$ is a path.

It remains to argue that we can encode the graph succinctly as a (small) circuit. Due to the fact that the graph is highly structured, this is not difficult, in principle. However, a problem arises because the size of the vertex set is not a power of 2 and this problem can *not* be fixed by just adding a sufficient number of independent vertices (they would always cause the diameter to become infinite). Instead, we take an arbitrary vertex and expand it to a clique whose size is chosen in such a way that the total number of vertices becomes a power of 2. This will not change the diameter of the graph. □

4 The Radius Problem

We start with $k = 1$.

Theorem 4.1. SUCCINCT-UNDIRECTED-1-RADIUS *is complete for* Σ_p^2.

Proof. The problem is in Σ_p^2 because we must check whether *there exists* a vertex such that *all vertices* are directly connected to it.

To prove completeness, let $L \in \Sigma_p^2$ be given and let R be a predicate from equation 2. For the reduction, let $x \in \Sigma^*$ be given. We map it to a circuit encoding the following graph: It is, once more, a split graph. This time, both the independent set and the clique have size 2^m.

There is an edge between a vertex z in the independent set and a vertex y in the clique iff $R(x, y, z)$ holds. Clearly, this graph can easily be encoded using a circuit (note that its size is, indeed, a power of 2).

To prove correctness, first let $x \in L$ and let y be a witness for this. Then from the vertex y of the clique we can reach all vertices in the clique in one step and we can also reach all elements of the independent set in one step. Now assume that the radius of the graph is 1 and let c be a center. Then c must be a member of the clique, because from vertices in the independent set we cannot reach other members of the independent set in one step. But, then, there must be an edge between $y = c$ and every member z of the independent set, which means that y is a witness for $x \in L$. $\square$

The more difficult case is $k \geq 2$.

Theorem 4.2. *For every $k \geq 2$, the problem* SUCCINCT-UNDIRECTED-k-RADIUS *is complete for Π_3^p.*

Proof. As before, membership of SUCCINCT-UNDIRECTED-k-RADIUS in Π_3^p is easy to see, so we describe only the reduction from an arbitrary language $L \in \Pi_3^p$. On input $x \in \Sigma^*$ we must construct a circuit encoding a graph G whose radius is at most k iff $x \in L$. This graph will be described in the following.

We start with the description of graphs G^w, one for each possible witness $w \in \Sigma^m$. Each graph starts with a simple path of length $k - 2$ (it is just a single vertex in case $k = 2$). Let us call the first vertex on this path the *potential center* and let us introduce the name p^w for it. Next, G^w contains two independent sets, each of size 2^m. The elements of the first set will be denoted z^w, with $z \in \Sigma^m$, the element of the second set will be denoted y^w, with $y \in \Sigma^m$. These sets are connected as follows: There is an edge between the last vertex of the simple path and every vertex of the first independent set. There is an edge between z^w and y^w iff $R(x, w, y, z)$ holds. Observe that we can reach all vertices in G^w from the potential center p^w in at most k steps iff $(\forall y \in \Sigma^m)(\exists z \in \Sigma^m)[R(x, w, y, z)]$.

The above construction ensures that $x \in L$ holds if, and only if, one of the potential centers is a k-center of its G^w. We will now extend the construction to ensure that (a) only potential centers can be, indeed, be k-centers of the whole graph and (b) that whenever a potential center is a k-center of some G^w, it is automatically a potential center of the whole graph G.

For these extensions, we first take the disjoint union of the graphs G^w. Then, we add one additional simple path of length $k - 1$, let us call it S, and connect all potential centers to one end of this simple path. We claim that the resulting graph, let us call it G', has the following property: Given any w, no vertex in all of G' other than possibly the potential center p^w has the property that all vertices G^w and also in the additional simple path can be reached in k steps. The reason for this is that vertices inside a G^w other than p^w cannot reach the end of the additional simple path in k steps, while vertices on the additional simple path cannot reach any y^w in k steps.

It remains to ensure that from the potential centers we can always reach all vertices in all *other* $G^{w'}$ in at most k steps. For this, we extend G' as follows, resulting in the final graph G: First, we add a clique of size 2^m, called the W-clique in the following, whose elements will be denoted $w \in \Sigma^m$. Each vertex w in this clique is directly connected to the potential center p^w. Additionally, for each vertex w in the clique and for every vertex $v \in G^{w'}$ with $w' \neq w$, we add an additional simple path from w to v whose length is exactly $k - 1$ (the path contains $k - 2$ vertices in addition to w and v). Let us call this path the *tube from w to v*. This means that from w we can now reach all vertices in all *other* $G^{w'}$ in at most $k - 1$ steps.

We claim that G has radius k iff $x \in L$. First, assume $x \in L$. Then there exists a witness $w \in \Sigma^*$ for this. By construction of the graph G, the potential center p^w is a k-center of the graph: We showed already that we can reach all vertices in G^w in k steps from p^w; we can reach all vertices on the path S in k steps; we can reach all vertices of the W-clique in at most 2 steps; we can reach all vertices of all other $G^{w'}$ in k steps via the tube from w to this vertex; and we can reach all vertices in any of the tubes in k steps by investing two steps to get to the start vertex w' of the tube and then investing the remaining $k-2$ steps to reach all vertices inside the tube.

For the other direction, assume that c is a k-center of G. First, it holds that c must be one of the potential centers: As we argued already for the graph G' we can reach the end of the path S only from these potential centers and the addition of the W-clique did not change this. Also, the vertices on the path S themselves also cannot be centers because from them we still cannot reach any y^w; the W-clique does not change this either. Second, if p^w is a k-center of all of G, then it must also be a k-center of G^w. For this, it suffices to argue that no path of length k starting at p^w that contains a vertex outside of G^w can end at any y^w. However, this is easy to see: Any such path would need to go through the W-clique. The shortest path from any member $w' \neq w$ of the W-clique to y^w has length $k-1$ (namely an appropriate tube) and the distance from p^c to any such w' is 2.

Although the construction of the graph is a bit involved, it can clearly the described succinctly using an appropriate circuit. As for the diameter problem, we face the problem that the vertex set is not readily a power of 2, but this can be fixed in the same way by "blowing up" any vertex to a clique until the next power of 2 has been reached. $\square$

References

[1] E. Hemaspaandra, L. Hemaspaandra, T. Tantau, and O. Watanabe. On the complexity of kings. Technical Report URCS-TR905, Department of Computer Science, University of Rochester, Rochester, NY, USA, December 2006. Revised, November 2007.