Solvability of Graph Inequalities

Marcus Schaefer
School of CTI
DePaul University
243 S. Wabash Avenue
Chicago, Illinois 60604, USA
mschaefer@cs.depaul.edu

Daniel Štefankovič
Department of Computer Science
University of Chicago
1100 East 58th Street
Chicago, Illinois 60637, USA
stefanko@cs.uchicago.edu

May 29, 2004

Abstract

We investigate a new type of graph inequality (in the tradition of Cvetković, Simić, and Capobianco) which is based on the subgraph relation and which allows as terms fixed graphs, graph variables with specified vertices, and the operation of identifying vertices. We present a simple graph inequality which does not have a solution, and show that the solvability of inequalities with only one graph variable and one specified vertex can be decided (in nondeterministic exponential time). The solvability of graph inequalities over directed graphs, however, turns out to be undecidable.

1 A Simple Graph Inequality

Consider this diagram:

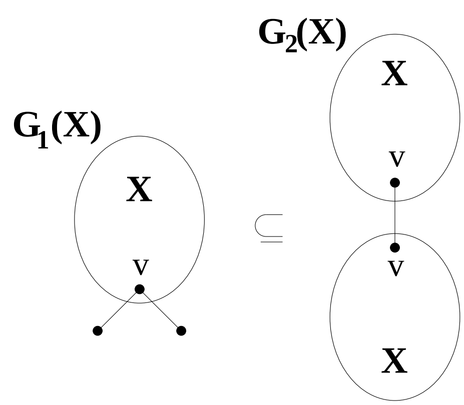

(1)

Figure 1: A Graph Inequality $G_1(X) \subseteq G_2(X)$

Is there a solution to this inequality? More precisely: is there an undirected graph X with a vertex v such that if we construct a graph $G_2(X)$ by taking two copies of X and connecting their v vertices by an edge, and a graph $G_1(X)$ by adding two new vertices to X and connecting them with v, then $G_1(X)$ occurs as a subgraph of $G_2(X)$?

A moment's reflection will show that the answer is yes: take X to be a path of length two together with an isolated vertex v. What happens if we restrict ourselves to connected graphs? Again the answer is yes: take a rooted infinite ternary tree, and connect its root by an edge to a new vertex v. What about finite, and connected graphs, the patient reader will ask? The answer in this case is no, there is no finite, connected graph X fulfilling (1), and this is the main result of this section.

Theorem 1.1 *There is no connected, finite solution of inequality (1).*

A simpler version of this theorem (for finite trees) was used in the first author's thesis [Sch99] to determine the computational complexity of the arrowing relation in Graph Ramsey Theory: deciding $F \rightarrow (T, K_n)$ is complete for the second level of the polynomial-time hierarchy (where F is a finite graph, T a finite tree of size at least two, and K_n the complete graph on n vertices).

Graph equations (more so than graph inequalities) have been studied for a while, and there are two survey papers dating back to the late seventies [CS77, CS79, Cap79]. The equalities and inequalities considered in these papers are more general in that they allow arbitrary operations on graphs such as complementation, tensor products, and squaring. Capobianco, Losi and Riley, for example, showed that there are no (nontrivial) trees whose square is the same as their complement [CLR89]. The more general question of which graphs fulfill $G^2 = \overline{G}$ is still open [BST94], but it is known that the equation has infinitely many solutions [CK95].

We conclude this section with a proof of Theorem 1.1. Section 2 contains a generalization of this result: the solvability of graph inequalities with only one variable having one specified vertex can be decided. In Section 3 we show that a natural generalization of graph inequalities leads to an undecidable solvability problem. Section 4 contains stronger results for graph inequalities over directed graphs: while the solvability of directed graph inequalities with only one variable and one specified vertex remains decidable, we can show that the solvability of directed graph inequalities is undecidable (even with at most three variables, and two specified vertices for each variable).

Before we begin the proof we introduce some standard notation [Die97]. We write $G = (V, E)$ for a graph G with *vertex set* $V = V(G)$ and an *edge set* $E = E(G)$. The edge between vertices $u, v \in V$ is written as (u, v). The *order* of a graph is defined as $|V(G)|$, the *size* $|G|$ as $|E(G)|$. A graph is *finite* if it has finite order, and *connected* if there is a path between any two of its vertices.

Proof of Theorem 1.1. Let X be a minimal solution of the inequality. Denote the copies of X in $G_2(X)$ by $X_i,\ i = 1, 2$. An *element* of X is either its edge or vertex. Given an element x of X we denote the corresponding element of X_i by x_i.

Let ϕ be the embedding of $G_1(X)$ into $G_2(X)$. Clearly $(v_1, v_2) \in \text{Im } \phi$, since otherwise $G_1(X)$ would map into X_1 or X_2. Assume that there is an edge $e \in X$ such that neither e_1 nor e_2 is in Im ϕ. Let Y be the connected component of $X - \{e\}$ containing v. From the connectedness of $G_1(X)$ it follows that Im $\phi \subseteq G_2(Y)$. Now the restriction of ϕ to $G_1(Y)$ is an embedding of $G_1(Y)$ into $G_2(Y)$, contradicting the minimality of X.

Thus for every $e \in X$ either e_1 or e_2 is in Im ϕ. Note that this implies that for every vertex $u \in X$ either u_1 or u_2 is in Im ϕ. Let Y_i be the subgraph of X corresponding to Im $\phi \cap X_i$ (as a subgraph of X_i). Then for each $e \in X$ either $e \in Y_1$, or $e \in Y_2$. We know that

$$Y_1 \cup Y_2 = X \tag{2}$$

$$|V(Y_1)| + |V(Y_2)| = |V(\text{Im } \phi)| = |V(G_1(X))| = |V(X)| + 2 \tag{3}$$

$$|E(Y_1)| + |E(Y_2)| = |E(\text{Im } \phi)| - 1 = |E(G_1(X))| - 1 = |E(X)| + 1 \tag{4}$$

The first equality in Equation (4) follows from the fact that $(v_1, v_2) \in \text{Im } \phi$, but $(v_1, v_2) \notin \text{Im } \phi \cap (X_1 \cup X_2)$. From (2), (3), (4) we conclude that $|V(Y_1 \cap Y_2)| = 2$ and $|E(Y_1 \cap Y_2)| = 1$ which implies that the intersection of Y_1 and Y_2 is a single edge f. We know that $v \in V(Y_1) \cap V(Y_2)$ and hence $f = (v, u)$ for some $u \in V(X)$. Figure 2 illustrates the situation.

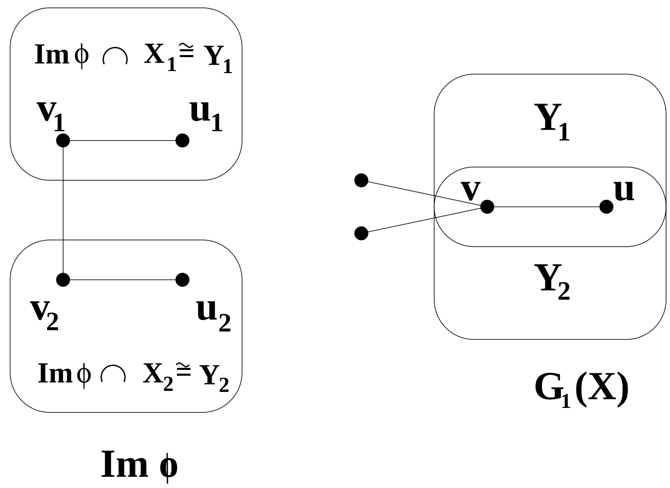

Figure 2: Im ϕ and $G_1(X)$.

Let a_i be the number of vertices from $V(Y_i) \setminus \{u, v\}$ which have degree 1 in X. Let b be 1 if u has degree 1 in X and 0 otherwise. The number of vertices of degree 1 in $G_1(X)$ is $a_1 + a_2 + b + 2$. The number of vertices of degree 1 in Im ϕ is at most $a_1 + a_2 + b + 1$. Hence Im ϕ and $G_1(X)$ are not isomorphic, a contradiction. ■

2 Decidability of Graph Inequalities

We could now start considering all kinds of diagrams involving graphs, vertices, edges and the subgraph relationship. How hard is it to settle these questions? In this section we will show that the solvability of graph inequalities of the type presented in the previous section, i.e., having only one graph variable with one specified vertex, is decidable. This will follow from an (exponential) upper bound on the size of a minimal solution (if there is one). This result will be complemented by the undecidability result of the next section.

Let us formalize the question. A graph variable X with a set of specified vertices $v_1, \ldots, v_m$ represents an unknown finite, connected graph whose vertex set includes vertices $v_1, \ldots, v_m$. Given several graph variables $X_1, \ldots, X_n$, and a graph G we can construct a graph term $G(X_1, \ldots, X_n)$ (called *gterm*) by taking several copies of each X_i, and identifying some specified vertices of the copies with some vertices of G. Since we are working with connected graphs we require $G(X_1, \ldots, X_n)$ to be connected (for any assignment of connected graphs to $X_1, \ldots, X_n$). Note that G itself does not have to be connected and that if $G(X_1, \ldots, X_n)$ is connected for some assignment of connected graphs to $X_1, \ldots, X_n$ then it is connected for all assignments.

Given two such gterms $G_1(X_1, \ldots, X_n)$, $G_2(X_1, \ldots, X_n)$ we can ask whether there exists an assignment of connected finite graphs to the variables $X_1, \ldots, X_n$ such that $G_1(X_1, \ldots, X_n)$ is a subgraph of $G_2(X_1, \ldots, X_n)$. A question of this type we call a *graph inequality*.

For the rest of this section we will consider the simplest possible case of a graph inequality: only one variable, X, with one specified vertex v. Let $G_1(X)$ be a gterm consisting of a connected graph H and a copy of X attached with v to each vertex of a multisubset $I = \{i_1, \ldots, i_\ell\}$ of vertices of H. Similarly construct $G_2(X)$ from a connected graph F and a multisubset $J = \{j_1, \ldots, j_k\}$ of vertices of F. The copy of X in $G_2(X)$ attached to j_r, $(1 \le r \le k)$ is called $X_{(r)}$ and the copy of X in $G_1(X)$ attached to i_r, $(1 \le r \le \ell)$ is called $X_{[r]}$. If there is only one copy of X in $G_1(X)$ we call it X.

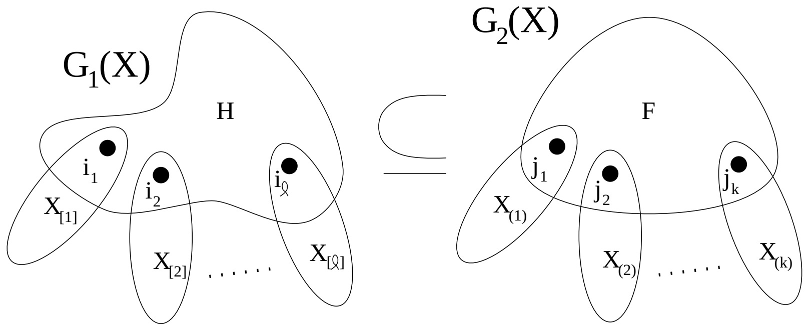

(5)

Figure 3: Inequality $G_1(X) \subseteq G_2(X)$

Theorem 2.1 *If* (5) *has a solution X, then it has a solution of size at most* $|F|(1+k)^{|H|}$.

The upper bound on the size of a minimal solution is exponential in the size of the equality, hence to decide solvability we just have to test all graphs up to that size, something which can be done in nondeterministic exponential time (**NEXP**).

Corollary 2.2 *The solvability of graph inequalities of type* (5) *can be decided in* **NEXP**.

We do not know the precise computational complexity of the decision problem. It is at least **NP**-hard, since we can ask whether a graph contains a clique.

At the core of the proof are Lemmas 2.5 and 2.7 which show that for a minimal solution to the graph inequality (if it exists) we can assume that all of the vertices of I are mapped to vertices of F. This reduces the problem to a simpler variant (namely the images of vertices from I are prescribed) dealt with by Lemma 2.4 (based on the representation result of Lemma 2.3).

First we characterize solutions of inequalities (with prescribed mapping) where on the left side there is only one copy of X and v has to map to a vertex w of F on the right side:

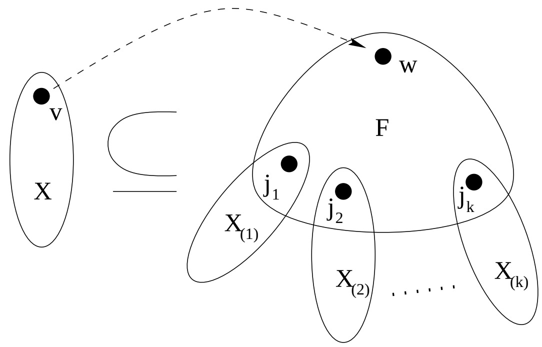

(6)

Figure 4: Inequality $X \subseteq G_2(X),\ v \to w$

If $w \in J$ then any connected graph is a solution. Now assume $w \notin J$. Let Σ be the alphabet consisting of the numbers $1, ..., k$. For each word α from Σ^* take a copy $F^{(\alpha)}$ of F. For every $\alpha \in \Sigma^*$ and $a \in \Sigma$ identify $w^{(\alpha a)}$ and $j_a^{(\alpha)}$. The resulting infinite graph is called F^∞ (see Figure 5).

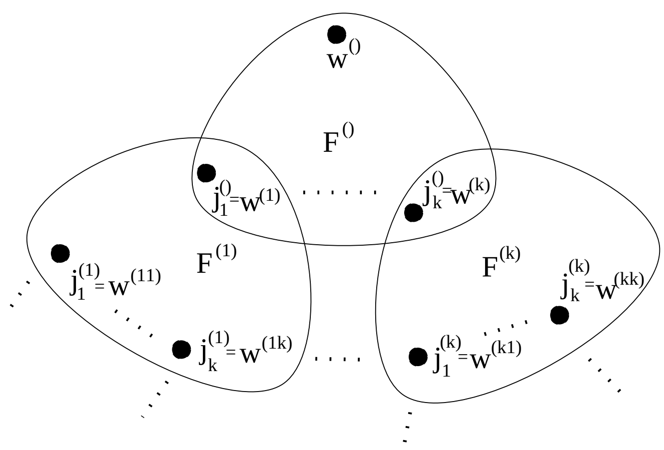

Figure 5: F^∞

Lemma 2.3 *Assume that $w \notin J = \{j_1, \ldots, j_k\}$. Then the solutions of inequality (6) are precisely the subgraphs X of F^∞ with $v = w^{()}$ such that*

$$\textit{for any edge } e \textit{ in } F, \textit{ any } \alpha \in \Sigma^*,\ a \in \Sigma \textit{ if the edge } e^{(a\alpha)} \textit{ is in } X, \textit{ then } e^{(\alpha)} \textit{ is in } X. \tag{7}$$

Proof. If X is a subgraph of F^∞ satisfying condition (7) then X is a solution of the inequality via mapping ϕ:

$$\phi(x^{()}) = x$$
$$\phi(x^{(a\alpha)}) = x^{(\alpha)}_{(a)}$$

If X is a solution of the inequality via mapping $\phi : X \to G_2(X)$ then define

$$Y^{()} = \phi^{-1}(F)$$
$$Y^{(a\alpha)} = \phi^{-1}(Y^{(\alpha)}_{(a)})$$

where $Y_{(a)}^{(\alpha)}$ is the copy of $Y^{(\alpha)}$ in $X_{(a)}$ in $G_2(X)$. If e is an edge of X with distance d from v then it either must map to F or to some edge f in some $X_{(r)}$ which has strictly smaller distance from $v_{(r)}$ than d. Edges adjacent to v must be mapped to F and hence they are in $Y^{()}$. By induction it follows that

$$X = \bigcup_{\alpha \in \Sigma^*} Y^{(\alpha)}$$

Clearly $Y^{(\alpha)}$ is a subgraph of F via $\phi^{|\alpha|+1}$ for any $\alpha \in \Sigma^*$. The element of $Y^{(\alpha)}$ corresponding to $x \in F$ is called $x^{(\alpha)}$. By induction it follows that $w^{(\alpha a)} = j_a^{(\alpha)}$ for any $\alpha \in \Sigma^*,\ a \in \Sigma$. From the definition of Y's if $e^{(a\alpha)}$ is in X then the edge $e^{(\alpha)}$ is also in X for any $\alpha \in \Sigma^*,\ a \in \Sigma$. Hence X is a subgraph of F^∞ satisfying (7). ∎

Solving systems of simple graph inequalities is useful in solving more complicated inequalities.

Lemma 2.4 *If a system of inequalities with prescribed mappings*

$$H_1 \subseteq X,\ h_1 \to v\ ;\ \ldots\ ;\ H_m \subseteq X,\ h_m \to v \tag{8}$$

$$X \subseteq F_1(X),\ v \to w_1\ ;\ \ldots\ ;\ X \subseteq F_n(X),\ v \to w_n \tag{9}$$

has a solution, then it has a solution of size at most $|F_1|(1+k_1)^M$ where k_1 is the number of copies of X in F_1 and $M := \max\{|H_1|, \ldots, |H_m|\}$ assuming that the graphs $H_1, \ldots, H_m$ are connected.

Proof. Let X be a minimal solution of the system. Let e be an edge of X whose distance d from v is maximal. Assume that $d > M$. If we remove the edge e then $X' = X - \{e\}$ still satisfies inequalities (8), because no edge of any $H_i,\ (1 \le i \le m)$ can map to e. If X satisfies (6) for $F = F_i\ (1 \le i \le n)$ then by Lemma 2.3 it is a subgraph of F^∞ with $v = w^{()}$ and it satisfies condition (7). Let $e = f^{(\alpha)}$. Clearly X' is also a subgraph of F^∞ and the condition is still satisfied, because $\mathrm{dist}(v, f^{(a\alpha)}) > \mathrm{dist}(v, f^{(\alpha)})$ and hence $f^{(a\alpha)} \notin X'$ for any $a \in \Sigma$. Therefore X' satisfies inequalities (9), a contradiction to the minimality of X.

Thus $\mathrm{dist}(v, e) \le M$. The size of the subgraph of F_1^∞ consisting of edges within distance M from v is bounded by $|F_1|(1+k_1)^M$. ∎

Now we return to inequality (5).

Lemma 2.5 *If there is more than one copy of X on the left side of inequality (5) then every $i_r = v_{[r]},\ (1 \le r \le \ell)$ must map to a vertex of F.*

Proof. Suppose, for example, that i_1 maps into some $X_{(r)} - \{j_r\}$. Let P be a path from i_1 to i_2. Graphs $X_{[1]}$ and $X_{[2]} \cup P$ share only vertex i_1. Hence the image of at least one of them does not contain j_r and since j_r is a cutvertex of G_2 that image must be contained in $X_{(r)} - \{j_r\}$, which is impossible, since there are more vertices in X_1 or in $X_2 \cup P$ than in $X_{(r)} - \{j_r\}$. ∎

Lemma 2.6 *If X is a solution of an inequality (5) via mapping $\psi : G_1(X) \to G_2(X)$ then there exists a mapping $\phi : G_1(X) \to G_2(X)$ such that $\phi(i) = \psi(i)$ and as many copies of X in i as possible are mapped to copies of X in $\phi(i)$ for every $i \in I$.*

Proof. Consider a bipartite graph B with partitions I and J where i_r is connected to j_s iff $\psi(i_r) = j_s$. W.l.o.g. assume that $\{(i_r, j_r); 1 \le r \le t\}$ is a maximal matching of B.

We need to show that there exists ϕ such that $X_{[r]}$ maps to $X_{(r)}$ for $1 \le r \le t$. Let $Y^1, \ldots, Y^q$ be the connected components of $X - \{v\}$. Let ϕ be a mapping such that

$$\sum_{r=1}^{t} \sum_{j=1}^{q} \left| \phi(Y_{[r]}^j) \cap Y_{(r)}^j \right| \tag{10}$$

is maximal. If for some r, j:

$$\phi\big(Y^j_{[r]}\big) \neq Y^j_{(r)}$$

then clearly $\phi\big(Y^j_{[r]}\big) \cap Y^j_{(r)} = \emptyset$, otherwise $\phi\big(Y^j_{[r]}\big)$ would have to contain j_r. Now we can change ϕ in such a way that $Y^j_{[r]}$ will be mapped to $Y^j_{(r)}$ and $\phi^{-1}\big(Y^j_{(r)}\big)$ will be mapped to $\phi\big(Y^j_{[r]}\big)$. This increased (10), a contradiction. Hence ϕ maps $X_{[r]}$ to $X_{(r)}$, $(1 \leq r \leq t)$. ■

We prove an analog of Lemma 2.5 for inequalities where X occurs only once on the left side of (5).

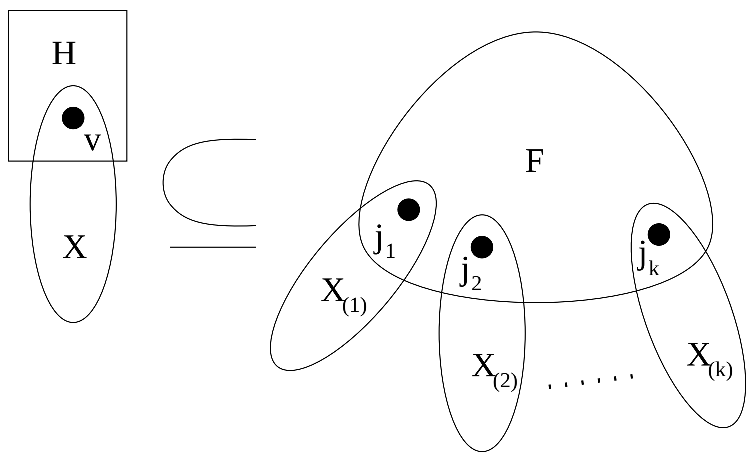

(11)

Lemma 2.7 *If (11) has a solution, then it has a solution X via a mapping ϕ which maps $v = i_1$ to a vertex of F.*

Proof. Suppose that there is no solution of inequality (11) such that v maps to a vertex of F, but there is a solution in which v maps into a vertex of $X_{(1)} - \{j_1\}$. Then clearly the following inequality with the condition that v must map to some $u \in X_{(1)}$ has a solution:

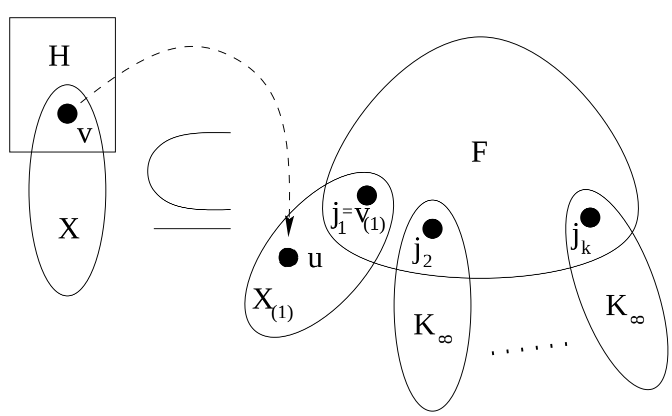

(12)

If $u = j_1$ then by Lemma 2.6 there is ϕ such that X is mapped to $X_{(1)}$. Therefore we can replace K_∞'s in (12) by $K_{|H|}$'s, since only H is mapped to $G_2(X) - X_{(1)}$. This, however, implies that $X = K_{|H|}$ is a solution of (11) in which v maps to a vertex of F, a contradiction.

Thus $u \neq j_1$ for every solution of (12). Let X be a minimal solution of (12). Graphs H and X share only v; moreover j_1 is a cutvertex of G_2 and hence either H or X must be mapped inside $X_{(1)} - \{j_1\}$. Since the later is not possible, H must be mapped inside $X_{(1)} - \{j_1\}$.

Now let $Y = \phi^{-1}\left(X_{(1)}\right) \cap X$ and $Z = \phi^{-1}\left(G_2(X) - (X_{(1)} - \{j_1\})\right)$. The common vertex of Y and Z is called $q = \phi^{-1}(j_1)$. Inequality (12) implies:

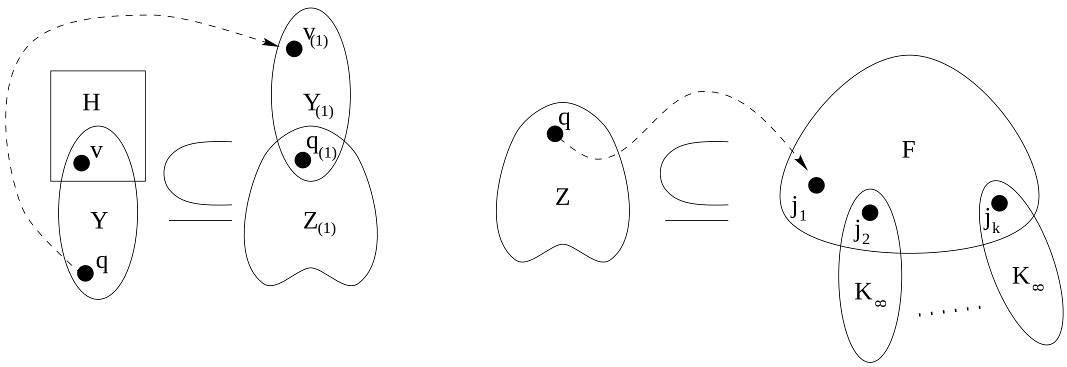

(13)

The second inequality follows directly from the definition. To see the first inequality, note that the graph on the left hand side is a subgraph of $X_{(1)}$ with q mapping to $j_{(1)}$, and that by definition of Y and Z the right hand side contains $X_{(1)}$ with $j_{(1)}$ of $X_{(1)}$ mapping to v of Y.

If in the first inequality v was mapped outside of $Y_{(1)}$ then the shortest path from q to v would have to map to a longer path, which is not possible. Hence v maps inside $Y^{(1)}$. Combining the two inequalities of (13) we get that Y satisfies inequality (12). This contradicts the minimality of X. ∎

We can now complete the proof of Theorem 2.1 by showing a bound on the size of a minimal solution (if there is one) of graph inequalities with one variable and one specified vertex.

Proof of of Theorem 2.1. From Lemmas 2.5 and 2.7 it follows that we only need to consider solutions in which every i_r, $(1 \le r \le \ell)$ maps to a vertex of F. For each such mapping ϕ, using Lemma 2.6, we can assume that if $i \in I$ maps to a vertex $j \in J$ then as many copies of X in i as possible map to copies of X in j.

Let

$$G'_1(X) \subseteq G'_2(X), \; = i_1 \to \phi\left(i_1\right), \ldots, i_\ell \to \phi\left(i_\ell\right) \tag{14}$$

be the inequality with prescribed mappings obtained by removing those $X_{[r]}$'s and $X_{(r)}$'s which are already taken care of by Lemma 2.6. Notice that now no i'_r, $(1 \le r \le \ell')$ maps to a j'_s, $(1 \le s \le k')$.

Let X' be a solution of (14) with mapping ψ. If $\psi\left(X'_{[r]}\right) \cap X'_{(s)} \neq \emptyset$ then some vertex from $X'_{[r]} - \{i'_r\}$ must map to j'_s. Since j'_s is a cutvertex, no other part of $G'_1(X')$ can map to $X'_{(s)}$. If for each $X'_{[r]}$, $1 \le r \le \ell'$ and H we take the set of objects (edges and $X'_{(s)}$'s) to which it is mapped, then these sets are disjoint.

There are only finitely many partitions of the objects of $G'_2(X)$ into $\ell+1$ disjoint sets. For each such partition we get a system of inequalities with prescribed mappings as in Lemma 2.4, which has a solution of size at most $|F|(1+k)^{|H|}$ (if it has one). ■

Note that using previous Lemmas we can easily prove Theorem 1.1. If there was a solution of (1) then by Lemma 2.7 there is a solution such that v from $G_1(X)$ maps to one of the v's in $G_2(X)$. By looking at the degrees of v's we see that this is not possible.

We conclude this section with a technical result that allows us to combine several inequalities with prescribed mappings. This lemma will be needed in the next section.

Lemma 2.8 *For any system of inequalities with prescribed mappings*

$$H_1 \subseteq X,\ h_1 \to v\ ;\ \ldots\ ;\ H_m \subseteq X,\ h_m \to v$$
$$X \subseteq F_1(X),\ v \to w_1\ ;\ \ldots\ ;\ X \subseteq F_n(X),\ v \to w_n$$

there is a single inequality which has the same set of solution as the system.

Proof. Consider the following inequality:

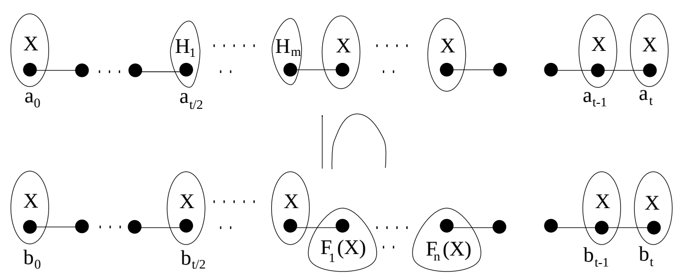

(15)

By Lemma 2.5 a_0 and a_t have to map to F. Clearly the a_0, a_t path of H in $G_1(X)$ has to map to a path in F in $G_2(X)$. If $t > 2(m+n+\max\{F_1,\ldots,F_n\})$ then the only path of length t in F is the b_0, b_t path. It follows that a_i maps to b_i, $(0 \le i \le t)$, because a_{t-1} cannot map to a_1. Hence X is a solution of (15) iff it is a solution of the system. ■

3 Undecidability of Graph Inequalities

The result of the last section might suggest that there is a general method to decide the solvability of graph inequalities. While we have to leave this question open for the time being, we do want to sketch a proof that a natural generalization of the problem turns out to be undecidable. We consider a logical language whose atoms are graph inequalities as above, i.e., diagrams involving graphs with labeled vertices, additional edges and vertices, and one occurrence of the subgraph relationship. We then build more complex formulas by allowing logical operators $\wedge$ (and), and $\neg$ (not), and quantifiers over graphs (and labeled vertices). We will not formally describe the semantics of this language since it is straightforward; the only point worth mentioning is that we assume vertices with different labels in the same graph to be different.

We will next show that formulas of this type are not decidable. More precisely we will show that this is even the case if we restrict the quantifiers in the formulas to be only existential or bounded (i.e., of the form ($\forall F \subseteq G$) or ($\exists F \subseteq G$)). Since formulas involving only bounded quantifiers are decidable (the bounds have to be explicit graphs, hence we can try all possible combinations), this is a reasonably sharp result on the complexity of graph inequalities. The main open problem of interest of course is whether the problem is undecidable in case we only allow existential quantifiers (and no bounded quantifiers at all). We will mention some interesting related problems in the conclusion.

Theorem 3.1 *The solvability of graph diagrams with Boolean operators, existential quantifiers and bounded quantifiers is not decidable.*

Proof. We will show the undecidability result by reducing the word problem for Semi-Thue systems to it (see for example [HU79]). Over an alphabet A a Semi-Thue system is a set of productions $x \Rightarrow y$ ($x, y \in A^*$) meaning that x can be transformed into y. The word problem for a Semi-Thue system is to decide whether given two words x and y there is a series of productions which applied to substrings of the words, transforms x into y.

We will represent the letters of the alphabet as paths of different lengths. A word will be coded as a path to which are attached further paths coding the letters of the word. A sequence of words will be coded in a similar way. We will then have to find a way to verify that such a sequence results from legal applications of the productions.

Fix a Semi-Thue system $(x_i \Rightarrow y_i)_{i \leq n}$ over some alphabet A, and suppose we are given two words x and y. The following diagram gives an example of how we represent words, in this case the word 21130 (Figure 6).

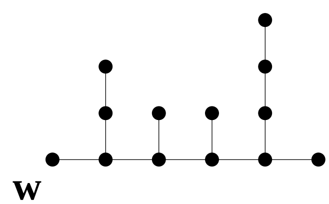

Figure 6: Representing the word 21130

The initial vertex w is used to link the word up in a sequence of words. In the manner depicted by the diagram we associate graphs X_i, Y_i, X and Y with the words x_i, y_i, x and y.

Assume that for all $A \subseteq G$ the following diagram (Figure 7) is true:

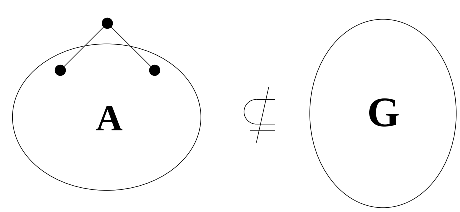

Figure 7: Forcing a tree

Then G does not contain any cycles, and therefore is a tree. Furthermore by excluding $K_{1,4}$ we can easily assure that G has maximal degree at most three. We now set up G to code the initial and final word. We do this by saying that there is a $A \subseteq G$ which fulfills the following diagram (Figure 8).

Note that for the diagram to be true w_X has to be mapped to u and w_Y to v (G is a tree). Hence G will contain a path from u to v. For each vertex w on that path let G_w be the graph attached to the path (if none, then just w). With the previous diagram we have assured that G_{w_X} codes x and

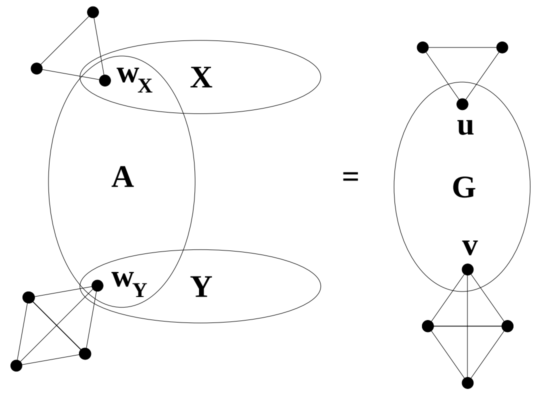

Figure 8: Forcing x and y

G_{w_Y} codes y. We only have to verify now that the transitions between words as coded by G are legal according to the system of productions given. We do this by saying that for any $A, B, C, D \subseteq G$ for which the following diagram is true (Figure 9),

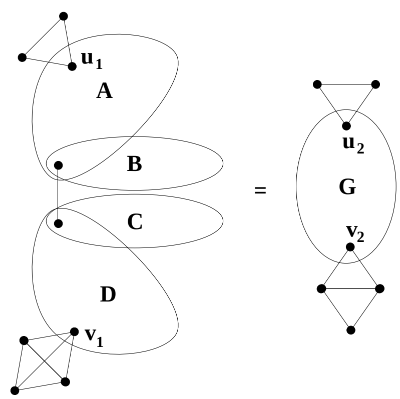

Figure 9: Transition from B to C

then there are $S, E, B', C' \subseteq G$ for which the next diagram is true (Figure 10), and such that $B' = X_i$ and $C' = Y_i$ for some $i \leq n$.

It is straightforward to check that in this manner we have encoded the original word problem: there is a G fulfilling all these conditions, if and only if there is a solution to the word problem. Hence the word problem can be written as a graph inequality with one existential quantifier and some bounded quantifiers. ■

4 Directed Graph Inequalities

So far we have only considered undirected graphs. What happens if we change the universe of graph inequalities to directed (or colored) graphs? Call these variants *directed (colored)* graph inequalities, respectively.

In the case of one variable with one specified vertex we can obtain the same result as in Theorem 2.1. As a matter of fact, the lemmas and proofs needed for that theorem can be used without modification.

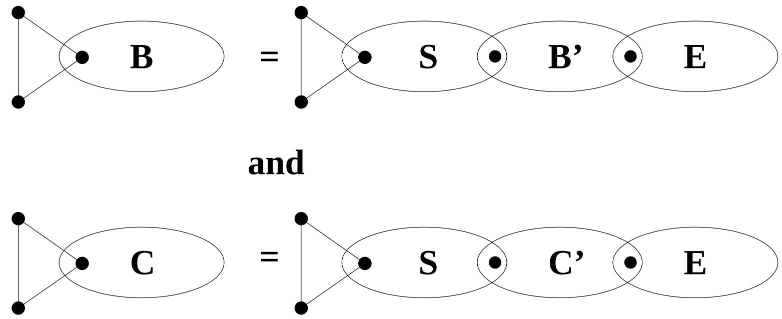

Figure 10: Application of production B' to C'

Theorem 4.1 *For directed (or colored) graphs, if* (5) *has a solution X then it has a solution of size at most* $|F|(1+k)^{|H|}$.

As above, this implies that the problem is decidable in **NEXP**.

The complexity of the undecidability proof in Section 3 stemmed from the difficulty of coding the alphabet: we had to use special devices to code letters, and then use bounded quantifiers to verify that the coding was correct. Allowing the edges in the graph to be directed, however, makes these constructions unnecessary.

Theorem 4.2 *The solvability of directed (colored) graph inequalities is undecidable.*

The problem remains undecidable even if we limit it to three variables with two specified vertices each. We only consider directed graphs, since the treatment for graphs with two colors is identical.

Proof. We will translate Post's Correspondence Problem (PCP) into a directed graph inequality. Since the former problem is known to be Turing-complete [HU79] this shows the undecidability of directed graph inequalities.

Post's Correspondence Problem asks, whether given a list of pairs of words $(p_i, q_i)_{1\leq i\leq n}$ there is a list of indices $i_1, \ldots, i_m$ such that $p_{i_1} \ldots p_{i_m} = q_{i_1} \ldots q_{i_m}$. PCP can be translated into a question about context-free grammars as follows: consider two grammars

(*i*) $S_1 \rightarrow \underline{i}\, S_1 p_i \mid \underline{i}\, p_i \ (1 \leq i \leq n)$,

(*ii*) $S_2 \rightarrow \underline{i}\, S_2 q_i \mid \underline{i}\, q_i \ (1 \leq i \leq n)$,

where $\underline{i}$ is a prefix-encoding of the number i. The original problem has a solution, if and only if the two grammars have a word in common, i.e., there is a word w such that $S_1 \rightarrow^* w$ and $S_2 \rightarrow^* w$.

Consider a context-free grammar with productions over the alphabet $\{0,1\}$ and one nonterminal symbol S. Every production has S on the left-hand side, and a (nonempty) string of letters and at most one occurrence of S on the right-hand side.

We will code zeroes and ones by the direction of edges, an outgoing edge coding a 0 (for a string starting in the vertex), an incoming edge a 1. Let G_a be the path corresponding to the string a (for an example see Figure 11).

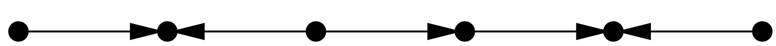

Figure 11: G_{01001}

A production is either of the form $S \rightarrow aSb$, where $ab \in \{0,1\}^+$, or of the form $S \rightarrow a$, where $a \in \{0,1\}^+$. We assume that there is always a production of the second kind.

Construct a graph inequality as follows: the left-hand side contains a graph variable X_S with two special vertices u_S and v_S. The right-hand side has two special vertices u'_S and v'_S. For every production of the form $S \to aSb$ include G_a starting in u'_S, and ending in the u_S vertex of a new copy of X_S, and G_b starting in the v_S vertex of X_S and ending in v'_S. For every production of the form $S \to a$ include G_a starting in u'_S and ending in v'_S.

If we require that u and v be mapped to u' and v', respectively, then a solution to the inequality corresponds to a word in the language described by the grammar, and, vice versa, every word in the language gives rise to a solution of the graph inequality.

For an example see Figure 12 which shows the graph inequality belonging to the system $S \Rightarrow 0S100 \mid 10S11 \mid 11S00 \mid 0100 \mid 1011 \mid 1100$.

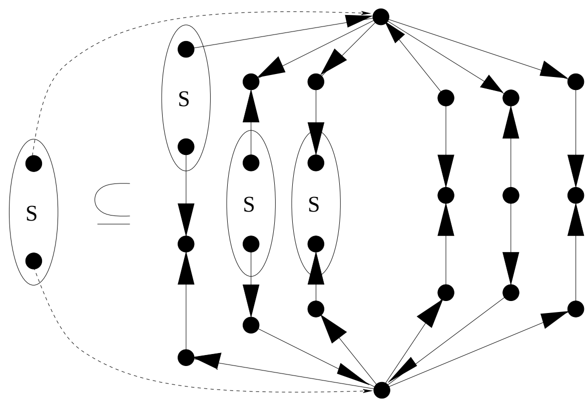

Figure 12: Graph Inequality for Semi-Thue System

We will first prove the claim that for every word in the language there is a corresponding solution of the graph inequality in a stronger form: for each n there is a graph G_S such that

(i) G_S solves the inequality (with u, v mapping to u', v'), and

(ii) there is a path G_w between u_S and v_S in G_S for every word w that can be derived in n steps from S.

We prove this statement by induction on n. For $n = 1$ let G_S consist of all paths G_a for which $S \to a$ is a production, and identify their starting vertices (calling it u_S), and their end vertices (calling it v_S). For the induction step, assume we have a graph G'_S with vertices u'_S and v'_S fulfilling the induction hypothesis for n. Build G_S with vertices u_S and v_S by including for each production $S \to aSb$ (new) copies of G_a, G'_S and G_b, and identifying u_S with the starting vertex of G_a, u'_S with the ending vertex of G_a, v'_S with the starting vertex of G_c, and v_S with the ending vertex of G_c. It is easy to show by induction that the graphs so constructed fulfill (i) and (ii) .

For the other direction suppose that there is a solution G_S to the graph inequality. We will show that for any path P from u_S to v_S in G_S there is a word w such that $S \to^* w$ and $P = G_w$. Use induction on the length of the path: let P be a path of minimal length between u_S and v_S for which the assertion has not been proven yet. P has length at least one (since u_S and v_S are different vertices). Fix w such that $P = G_w$. Since G_S fulfills the inequality, P must be a subpath of the right-hand side of the inequality starting in u'_S and ending in v'_S. The way the right-hand side was constructed, P must therefore be a subpath in a graph corresponding to a particular production $S \to aSb$, or $S \to a$. In the later case $a = w$ and we are done. In the former case P consists of three parts corresponding to

a, S and b, respectively. Since a and b together have length at least one, we can apply the induction hypothesis to the subpath of P corresponding to S.

If we are given two grammars $\mathcal{G}_1, \mathcal{G}_2$ we can construct the inequalities for them as above, and ask whether there are graphs fulfilling them, and a path P from u_P to v_P which is subgraph both of X_{S_1} and X_{S_2} where u_P, and v_P have to be mapped to u_{S_i} and v_{S_i} $(i = 1, 2)$. Such a path corresponds to a word w which can be derived in both grammars. We are left with the task of combining the inequalities into a single inequality fulfilling the additional requirements on the u and v vertices.

Consider the directed graph inequality of Figure 13.

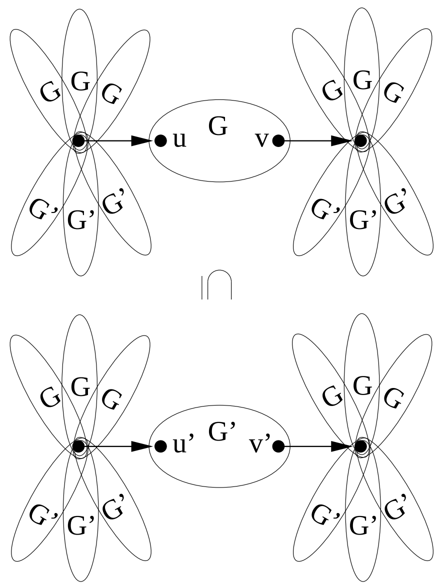

Figure 13: $G \subseteq G'$

We claim that if G and G' are solutions of this inequality, then G is a subgraph of G' such that u and v are mapped to u' and v' respectively (and, obviously, any such graphs are solutions to the inequality). To see this, suppose that one of the vertices at the heart of a sunflower does not map to its corresponding vertex. It then has to map to a labeled vertex, or into a G' or G, say G'. This is not possible, since such a vertex is at the heart of three copies of G', at most two of which can map outside the G', so there would have to be a full copy of G' within G', which is impossible. Hence the hearts of the sunflowers map to each other, and, in consequence, the copies of G map to the corresponding copies of G' while u and v map to u' and v'.

We have four equations altogether: $G_{S_i} \subseteq G_i$ (with G_i the right-hand sides constructed from the grammars), and $G_P \subseteq G_{S_i}$ $(i = 1, 2)$. We can extend the diagram above to incorporate all four inequalities: it will contain five sunflowers on each side of the inequality, between which the terms of the four inequalities are linked up; each sunflower will have three copies of each graph involved in the construction, hence the hearts of the sunflowers map to each other, as above. Thus we get a single directed graph inequality which has a solution if and only if the two grammars have a word in common. ■

5 Conclusion

Several questions remain open, the most nagging one being the complexity of deciding the solvability of (undirected) graph inequalities (without additional quantifiers and Boolean operators). It seems hard to translate the corresponding undecidability result for directed graph inequalities back to the undirected case. Another approach would be to strengthen the proof of the undirected undecidability result which required one existential quantifier and several alternations of bounded quantifiers. It seems likely that using a different problem for the reduction (for example Post's Correspondence Problem) one might get the language down to existential and bounded universal quantifiers only. Getting rid of that last layer of bounded quantifiers, and thereby settling the complexity of Boolean combinations of graph inequalities, seems harder. The language shown undecidable in Section 3 for example is powerful enough to code the edge reconstruction conjecture (in a more or less natural fashion). Hence a decision procedure would have come as a surprise. In the case of graph inequalities the situation is different: we do not know of any difficult open problem that can be phrased as a graph inequality, hence decidability might still be an option.

Question 1 *Is the solvability of graph inequalities (as defined in Section 2) decidable?*

A positive indication for decidability is that it seems difficult to force large solutions. If graph inequalities were undecidable, then the solution size would have to grow faster than any computable function. The best result we have been able to obtain so far shows that a quadratic lower bound is possible, a far cry from undecidability.

Theorem 5.1 *There is a graph inequality $G_1(X) \subseteq G_2(X)$ of size $O(n)$ such that the size of a minimal solution is $\Omega(n^2)$.*

Proof. Consider the following system of inequalities with prescribed mappings:

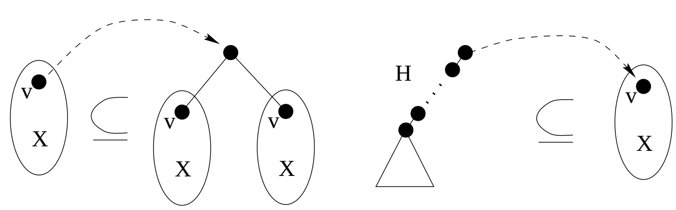

(16)

where H is a path of length n connected to a complete binary tree of depth $\log n$. Let B be the infinite binary tree with edges naturally labeled by strings from $\{0,1\}^+$. By Lemma 2.3 solutions of the first inequality are subgraphs X of B such that if edge $a\alpha$ is in X then edge α is also in X for any $a \in \{0,1\}$, $\alpha \in \{0,1\}^+$. From the second inequality it follows that for any solution X there is some $\alpha \in \{0,1\}^n$ such that for every $\beta \in \{0,1\}^{\log n}$ edge $\alpha\beta$ is in X. Hence for any suffix γ of α for every $\beta \in \{0,1\}^{\log n}$ edge $\gamma\beta$ is in X and therefore there are $\Omega(n^2)$ edges in X. Using Lemma 2.8 we combine (16) into a single inequality. ∎

Question 2 *Are there graph inequalities whose minimal solutions have at least exponential size?*

Our decidability result for graph inequalities with one variable (and one labeled vertex) shows that the computational complexity of the problem lies in **NEXP**. As we pointed out earlier, it is also **NP**-hard (since we can ask for a clique as subgraph, without even using the existential quantifier).

Question 3 *What is the computational complexity of deciding the solvability of one variable, one vertex graph inequalities? Is the problem* **NEXP***-complete?*

First steps towards generalizations of the decidability result would probably try to increase the number of specified vertices, then the number of variables. Also, can we decide Boolean combinations of graph inequalities?

One special case of Boolean combinations we can settle with the techniques from Section 2: graph equalities with one variable and one specified vertex.

Theorem 5.2 *The solvability of graph equalities with one variable and one specified vertex is decidable.*

Proof. Lemma 2.5 allows us to assume that variable X occurs at most once on each side of the equality (otherwise we can use Lemma 2.4 as in the proof of Theorem 2.1). If X does not occur on one of the sides we are done. If it occurs precisely once on each side, it is not too difficult to see that the equality is solvable, if the two graphs the variable is attached to are isomorphic (where the labeled vertices have to map to each other). The decision procedure outlined here is, again, in **NEXP**. ■

In the case of directed graph inequalities we have a tight separation of decidability and undecidability: one variable with one specified vertex is decidable, three variables with two specified vertices are not. While it might be interesting to find out what happens in the case of two variables, a more promising object of study should be the computational complexity of directed graph inequalities. The direction of the edges might help in encoding a problem complete for **EXP** or **NEXP**.

Question 4 *What is the computational complexity of deciding the solvability of one variable, one vertex directed (or colored) graph inequalities? Is the problem* **NEXP***-complete?*

Finally we would like to suggest that the question of computational complexity should also be an interesting one for the more general types of graph equalities and graph inequalities studied in the literature [CS79].

Acknowledgments. We would like to thank Laci Babai and Janos Simon for helpful discussions.

References

[BST94] Vladimir Baltić, Slobodan K. Simić, and Velibor Tintor. Some remarks on graph equation $G^2 = \overline{G}$. *Univ. Beograd. Publ. Elektrotehn. Fak. Ser. Mat.*, 5:43–48, 1994.

[Cap79] Michael F. Capobianco. Graph equations. *Annals New York Academy of Sciences*, 319:114–118, 1979.

[CK95] Michael Capobianco and Suh-Ryung Kim. More results on the graph equation $G^2 = \overline{G}$. In *Graph theory, combinatorics, and algorithms*, pages 617–628. Wiley, 1995.

[CLR89] Michael F. Capobianco, Karen Losi, and Beth Riley. $G^2 = \overline{G}$ has no nontrivial tree solutions. *Annals New York Academy of Sciences*, 555:103–105, 1989.

[CS77] Dragoš M. Cvetković and Slobodan K. Simić. Graph equations. In *Contributions to graph theory and its applications*, pages 40–56. Technische Hochschule Ilmenau, 1977.

[CS79] Dragoš M. Cvetković and Slobodan K. Simić. A bibliography of graph equations. *Journal of Graph Theory*, 3(4):311–324, 1979.

[Die97] Reinhard Diestel. *Graph Theory.* Springer, New York, 1997.

[HU79] John E. Hopcroft and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation.* Addison-Wesley, 1979.

[Sch99] Marcus Schaefer. *Completeness and Incompleteness.* PhD thesis, University of Chicago, June 1999.