Fixing the UPCC case of Split-or-Johnson

László Babai, 1-14-2017

1 The story

In November 2015 I announced the result that Graph Isomorphism can be tested in quasipolynomial time. The first version of the full paper was posted on arXiv in December 2015, and an updated version in January 2016 (arXiv:1512.03547). Below we refer to Version 2 (Jan 2016) of that paper as [Q2].

On January 1, 2017, Harald Helfgott pointed out an error in the timing analysis of one of the combinatorial Divide-and-Conquer tools, the "Split-or-Johnson" routine (SoJ). The error invalidated the quasipolynomial claim. A revised analysis of a slightly modified algorithm (adjusting a threshold parameter for optimal analysis) gave a running time bound of $\exp\exp(\widetilde{O}(\sqrt{\log n}))$ where the $\widetilde{O}$ hides a poly-loglog term. While still subexponential (less than $\exp(n^{\epsilon})$ for all positive ϵ) and thus considerably stronger than the previously known moderately exponential bound of $\exp(\widetilde{O}(\sqrt{n}))$ (Luks, 1983), this bound was not quasipolynomial. I announced the revised bound on my home page on January 4, 2017. Three days later I found a simple fix and announced this on January 9, 2017 both on the website and in a lecture I gave that day at Georgia Tech. So the quasipolynomial claim has been restored.

Meanwhile I have been working on a major revision of the arXiv posting, taking into account a large number of helpful comments from colleagues. The intention of the update is to improve the presentation, give more detailed proofs and analysis, and also to fix another, less dramatic, error that was pointed out to me by Jin-Yi Cai in Spring 2016; that error occurred in the "Design Lemma" algorithm (DL). My analysis of the modified algorithm for the DL was additionally verified by Gábor Tardos and by Harald Helfgott, both of whom also contributed many comments on how to improve the presentation.

Since the completion of the revision of the paper is still months away, and given the explosive interest in the recent (SoJ) update, I decided to post this update in advance of the revision of the full paper. I will also post the DL update in the near future.

This note requires some familiarity with [Q2].

The error was in Case 8(iii) ("uniprimitive (UPCC) case") of the "Bipartite Split-or-Johnson" algorithm (Sec. 7.5 in [Q2]) to which the proposed solution (Sec. 7.9 in [Q2]) involved a recursive call to the "UPCC Split-or-Johnson" algorithm. This recursive call caused an impermissible blow-up in the running time.

The present note shows how to replace this "malignant" recursive call by a "benign" one that leads to the same type of recurrence that occurs in many other places in the paper and leads to a quasipolynomial solution.

This update also subsumes Case 8(iv) ("Johnson case") of the "Bipartite Split-or-Johnson" algorithm; consequently, that case (Secs. 7.10 and 7.11 in [Q2]) is now eliminated.

2 Notation, terminology, and preliminary observations

We use the term "bipartite graph" to denote a triple of the form $X = (A, B; E)$ where $E \subseteq A \times B$. So the edges of such a graph are oriented from A to B. By the "degree" of the vertices in A we mean their out-degree, and for vertices in B their in-degree. We say that X is *semiregular* if every vertex in A has the same degree and every vertex in B has the same degree. The *trivial* bipartite graphs are the empty ($E = \emptyset$) and complete ($E = A \times B$) bipartite graphs.

By a *coherent configuration* (CC) we mean a classical one, as defined in Sec. 2.5 in [Q2]. We write $\mathfrak{X} = (V; R_1, \ldots, R_r)$ for a CC on vertex set V with edge-color classes $R_1, \ldots, R_r$. The R_i form a partition of $V \times V$. For $x, y \in V$ we write $c(x, y) = i$ if $(x, y) \in R_i$ and call $c(x, y)$ the color of the edge (x, y). We write $c(c) := c(x, x)$ and call it the color of vertex x.

A CC is *homogeneous* if all vertices have the same color. A CC is *primitive* (PCC) if it is homogeneous and all off-diagonal edge-colors define (strongly) connected graphs. A CC is *uniprimitive* (UPCC) if it is primitive and not a clique, i. e., its rank is $r \geq 3$, i. e., there are at least two off-diagonal edge-colors.

Notation 2.1 (Intersection numbers)**.** For a CC, we write p^k_{ij} for the intersection number defined as being, for any $(x, y) \in R_k$, the number of vertices z such that $c(x, z) = i$ and $c(z, y) = j$.

(These numbers are called "structure constants" in [Q2], in reference to the structure constants of the associated adjacency algebra. However, "intersection numbers" is the more common term.)

Notation 2.2. For a relation $R \subseteq V \times V$ on a set V, let $R(x)$ denote the set of out-neighbors of $x \in V$ in R, i. e., $R(x) = \{y \in V \mid (x, y) \in R\}$.

Proposition 2.3 (Semiregular neighborhoods)**.** *Let $\mathfrak{X} = (V; R_1, \ldots, R_r)$ be a CC. Then for all $x \in V$ and all colors i, j, k, the bipartite graph $Y = (R_i(x), R_j(x); R_k \cap (R_i(x) \times R_j(x)))$ is semiregular.*

Proof. Let $u \in R_i(x)$, so $c(x, u) = i$. Now the degree of u in color k in Y is p^i_{j,k^-} regardless of the choice of u. □

Definition 2.4 (twins)**.** Let $X = (A, B; E)$ be a bipartite graph. For $x \neq y \in A$ we say that x and y are *twins* if $E(x) = E(y)$; and $x \neq y \in B$ are twins if $E^-(x) = E^-(y)$ (where $E^- = \{(v, u) \mid (u, v) \in E\}$).

Twins in this sense are referred to as "strong twins" in [Q2]. We shall not need "weak twins."

Proposition 2.5 (Twins determined)**.** *Let $\mathfrak{X} = (V; R_1, \ldots, R_r)$ be a CC. Let $A, B \subseteq V$ be two vertex-color classes in $\mathfrak{X}$. Assume $R_i \subseteq A \times B$. Then for all pairs $x, y \in A$, $x \neq y$, the color $c(x, y)$ determines whether or not x, y are twins in the bipartite graph $X = (A, B; R_i)$.*

Proof. Let $c(x, y) = k$ and let the color of A be a, i. e., R_a is the diagonal of A. Now x, y are twins in X if $R_i(x) = R_i(y)$. The latter is equivalent to saying that $p^k_{i,i^-} = p^a_{i,i^-}$. This equation does not depend on the choice of x, y, only on k. □

Proposition 2.6 (Not all twins). *Let $X = (A, B; E)$ be a nontrivial semiregular bipartite graph. Then not all vertices in A are twins.*

Proof. Suppose all vertices in A are twins. Then all have the same neighborhood $N \subseteq B$. Now $N \neq \emptyset$ because X is not empty. But then the vertices in $B \setminus N$ have degree zero. But there are no such vertices by semiregularity, so $B = N$. But then X is the complete bipartite graph, contrary assumption. □

3 The fix

We are in Case 8(iii) of the Bipartite SoJ algorithm, Sec. 7.5.

We have a CC $\mathfrak{X}$ on $V_1 \cup V_2$ with two color classes of vertices, V_1 and V_2. The restriction of $\mathfrak{X}$ to V_i is $\mathfrak{X}_i$. We denote by $\mathfrak{X}_3$ the colored bipartite graph $(V_2, V_1; R_j \mid R_j \subseteq V_2 \times V_1)$. (This bipartite graph is oriented from V_2 to V_1, in deviation from the notation in [Q2].) $\mathfrak{X}_3$ arises by refining the coloring of a nontrivial bipartite graph and is therefore not monochromatic. We highlight this key fact.

Assumption 3.1. $\mathfrak{X}_3$ is **not monochomatic**, i. e., it has more than one color.

By "dominant color" among a set of pairs is a color that belongs to more than half of the pairs concerned.

The algorithm works with a threshold parameter $3/4 \leq \alpha < 1$; it is assumed that $n_2 \leq \alpha n_1$.

A *good coloring* of V_1 is a coloring where every vertex-color class has relative size $\leq \alpha$. A *good equipartition* of V_1 is a nontrivial equipartition (partition into equal parts) of a "large color class," i. e., a vertex-color class of relative size greater than α. Our goal is to find a canonical good coloring of V_1 or a canonical good partition of V_1, or find a canonically embedded Johnson graph on a large vertex-color class in V_1.

Observation 3.2. *We may assume $\mathfrak{X}_1$ is a UPCC.*

Proof. If $\mathfrak{X}_1$ is not homogeneous, we either have a good coloring of V_1 or we remove the complement of the dominant vertex-color and start over with updated α (Step 5 in Sec. 7.5 [Q2]). If $\mathfrak{X}_1$ is homogeneous but imprimitive, we have a canonical equipartition, goal achieved, exit. We need to rule out that $\mathfrak{X}_1$ is a clique configuration. This would make the neighborhood hypergraph on V_1 of each bipartite graph $(V_1, V_2; R_i^-)$ (where R_i is a color in $\mathfrak{X}_3$) a BIBD having n_1 vertices and n_2 hyperedges, so by Fisher's inequality $n_2 \geq n_1$, contrary our standing assumption. □

Note. All we actually need is that $\mathfrak{X}_1$ is homogeneous.

We state the defining assumption of the case under consideration (8(iii)).

Assumption 3.3. $\mathfrak{X}_2$ is a UPCC.

Our inductive goal is to obtain a canonical nontrivial semiregular bipartite graph $X' = (V_2', V_1; E')$ where $|V_2'| \leq n_2/2$. Significant progress occurs if either we reach one of our goals on V_1 or we find such a canonical bipartite graph.

Theorem 3.4. *Under Assumption 3.3, we can make significant progress (either obtain a good coloring of V_1 or reduce V_2 by half) at a multiplicative cost of n_2.*

The proof is based on the following lemma, the sole technical contribution of this note.

Lemma 3.5. *Let $R_i \subseteq V_2 \times V_1$ be a color class in $\mathfrak{X}_3$. Consider the bipartite graph $L = (V_2, V_1; R_i)$. Let $j \in \{1, 2\}$. Then there are no L-twins in V_j.*

Proof. Suppose for a contradictions that $x, y \in V_j$ are L-twins ($x \neq y$). Let $c(x, y) = k$. Then, by Prop. 2.5, all pairs of color k are twins. This means the twin equivalence relation includes the transitive closure of R_k. But $\mathfrak{X}_j$ is primitive, so R_k is connected, so this transitive closure is $V_j \times V_j$. Hence all of V_j is a twin equivalence class. But then, by Prop. 2.6 it follows that L is the complete bipartite graph, i. e., $R_i = V_2 \times V_1$, i. e., $\mathfrak{X}_3$ is monochromatic, contrary to Assumption 3.1. □

We now describe the algorithm that justifies Theorem 3.4.

Case 1. There is no dominant color in $\mathfrak{X}_3$.

In this case, fixing any $x \in V_2$ (multiplicative cost n_2) produces a good coloring of V_1, goal achieved, exit.

Case 2. Color R_3 is dominant in $\mathfrak{X}_3$.

Henceforth let $L = (V_2, V_1; R_3)$ (a bipartite graph).

Let us fix a vertex $x \in V_2$ (multiplicative cost n_2).

Let $W = R_3(x)$. So $|W| > n_1/2$.

If $|W| \leq \alpha n_1$, we have a good coloring, exit.

Assume now that $|W| > \alpha n_1$.

Let R_2 be a non-dominant color in $\mathfrak{X}_2$. Let $U = R_2(x)$, so $|U| < n_2/2$.

Let us consider the bipartite graph $H = (U, W; R_3 \cap (U \times W))$.

Claim 3.6. *L is a nontrivial semiregular bipartite graph.*

Proof. Semiregularity follows from Prop. 2.3. L is nonempty because for any $u \in U$ we have $|R_3(u)| > n_1/2$ and therefore $R_3(u)$ intersects W.

We claim that L is not complete. Indeed suppose it is. Then for any $u \in U$ we have $R_3(u) \supseteq W$. But $|R_3(u)| = |R_3(x)| = |W|$, hence $R_3(u) = W$. This makes x and u twins in the bipartite graph $(V_2, V_2; R_3)$, contradicting Lemma 3.5. □

So now we replace our bipartite graph by L. This is significant progress because $|U| < n_2/2$.

This completes the proof of Theorem 3.4.