Title
Small Spectral Gap in the Combinatorial Laplacian Implies Hamiltonian

Authors
Butler, Steve
Chung, Fan

Peer reviewed

Small Spectral Gap in the Combinatorial Laplacian Implies Hamiltonian

Steve Butler[1] and Fan Chung[2]

[1]Department of Mathematics, University of California, Los Angeles, CA 90095-1555, USA
 sbutler@math.ucsd.edu

[2]Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112, USA
 fan@math.ucsd.edu

Abstract. We consider the spectral and algorithmic aspects of the problem of finding a Hamiltonian cycle in a graph. We show that a sufficient condition for a graph being Hamiltonian is that the nontrivial eigenvalues of the combinatorial Laplacian are sufficiently close to the average degree of the graph. An algorithm is given for the problem of finding a Hamiltonian cycle in graphs with bounded spectral gaps which has complexity of order $n^{c\ln n}$.

Keywords: Hamiltonian, combinatorial Laplacian, spectral graph theory

1. Introduction

The problem of determining whether a graph is Hamiltonian (i.e., possesses a cycle visiting each vertex exactly once) is one of the most difficult classical problems in graph theory. While large numbers of sufficient conditions have been discovered they mostly involve relatively dense graphs. As an example, Dirac's Theorem states that if the degree at each vertex is at least $|V(G)|/2$ then the graph is Hamiltonian.

A breakthrough in the study of graphs which are Hamiltonian occurred thirty years ago when Komlós and Szemerédi [8] showed that almost surely a random graph is Hamiltonian. The technique involves the rotation of paths attributed to Posa [10]. The underlying idea is that if there is a path of maximal length then edges from the endpoints of the path must connect to the center of the path and can be used to *rotate* the path to create a new maximal path. This process of rotation can be repeated and we form a large collection of maximal paths and then with high probability one of these paths can be closed to form a cycle. Since almost surely the random graph is connected this in turn implies that in fact the path must go through all of the vertices (i.e., otherwise we could extend the path).

More recently Krivelevich and Sudakov [9] considered the problem of showing that a d-regular graph is Hamiltonian. They showed that if σ, the second largest absolute value of an eigenvalue of the adjacency matrix of a d-regular graph, satisfies

$$\sigma \le c\frac{(\log\log n)^2}{\log n(\log\log\log n)}d$$

for a constant c and n sufficiently large, then G is Hamiltonian. Their underlying approach was similar to that of Komlós and Szemerédi in that graphs with small σ are "random-like", that is have many of the same properties as one would expect a random graph to have.

In their paper, Krivelevich and Sudakov stated that "it seems plausible ... that our techniques can be used to prove Hamiltonicity of almost regular graphs (i.e., graphs in which all degrees are very close to the average degree)". Here we show how to establish such a result by using their approach with some discrepancy and discrete isoperimetric inequalities of graphs found in Chung [4]. Our main result is as follows.

Theorem 1.1. *Let G be a graph on n vertices with average degree d and $0 = \lambda_0 \le \lambda_1 \le \lambda_2 \le \cdots \le \lambda_{n-1}$ be the eigenvalues of the combinatorial Laplacian of G. There is a constant c so that if*

$$|d - \lambda_i| \le c\frac{(\log\log n)^2}{\log n(\log\log\log n)}d,$$

for $i \ne 0$ and n sufficiently large, then G is Hamiltonian.

NOTE. When G is regular the eigenvalues of the adjacency matrix and the combinatorial Laplacian ($L = D - A$, where D is the diagonal degree matrix and A is the adjacency matrix) are related by a reflection and a shift of d. From this it can be seen that Theorem 1.1 implies the result of Krivelevich and Sudakov. When G is not regular the eigenvalues of the two matrices are no longer related in a trivial manner and it becomes more important which spectrum is used. We have chosen the combinatorial Laplacian in that the spectrum is useful in controlling neighborhood expansion which plays an important role in the proof.

The problem of finding a Hamiltonian cycle (or the traveling salesman algorithms, TSP in short) is known to be NP-complete (Karp [7]). Currently there are implementations for the TSP algorithm for instances with up to 13000 cities (see Applegate, Bixby, Chvátal, and Cook [2] and Woeginger [11]). The exact TSP algorithm of Held and Karp [6] is based on dynamic programming with time complexity $O(c^n)$ for some $c < 2$. For the Euclidean TSP, Chang, Hwang, and Lee [3] have a sub-exponential algorithm with some constant $c > 1$.

In Section 4 we give a subexponential algorithm for finding a Hamiltonian cycle in graphs that satisfy the required spectral bounds. The correctness of the algorithm follows from the proof of Theorem 1.1.

2. Discrepancy Inequalities Using the Combinatorial Laplacian

To prove Hamiltonicity we first need to translate the constraints placed on the spectrum into constraints on how sets expand, i.e., discrepancy bounds. Krivelevich and

Sudakov [9] used discrepancy inequalities which apply for regular graphs (see Alon and Spencer [1, Chapter 9]), namely, for G a graph on n vertices, regular of degree d, with σ the largest absolute value of a nontrivial eigenvalue of the adjacency matrix, and if $e(X, Y)$ is the number of edges with one endpoint in X and one in Y while $e(X)$ is the number of edges with both endpoints in X for $X, Y \subseteq V$ then

$$\left| e(X, Y) - \frac{d}{n}|X||Y| \right| \leq \frac{\sigma}{n}\sqrt{|X|(n-|X|)|Y|(n-|Y|)}, \text{ and } \left| e(X) - \frac{d}{2n}|X|^2 \right| < \frac{\sigma}{2}|X|.$$

For a more general case we need a stronger result. The following are discrepancy bounds adopted from [4]. For completeness we also include a short proof.

Theorem 2.1. *Suppose that a graph G with n vertices has average degree d and the eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_{n-1}$ of the combinatorial Laplacian satisfy $|d - \lambda_i| \leq \theta$ for $i \neq 0$. Then for any two disjoint subsets, X and Y, of vertices in G*

$$\left| e(X, Y) - \frac{d}{n}|X||Y| \right| \leq \frac{\theta}{n}\sqrt{|X|(n-|X|)|Y|(n-|Y|)}, \tag{2.1}$$

and

$$\left| e(X) - \frac{d|X|(|X|-1)}{2n} \right| \leq \frac{\theta}{n}|X|(n-|X|/2). \tag{2.2}$$

Remark 2.2. If we consider the case $X = \{v\}$ and $Y = V \setminus \{v\}$ then (2.1) implies that $d(n-1)/n - \theta \leq d_v \leq d + \theta$. So when θ is small compared to d (as is the case in Theorem 1.1) the degrees fall in a small range and so the graph can be described as almost-regular.

Proof. Let φ_i be the orthonormal eigenvectors corresponding to the eigenvalues λ_i. Note that $\varphi_0 = \mathbf{1}/\sqrt{n}$ where $\mathbf{1}$ is the all 1s vector. We now use the eigenvectors to decompose the characteristic vectors ψ_X and ψ_Y by $\psi_X = \sum_i a_i \varphi_i$, and $\psi_Y = \sum_i b_i \varphi_i$. We then have

$$e(X, Y) = \langle \psi_X, (dI + A)\psi_Y \rangle$$

$$= d\langle \psi_X, \psi_Y \rangle + \langle \psi_X, (D - L)\psi_Y \rangle$$

$$= d\sum_i a_i b_i - \sum_{i \neq 0} \lambda_i a_i b_i$$

$$= da_0 b_0 - \sum_{i \neq 0}(d - \lambda_i)a_i b_i.$$

Since $a_0 = \langle \psi_X, \varphi_0 \rangle = |X|/\sqrt{n}$ and $b_0 = \langle \psi_Y, \varphi_0 \rangle = |Y|/\sqrt{n}$ we have upon rearranging

$$\left| e(X, Y) - \frac{d}{n}|X||Y| \right| = \left| \sum_{i \neq 0}(d - \lambda_i)a_i b_i \right|$$

$$\leq \theta \sum_{i \neq 0}|a_i||b_i|$$

$$\leq \theta \sqrt{\sum_{i\neq 0} a_i^2 \sum_{j\neq 0} b_j^2}$$

$$= \theta \sqrt{(|X| - |X|^2/n)(|Y| - |Y|^2/n)}$$

$$= \frac{\theta}{n} \sqrt{|X|(n - |X|)|Y|(n - |Y|)},$$

establishing (2.1).

Let $x = |X|$ and let $X' \subset X$ denote a set with $|X'| = \lfloor x/2 \rfloor$. Since

$$\sum_{\substack{X' \subseteq X \\ |X'| = x'}} e(X', X \setminus X') = \binom{x-2}{x'-1} e(X, X),$$

we have

$$\binom{x-2}{x'-1}\left| e(X,X) - \frac{d}{n}x(x-1) \right| = \left| \sum_{\substack{X' \subseteq X \\ |X'| = x'}} e\left(X', X \setminus X'\right) - \binom{x-2}{x'-1}\frac{d}{n}x(x-1) \right|$$

$$= \left| \sum_{\substack{X' \subseteq X \\ |X'| = x'}} \left(e\left(X', X \setminus X'\right) - \frac{d}{n}\left\lfloor \frac{x}{2} \right\rfloor \left\lceil \frac{x}{2} \right\rceil \right) \right|$$

$$\leq \binom{x}{x'}\frac{\theta}{n}\sqrt{x'(n-x')(x-x')(n-x+x')},$$

where the last step follows by $\binom{x}{x'}$ applications of (2.1). Rearranging we have

$$\left| e(X,X) - \frac{d}{n}x(x-1) \right| = \frac{\binom{x}{x'}}{\binom{x-2}{x'-1}}\frac{\theta}{n}\sqrt{x'(n-x')(x-x')(n-x+x')} \leq \frac{2\theta}{n}x(n-x/2),$$

since $e(X,X) = 2e(X)$ this establishes (2.2). ∎

2.1. Some Properties of Edge Expansion

We use Theorem 2.1 in the proof of Theorem 1.1 in two ways. The first is to show that we have good neighborhood expansion and the second is to show that the graph is connected. The following corollary contains the equivalent of [9, Propositions 2.1-2.5]. We let $N(X) = \{v \in V \setminus X : v \sim x \in X\}$ denote the neighborhood of X.

Corollary 2.3. *Let G be a graph on n vertices with average degree d and $|d - \lambda_i| \leq \theta$ for the non-trivial eigenvalues of the combinatorial Laplacian. Further we will assume that $8\theta < d$ and that $X, Y \subseteq V$. Then the following hold:*

(a) *if $|X| \leq \theta n/d$ then $e(X) \leq \frac{3}{2}\theta|X|$.*

(b) if $|X| \leq \dfrac{\theta^2 n}{d^2}$ then $|N(X)| > \dfrac{(d/2 - 4\theta)^2}{3\theta^2}|X|$.

(c) if $|X| > \theta^2 n/d^2$ then $|N(X)| > \dfrac{n}{2} - |X|$.

(d) if $X \cap Y = \emptyset$ and $e(X, Y) = 0$, then $|X||Y| < \dfrac{\theta^2 n^2}{d^2}$.

(e) G is connected.

Proof. For (a) we use (2.2) and the assumptions to get

$$e(X) \leq \frac{d|X|(|X| - 1)}{2n} + \frac{\theta}{n}|X|(n - |X|/2) \leq \frac{\theta|X|}{2} + \theta|X| = \frac{3}{2}\theta|X|.$$

For (b) we use (a), i.e., $e(X) \leq \frac{3}{2}\theta|X|$, and the remark following Theorem 2.1 to get

$$e(X, N(X)) = \sum_{x \in X} d_x - 2e(X) \geq (d(n-1)/n - \theta)|X| - 3\theta|X| > (d/2 - 4\theta)|X|.$$

Now we use (2.1) to conclude

$$(d/2 - 4\theta)|X| < e(X, N(X)) < \frac{d}{n}|X||N(X)| + \theta\sqrt{|X||N(X)|}. \qquad (2.3)$$

If $N(X) \leq \dfrac{(d/2 - 4\theta)^2}{3\theta^2}|X|$ then we would have

$$\frac{d}{n}|X||N(X)| + \theta\sqrt{|X||N(X)|} \leq \frac{d(d/2 - 4\theta)^2|X|^2}{3\theta^2 n} + \frac{\theta|X|(d/2 - 4\theta)}{\sqrt{3\theta}}$$

$$\leq \frac{\theta^2 n}{d^2} \frac{(d/2 - 4\theta)^2 d|X|}{3\theta^2 n} + \frac{(d/2 - 4\theta)|X|}{\sqrt{3}}$$

$$= \frac{(d/2 - 4\theta)^2|X|}{3d} + \frac{(d/2 - 4\theta)|X|}{\sqrt{3}}$$

$$< \frac{(d/2 - 4\theta)|X|}{3} + \frac{(d/2 - 4\theta)|X|}{\sqrt{3}}$$

$$< (d/2 - 4\theta)|X|,$$

using that $8\theta < d$ in going to the last line, which is a contradiction to (2.3), establishing (b).

For (c) we let $Y = V \setminus (X \cup N(X))$ and use (2.1) to get

$$0 = e(X, Y) \geq \frac{d|X||Y|}{n} - \theta\sqrt{|X||Y|\left(1 - \frac{|Y|}{n}\right)},$$

which upon rearranging gives

$$\frac{|Y|}{1 - |Y|/n} \leq \frac{\theta^2 n^2}{d^2|X|} < n.$$

This implies that $|Y| < n/2$ and so $|N(X)| = n - |X| - |Y| > n/2 - |X|$.

Part (d) is again a consequence of (2.1), i.e.,

$$0 = e(X, Y) > \frac{d|X||Y|}{n} - \theta\sqrt{|X||Y|}$$

and then rearrange.

Finally, for (e), if G is disconnected then G has a connected component X of size $|X| < n/2$. Since $N(X) = \emptyset$, it follows from part (c) that $|X| \leq \theta^2 n/d^2$, but this now contradicts (b). ∎

3. Proof of Theorem 1.1

The remainder of the proof of Theorem 1.1 now proceeds as in [9] with small changes to various constants. We will give a detailed outline of the proof and refer the interested reader to [9] for details. A brief outline of the proof is as follows: Start with a path in the graph of maximal length and consider "rotations" of this path in G, some rotation of the path can be closed to form a cycle which by maximality of the length and G being connected implies that the graph is Hamiltonian.

Let $P = (v_1, v_2, \ldots, v_m)$ be a path of maximal length in G. If v_i is adjacent to v_m (abbreviated $v_i \sim v_m$) for some i then another path of maximal length is given by $P' = (v_1, \ldots, v_i, v_m, v_{m-1}, \ldots, v_{i+1})$. We say that P' is a rotation of P with fixed endpoint v_1, pivot v_i and broken edge $v_i \sim v_{i+1}$. We can now pivot P' to get P'' and so on. For $t \geq 0$ we let $S_t = \{v \in \{v_2, \ldots, v_m\} : v$ is the endpoint of a path obtainable from P by at most t rotations with fixed endpoint v_1 and all broken edges in $P\}$. It is easy to see that $S_t \subseteq S_{t+1}$ and by maximality of P that all edges incident to a vertex in S_t have their second endpoint in P.

By examining subsets S_t it can be shown that the sets have good expansion, namely, we have the following.

Lemma 3.1. ([9, Proposition 3.1]) *For* $t \geq 0$, $|S_{t+1}| \geq \frac{1}{2}|N(S_t)| - \frac{3}{2}|S_t|$.

Now let
$$t_0 = \left\lceil \frac{\log n - 2\log(d/\theta)}{2\log(d/\theta) - 10} \right\rceil + 2.$$

By Corollary 2.3(b), as long as $|S_t| \leq \theta^2 n/d^2$ then $|N(S_t)| \geq (d/2 - 4\theta)^2|S_t|/3\theta^2$. In particular, if d is large compared to θ (say $d > 1000\theta$) then it follows that when $|S_t| \leq \theta^2 n/d^2$,

$$\frac{|S_{t+1}|}{|S_t|} \geq \frac{(d/2 - 4\theta)^2}{7\theta^2}.$$

In particular, after at most $t_0 - 2$ steps we must have that $|S_t| > \theta^2 n/d^2$. By Corollary 2.3(c) when we take one more step we will have

$$|S_{t+1}| \geq \frac{1}{2}|N(S_t)| - \frac{3}{2}|S_t| \geq \frac{1}{2}\left(\frac{n}{2} - |S_t|\right) - \frac{3}{2}|S_t| = \frac{n}{4} - 2|S_t| \geq \frac{n}{4} - 2|S_{t+1}|,$$

implying $|S_{t+1}| \geq n/12$. Applying Corollary 2.3(d) we then have that $|N(S_{t+1})| = n - o(n)$. Applying Corollary 2.3(c) we can similarly conclude that

$$|S_{t+2}| \geq \frac{1}{2}(n - o(n)) - \frac{3}{2}|S_{t+2}|,$$

and thus $|S_{t+2}| \geq (1 - o(1))n/5 > n/6$ for n sufficiently large.

Let $B(v_1) = S_{t_0}$ and $A_0 = B(v_1) \cup \{v_1\}$. For each $v \in B(v_1)$ we can repeat the above argument to get a set $B(v)$, $|B(v)| > n/6$, of endpoints of maximum length paths with endpoint v. Note that each endpoint in $B(v)$ was obtained by at most $2t_0$ rotations of P. Also note that since each vertex in $B(v_1)$ is in P, we see that P must contain some constant fraction of the vertices. We now have that for each $a \in A_0$, $b \in B(a)$ there is a maximal length path $P(a,b)$ which is obtainable from P by at most $\rho = 2t_0$ rotations.

We return to the initial path P and direct it and divide it into 2ρ disjoint segments $I_1, \ldots, I_{2\rho}$ each of length at least $\lfloor n/12\rho \rfloor$. Since each path $P(a,b)$ is obtainable from P by at most ρ rotations there are at least ρ of the segments untouched, we will call these segments unbroken in $P(a,b)$. Each unbroken segment has two orientations, an absolute orientation from P and a relative orientation from $P(a,b)$ (where we direct that path from a to b).

Let $k = 2\max\{1, \lceil 500\theta\rho/d \rceil\}$. We consider sequences $\sigma = I_{i_1}, \ldots, I_{i_k}$ of k unbroken segments of P which occur in this order in $P(a,b)$ where σ specifies not only the order of segments in $P(a,b)$ but also their relative orientation. We say then that $P(a,b)$ contains σ. Note that as $P(a,b)$ has at least ρ unbroken segments I_{i_j}, $P(a,b)$ contains at least $\binom{\rho}{k}$ sequences σ. For a given σ we denote by $L(\sigma)$ the set of all pairs $a \in A_0$, $b \in B(a)$ for which the path $P(a,b)$ contains σ.

The total number of possible sequences σ is at most $(2\rho)_k 2^k$. Therefore by averaging we obtain that there exists a sequence σ_0 for which

$$|L(\sigma_0)| \geq \frac{n^2}{36} \frac{\binom{\rho}{k}}{(2\rho)_k 2^k} > \frac{n^2}{36} \left(\frac{\rho - k}{2\rho - k} \right)^k \frac{1}{k! 2^k}.$$

When $k \leq \rho/2$ then $(\rho - k)/(2\rho - k) \geq 1/3$ and it follows that there exists a sequence σ_0 for which $|L(\sigma_0)| \geq n^2/(k! 6^{k+2})$. We fix such a sequence and let $\alpha = \frac{1}{(k! 6^{k+2})}$. Let $\hat{A} = \{a \in A_0 : L(\sigma_0) \text{ contains at least } \alpha n/2 \text{ pairs with } a \text{ as the first element}\}$. Then $|\hat{A}| \geq \alpha n/2$. For each $a \in \hat{A}$, let $\hat{B}(a) = \{b \in B(A) : (a,b) \in L(\sigma_0)\}$. The definition of $\hat{A}$ guarantees that $|\hat{B}(a)| \geq \alpha n/2$.

Let C_1 be the union of the first $k/2$ segments of σ_0, in the fixed order and with the fixed relative orientation in which they occur along any of the paths $P(a,b)$, $(a,b) \in L(\sigma_0)$. Let C_2 be the union of the last $k/2$ segments of σ_0. Notice that

$$|C_i| \geq \frac{k}{2} \left\lfloor \frac{n}{12\rho} \right\rfloor \geq \frac{500\theta\rho}{d} \left\lfloor \frac{n}{12\rho} \right\rfloor > \frac{40\theta n}{d}. \tag{3.1}$$

Given a path P_0 and a set $S \subseteq \{v_1, \ldots, v_m\}$, a vertex $v \in S$ is called an interior point of S with respect to P if both neighbors of v along P are in S. The set of all interior points of S will be denoted by $\text{int}(S)$. Using (2.1) along with (3.1) we can get the following result.

Lemma 3.2. ([9, Proposition 3.2]) *The set C_1 contains a subset C_1' with $|\text{int}(C_1')| \geq nk/(48\rho)$ so that every vertex $v \in C_1'$ has at least 14θ neighbors in $\text{int}(C_1')$. A similar statement holds for C_2.*

We now fix the sets C_1' and C_2'. Using Corollary 2.3(d) it can then be shown that the terminal vertices connect back to the interior of the half paths. We remark that it is at this step where our bound for the spectral gap is needed.

Lemma 3.3. ([9, Propositions 3.3 and 3.4]) *There is a vertex $\hat{a} \in \hat{A}$ connected by an edge to $\text{int}(C_1')$. Similarly there is a vertex $\hat{b} \in \hat{B}(\hat{a})$ connected by an edge to $\text{int}(C_2')$.*

Let x be a vertex separating C_1' and C_2' along $P(\hat{a}, \hat{b})$. This splits the path into two half paths P_1 and P_2 obtained by splitting $P(\hat{a}, \hat{b})$ at x. We now fix x and rotate each of the P_i insuring that we only use vertices in their respective $\text{int}(C_i')$ as pivots. If we can show that the corresponding sets of endpoints are sufficiently large then Corollary 2.3(d) will show that there is an edge between V_1 and V_2. This edge closes a path of maximal length to a cycle. As G is connected, by Corollary 2.3(e), any non-Hamilton cycle can be extended to a path covering some additional vertices. Therefore the assumption about the maximality of P_0 implies that P_0 is a Hamilton path, and thus the above created cycle is Hamilton as well.

So now consider P_1. Let $T_i = \{v \in C_1' \setminus \{x\} : v$ is the endpoint of a path obtainable from P_1 by i rotations with fixed endpoint x, all pivots in $\text{int}(C_1')$ and all broken edges in $P_1\}$. Examining the sets T_i it can be shown using (2.1) that when $|T_i|$ is small it contains a subset which expands well, in particular, within $\log(\lambda n/d)$ steps we can have $|T_i| \geq \lambda n/d$.

Lemma 3.4. ([9, Proposition 3.5]) *There exists an i for which $|T_i| \geq \lambda n/d$.*

Hence, the set V_1 of endpoints of all rotations of P_1 has cardinality $|V_1| \geq \lambda n/d$. By Lemma 3.4, $\hat{b}$ has a neighbor in $\text{int}(C_2')$, the same argument can be carried out for P_2 to show that the set V_2 of endpoints of its rotations has at least $\lambda n/d$ vertices as well. Then by Corollary 2.3(d) there is an edge connecting V_1 and V_2 and thus closing the cycle. This completes the proof of Theorem 1.1.

4. An Algorithm for Finding Hamiltonian Cycles

The proof given above is algorithmic. Namely, it gives a procedure to find a Hamiltonian cycle in a graph which is guaranteed to be successful when the spectral gap is sufficiently small. The procedure can be described as follows:

Algorithm:

Step 0: Start with a graph G and an initial long path P (such a path can be found by starting at a single vertex and extending the path as far as possible in both directions).

Step 1: Fix one end of P and look at the rotations of P to form the sets $S_t = \{v \in V(P) : v$ is the endpoint of a path obtainable from P by at most t rotations with fixed endpoint and all broken edges in $P\}$ for $1 \leq t \leq t_0 = \left\lceil \frac{\log n - 2\log(d/\theta)}{2\log(d/\theta) - 10} \right\rceil + 2$ as described in Section 3. Then for each vertex in each S_t we similarly rotate. We now have a large collection of rotations of our initial long path each of which differs from P by

at most $2t_0$ rotations. Note if the vertex of an endpoint on any rotation is adjacent to some vertex off the path, we extend the path P to a path P' which is at least one vertex longer and start Step 1 over.

Step 2: We divide up the initial path into $4t_0$ segments and find a sequence σ_0 of $k = 2\max\{1, \lceil 500\theta\rho/d \rceil\}$ segments with orientation so that the number of rotated paths preserving σ_0 is $c'n^2$ for $c' = 1/(k!6^{k+2})$. Using this orientation we will find a path P' so that we can fix a central point on the path and rotate the two halves so that the respective collection of endpoints will contain many vertices on the graph.

Step 3: If any of these endpoints are adjacent to vertices not in the original path we can again extend the path to a new path P'' which is at least one vertex longer and then return to Step 1. Otherwise, there must be some edge connecting two endpoints of some path. In this case we have formed a cycle. If the cycle does not contain all vertices, then since the graph is connected we can find a vertex not on the original path adjacent to the cycle and thus find a longer path than what we started with, and return to Step 1. Otherwise, we have a cycle containing all of the vertices, i.e., we have found a Hamiltonian cycle and we stop the procedure.

The existence of such a σ_0 in Step 2 follows from an averaging argument and the fact that we are dealing with many rotations which in turn follows from Lemma 3.1. In Step 3 we used Lemma 3.3 to guarantee that the set of respective endpoints of the two half-paths is large and then used properties of edge discrepancy to close the cycle. We then have a cycle which we can either use to find a longer path and repeat the process or the cycle is Hamiltonian and we are done.

Examining the steps we note that given an initial path we can describe a sequence of rotations by listing the series of broken edges. In particular, since our path has some fraction of the edges it takes at most n^{2t_o} to describe our sets of rotations. Since t_0 is of order $\ln n$ this will take order $n^{c\ln n}$. For Step 2 when we consider the possible σ_0 there are of order $t_0^{t_0}$ such σ_0, and in the worst case we have to compare every σ_0 with every rotated path. In Step 3 we again rotate half-paths which will again have order of complexity bounded by $n^{c\ln n}$. In the worst case we have to repeat the argument n times and so this will give us one last factor of n. So in total we have computational complexity of order $n^{c\ln n}$.

References

1. Alon, N., Spencer, J.H.: The Probabilistic Method, 2nd Ed. Wiley, New York (2000)
2. Applegate, D., Bixby, R., Chvátal, V., Cook, W.: On the solution of traveling salesman problems. Doc. Math. 3, 645–656 (1998)
3. Chang, H.-W., Hwang, F.K., Lee, J.S.: The searching over separators strategy to solve some NP-hard problems in subexponential time. Algorithmica 9, 398–423 (1993)
4. Chung, F.: Discrete isoperimetric inequalities. In: Surveys in Differential Geometry, vol. IX, pp. 53–82. International Press, Somerville (2004)
5. Dirac, G.A.: Hamiltonian circuits and long circuits. Ann. Discrete Math. 3, 75–92 (1978)
6. Held, M., Karp, R.M.: The traveling-salesman problem and minimum spanning trees. Operation Res. 18, 1138–1162 (1970)
7. Karp, R.M.: Reduciblity among combinatorial problems. In: Miller, R.E., Thatcher, J.W. (eds.) Comoplexity of Computer Computations, pp. 85–103. Plenum, New York (1972)

8. Komlós, J., Szemerédi, E.: Hamilton cycles in random graphs. In: Infinite and finite sets,
    vol. II, pp. 1003–1010. North-Holland, Amsterdam (1975)

9. Krivelevich, M., Sudakov, B.: Sparse pseudo-random graphs are Hamiltonian. J. Graph
    Theory 42, 17–33 (2003)

10. Pósa, L.: Hamiltonian circuits in random graphs. Discrete Math. 14, 359–364 (1976)

11. Woeginger, G.J.: Exact algorithms for NP-hard problems: a survey. In: Combinatorial
    Optimization—Eureka, You shrink!, pp. 185–207. Springer, New York (2003)