

On the integrality gap of the subtour LP for the 1,2-TSP

Jiawei Qian · Frans Schalekamp ·
David P. Williamson · Anke van Zuylen

Received: 14 June 2012 / Accepted: 24 October 2014 / Published online: 12 November 2014
© Springer-Verlag Berlin Heidelberg and Mathematical Optimization Society 2014

Abstract In this paper, we study the integrality gap of the subtour LP relaxation for the traveling salesman problem in the special case when all edge costs are either 1 or 2. For the general case of symmetric costs that obey triangle inequality, a famous conjecture is that the integrality gap is $4/3$. Little progress towards resolving this conjecture has been made in 30 years. We conjecture that when all edge costs $c_{ij} \in \{1, 2\}$, the integrality gap is $10/9$. We show that this conjecture is true when the optimal subtour LP solution has a certain structure. Under a weaker assumption, which is an analog of a recent conjecture by Schalekamp et al., we show that the integrality gap is at most $7/6$. When we do not make any assumptions on the structure of the optimal subtour LP

A preliminary version of this paper appeared in LATIN 2012: Theoretical Informatics [19]. Some results in the current paper are stronger, using the same techniques as in [19]. The first author was supported in part by NSF grant CCF-1115256. This work was carried out in part while the third author was on sabbatical at TU Berlin, and the fourth author was at Max-Planck-Institut für Informatik in Saarbrücken, Germany. The third author was supported in part by the Berlin Mathematical School, the Alexander von Humboldt Foundation, and NSF Grant CCF-1115256.

J. Qian
Bank of America Merrill Lynch, Chicago, IL 60661, USA
e-mail: jq35@cornell.edu

F. Schalekamp · A. van Zuylen (✉)
Department of Mathematics, College of William and Mary, Williamsburg, VA 23185, USA
e-mail: anke@wm.edu

F. Schalekamp
e-mail: frans@wm.edu

D. P. Williamson
School of Operations Research and Information Engineering, Cornell University,
Ithaca, NY 14853, USA
e-mail: dpw@cs.cornell.edu

solution, we can show that integrality gap is at most $5/4$; this is the first bound on the integrality gap of the subtour LP strictly less than $4/3$ known for an interesting special case of the TSP. We show computationally that the integrality gap is at most $10/9$ for all instances with at most 12 cities.

Keywords Traveling salesman problem · Subtour elimination · Linear programming · Integrality gap

Mathematics Subject Classification 90C05 · 90C27 · 05C70

1 Introduction

The traveling salesman problem (TSP) is one of the most well studied problems in combinatorial optimization. Given a set of cities $\{1, 2, \dots, n\}$, and distances $c(i, j)$ for traveling from city i to j , the goal is to find a tour of minimum length that visits each city exactly once. An important special case of the TSP is the case when the distance forms a metric, i.e., $c(i, j) \leq c(i, k) + c(k, j)$ for all i, j, k , and all distances are symmetric, i.e., $c(i, j) = c(j, i)$ for all i, j . The symmetric TSP is known to be NP-hard, even if $c(i, j) \in \{1, 2\}$ for all i, j [18]; note that such instances trivially obey the triangle inequality. Such instances are also known to be APX-hard; that is, there is no α -approximation algorithm for the problem for some $\alpha > 1$ unless $P = NP$.

The metric TSP can be approximated to within a factor of $\frac{3}{2}$ using an algorithm by Christofides [7] from 1976. The algorithm combines a minimum spanning tree with a matching on the odd-degree nodes to get an Eulerian graph that can be shortcut to a tour; the analysis shows that the minimum spanning tree and the matching cost no more than the optimal tour and half the optimal tour respectively. Better results are known for several special cases, but, surprisingly, no progress has been made on approximating the general symmetric TSP in more than 30 years. A natural direction for trying to obtain better approximation algorithms is to use linear programming. The following linear programming relaxation of the traveling salesman problem was used by Dantzig et al. [8]. For simplicity of notation, we let $G = (V, E)$ be a complete undirected graph on n nodes. In the LP relaxation, we have a variable $x(e)$ for all $e = (i, j)$ that denotes whether we travel directly between cities i and j on our tour. Let $c(e) = c(i, j)$, and let $\delta(S)$ denote the set of all edges with exactly one endpoint in $S \subseteq V$. Then the relaxation is

$$\begin{aligned} & \text{Min } \sum_{e \in E} c(e)x(e) \\ (\text{SUBT}) \quad & \text{subject to: } \sum_{e \in \delta(i)} x(e) = 2, \quad \forall i \in V, & (1) \\ & \sum_{e \in \delta(S)} x(e) \geq 2, \quad \forall S \subset V, \quad 3 \leq |S| \leq |V| - 3, & (2) \\ & 0 \leq x(e) \leq 1, \quad \forall e \in E. & (3) \end{aligned}$$

The first set of constraints (1) are called the *degree constraints*. The second set of constraints (2) are sometimes called *subtour elimination constraints* or sometimes just *subtour constraints*, since they prevent solutions in which there is a subtour of just the nodes in S . As a result, the linear program is sometimes called the *subtour LP*. It has been shown by Wolsey [24] (and later Shmoys and Williamson [22]) that Christofides' algorithm finds a tour of length at most $\frac{3}{2}$ times the optimal value of the subtour LP; these proofs show that the minimum spanning tree and the matching on odd-degree nodes can be bounded above by the value of the subtour LP, and half the value of the subtour LP, respectively. This implies that the integrality gap, the worst case ratio of the length of an optimal tour divided by the optimal value of the LP, is at most $\frac{3}{2}$. However, no examples are known that show that the integrality gap can be as large as $\frac{3}{2}$; in fact, no examples are known for which the integrality gap is greater than $\frac{4}{3}$. A well known conjecture states that the integrality gap is indeed $\frac{4}{3}$; see (for example) Goemans [10].

Recently, progress has been made in several directions, both in improving the best approximation guarantee and in determining the exact integrality gap of the subtour LP for certain special cases of the symmetric TSP. In the *graph-TSP*, the costs $c(i, j)$ are equal to the shortest path distance in an underlying unweighted graph. If the graph is cubic and 3-connected, Gamarnik et al. [9] show an approximation algorithm with guarantee slightly better than $\frac{3}{2}$. Oveis Gharan et al. [17] show that the graph-TSP can be approximated to within $\frac{3}{2} - \epsilon$ for a small constant $\epsilon > 0$ for all graphs. Boyd et al. [6], and Aggarwal et al. [1] independently give a $\frac{4}{3}$ -approximation algorithm if the underlying graph is cubic. Mömke and Svensson [15] improve these results by giving a 1.461-approximation for the graph-TSP and an $\frac{4}{3}$ -approximation algorithm if the underlying graph is subcubic. Mucha [16] improves the analysis of the Mömke–Svensson algorithm to a $\frac{13}{9}$ -approximation algorithm, and Sebő and Vygen [21] combine the ideas of Mömke and Svensson [15] with an algorithm based on a carefully chosen ear decomposition of the graph to obtain a $\frac{7}{5}$ -approximation algorithm. All of these α -approximation algorithms imply a corresponding upper bound of α on the integrality gap for the corresponding graph-TSP instances.

In Schalekamp et al. [20], three of the authors of this paper resolve a related conjecture. A 2-matching of a graph is a set of edges such that no edge appears twice and each node has degree two, i.e., it is an integer solution to the LP (*SUBT*) with only constraints (1) and (3). Note that a minimum-cost 2-matching thus provides a lower bound on the length of the optimal TSP tour. A minimum-cost 2-matching can be found in polynomial time using a reduction to a certain minimum-cost matching problem. Boyd and Carr [5] conjecture that the worst case ratio of the cost of a minimum-cost 2-matching and the optimal value of the subtour LP is at most $\frac{10}{9}$. This conjecture was proved to be true by Schalekamp et al. and examples are known that show this result is tight.

Unlike the techniques used to obtain better results for the graph-TSP, the techniques of Schalekamp et al. work on general weighted instances that are symmetric and obey the triangle inequality. However, their results only apply to 2-matchings and it is not clear how to enforce global connectivity on the solution obtained by their method. A potential direction for progress on resolving the integrality gap for the subtour LP is a conjecture by Schalekamp et al. that the worst-case integrality gap is attained for instances for which the optimal subtour LP solution is a basic

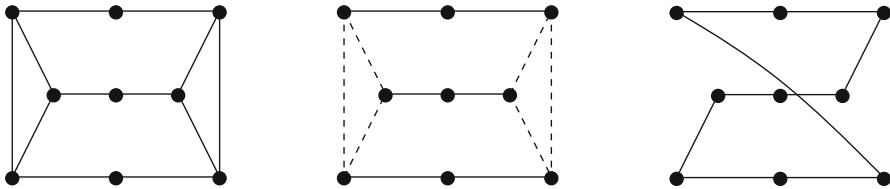


Fig. 1 Illustration of the worst example known for the integrality gap for the 1,2-TSP. The figure on the *left* shows all edges of cost 1. The figure in the *center* gives the subtour LP solution, in which the *dotted* edges have value $\frac{1}{2}$, and the *solid* edges have value 1; this is also an optimal fractional 2-matching. The figure on the *right* gives the optimal tour and the optimal 2-matching

solution to the linear program obtained by dropping the subtour elimination constraints.

In this paper, we turn our attention to the 1,2-TSP, where $c(i, j) \in \{1, 2\}$ for all i, j . Note that bounding the cost of enforcing connectivity is relatively easy in this class of problems, since we may connect any two components for an increase in cost of at most 2. Papadimitriou and Yannakakis [18] show how to approximate 1,2-TSP within a factor of $\frac{11}{9}$ by computing a minimum-cost 2-matching and merging its cycles into a tour. In addition, they show a ratio of $\frac{7}{6}$ if they start with a minimum-cost 2-matching that has no cycles of length 3. Bläser and Ram [4] improve this ratio and the best known approximation factor of $\frac{8}{7}$ is given by Berman and Karpinski [3].

We do not know a tight bound on the integrality gap of the subtour LP even in the case of the 1,2-TSP. As an upper bound, we appear to know only that the gap is at most $\frac{3}{2}$ via Wolsey's result. There is an easy nine city example showing that the gap must be at least $\frac{10}{9}$; see Fig. 1. This example has been extended to a class of instances on $9k$ nodes for any positive integer k by Williamson [23]. The contribution of this paper is to begin a study of the integrality gap of the 1,2-TSP, and to improve our state of knowledge for the subtour LP in this case. We prove an upper bound on the integrality gap for the subtour LP of $\frac{5}{4}$, which is the first bound on the integrality gap with value less than $\frac{4}{3}$ for a natural class of TSP instances. Under an analog of a conjecture of Schalekamp et al. [20], we show that the integrality gap is at most $\frac{7}{6}$, and with an additional assumption on the structure of the solution, we can improve this bound to $\frac{10}{9}$. We describe these results in more detail below.

All the known approximation algorithms since the initial work of Papadimitriou and Yannakakis [18] on the problem start by computing a minimum-cost 2-matching. However, the example of Fig. 1 shows that an optimal 2-matching can be as much as $\frac{10}{9}$ times the value of the subtour LP for the 1,2-TSP, so we cannot directly replace the bound on the optimal solution in these approximation algorithms with the subtour LP in the same way that Wolsey did with Christofides' algorithm in the general case. Using the result of Schalekamp et al. [20] and a new lemma that relates part of the analysis of Papadimitriou and Yannakakis [18] to the subtour LP bound, we obtain a preliminary upper bound on the integrality gap of the subtour LP for the 1,2-TSP of $\frac{7}{9} \cdot \frac{10}{9} + \frac{4}{9} = \frac{106}{81} \approx 1.3086$.

To improve this upper bound to $\frac{5}{4}$, we first show stronger results in some cases. A fractional 2-matching is a basic optimal solution to the LP (*SUBT*) with only constraints

(1) and (3). Schalekamp et al. [20] have conjectured that the worst-case integrality gap for the subtour LP is obtained when the optimal solution to the subtour LP is an extreme point of the fractional 2-matching polytope. We show that if this is the case for 1,2-TSP then we can find a tour of cost at most $\frac{7}{6}$ the cost of the fractional 2-matching, implying that the integrality gap is at most $\frac{7}{6}$ in these cases. Next, we show that if this optimal solution to the fractional 2-matching problem has a certain structure, then we can find a tour of cost at most $\frac{10}{9}$ times the cost of the fractional 2-matching, implying an upper bound on the integrality gap of $\frac{10}{9}$ for these cases. Figure 1 shows that this result is tight.

We then use the previous arguments to show that one can construct a tour of cost at most $\frac{5}{4}$ times the subtour LP value. To do this, we prove that we can assume without loss of generality that the optimal value of the subtour LP is less than $n + 1$, where n denotes the number of nodes. Combined with a more careful analysis based on the results obtained before, we obtain our main result. The results above all lead to polynomial-time algorithms, though we do not state the exact running times.

Finally, we perform computational experiments to show that the integrality gap is at most $\frac{10}{9}$ for $n \leq 12$. We conjecture that the integrality gap is in fact exactly $\frac{10}{9}$.

We note that the upper bound on the integrality gap for general 1,2-TSP instances presented in this paper is stronger than the bound that appeared in a preliminary version of this paper [19] of $\frac{19}{15}$. In the time between publication of the preliminary version and the current revision, Mnich and Mömke [14] obtained an upper bound of $\frac{5}{4}$ on the integrality gap for 1,2-TSP instances that have the additional property of being “fractionally Hamiltonian”, which means that the optimal objective value of the subtour LP is equal to the number of nodes in the instance. In this version, using the same techniques as in the preliminary version, we show an unconditional upper bound on the integrality gap of $\frac{5}{4}$, and a bound of $\frac{26}{21}$ for fractionally Hamiltonian instances.

The remainder of this paper is structured as follows. Section 2 contains preliminaries and a first general bound on the integrality gap for the 1,2-TSP. We show how to obtain stronger bounds if the optimal subtour LP solution is a fractional 2-matching in Sect. 3. In Sect. 4, we combine the arguments from the previous sections and show that the integrality gap without any assumptions on the structure of the subtour LP solution is at most $\frac{5}{4}$. We describe our computational experiments in Sect. 5. Finally, we close with a conjecture on the integrality gap of the subtour LP for the 1,2-TSP in Sect. 6.

2 Preliminaries and a first bound on the integrality gap

We will work extensively with 2-matchings and fractional 2-matchings; that is, extreme points x of the LP (*SUBT*) with only constraints (1) and (3), where in the first case the solutions are required to be integer. For convenience we will abbreviate “fractional 2-matching” by F2M and “2-matching” by 2M. The basic solutions of the F2M polytope have the following well-known structure (attributed to Balinski [2]). Each connected component of the support graph (that is, the edges e for which $x(e) > 0$) is either a cycle on at least three nodes with $x(e) = 1$ for all edges e in the cycle, or consists of odd-sized cycles with $x(e) = \frac{1}{2}$ for all edges e in the cycle connected by paths of edges e with $x(e) = 1$ for each edge e in the path (the center figure in Fig. 1 is an example). We call the former components *integer components* and the latter *frac-*

tional components. In a fractional component, we call a path of edges e with $x(e) = 1$ a 1 -*path*. The edges e with $x(e) = \frac{1}{2}$ in cycles are called *cycle edges*. An F2M with a single component is called *connected*, and we call a component 2 -*connected* if, in the graph induced by the component, the sum of the x -values on the edges crossing any cut is at least 2 . We let n denote the number of nodes in an instance.

As mentioned in the introduction, Schalekamp et al. [20] have shown the following.

Theorem 1 (Schalekamp et al. [20]) *If edge costs obey the triangle inequality, then the cost of an optimal 2-matching is at most $\frac{10}{9}$ times the value of the subtour LP.*

It is not hard to show that this immediately implies an upper bound of $\frac{13}{9}$ on the integrality gap of the subtour LP for the 1,2-TSP: we can just compute a minimum cost 2-matching at a cost of $\frac{10}{9}$ the value of the subtour LP, remove the most expensive edge from each cycle, which gives a collection of node-disjoint paths, and add edges of cost 2 to combine these paths into a tour. Each cycle has at least three edges; at worst, we remove an edge of cost 1 from each cycle and then need an edge of cost 2 to patch the paths into a tour. Thus the overall cost increases by at most $\frac{1}{3}n$, giving a tour, the cost of which can be bounded by $\frac{10}{9} + \frac{1}{3} = \frac{13}{9}$ times the value of the subtour LP, because n is a lower bound on the value of the subtour LP.

The algorithm of Papadimitriou and Yannakakis [18] improves on this idea, by cleverly merging the cycles of the optimal 2M solution. We summarize the properties of their algorithm that we will use. First, observe that we can assume without loss of generality that the optimal 2M solution consists of a number of cycles with only edges of cost 1 (“pure” cycles) and at most one cycle which has one or more edges of cost 2 (the “non-pure” cycle), by deleting the edges of cost 2 and combining the resulting disjoint paths into a single cycle. Moreover, if i is a node in the non-pure cycle which is incident on an edge of cost 2 in the cycle, then there can be no edge of cost 1 connecting i to a node in a pure cycle (since otherwise, we can merge the non-pure cycle with a pure cycle without increasing the cost).

The Papadimitriou–Yannakakis algorithm solves the following bipartite matching problem: On one side we have a node for every pure cycle, and on the other side, we have a node for every node in the instance. There is an edge from pure cycle C to node i , if $i \notin C$ and there is an edge of cost 1 from i to some node in C . Let r be the number of pure cycles that are unmatched in a maximum cardinality bipartite matching. Papadimitriou and Yannakakis show how to “patch together” the matched cycles. We refer the reader to their original paper [18] for more details. The resulting cycles are then combined into a tour of cost at most

$$\frac{7}{9}OPT(2M) + \frac{4}{9}n + \frac{1}{3}r, \quad (4)$$

where $OPT(2M)$ is the cost of an optimal 2M solution.¹

¹ In [18], $OPT(2M)$ is expressed as $n+k$, where k is the number of edges of cost 2 in the optimal 2M solution. The number of unmatched pure cycles is denoted by r_2 . The bound given by [18] is $n+k + \frac{2}{9}(n-n_2-k) + r_2$, where n_2 is a quantity that is lower bounded by $3r_2$. Therefore, the bound in [18] can be upper bounded by $\frac{7}{9}(n+k) + \frac{4}{9}n + \frac{1}{3}r_2$.

We now show how to convert this bound into a bound in terms of the optimal value to *SUBT*.

Lemma 1 *Let r be the number of pure cycles that are unmatched in a maximum cardinality bipartite matching instance defined by Papadimitriou and Yannakakis. Then*

$$OPT(SUBT) \geq n + r.$$

Proof We note that for a bipartite matching instance, the size of the minimum cardinality vertex cover is equal to the size of the maximum matching. We use this fact to construct a feasible dual solution to the subtour LP that has value $n + r$. Let \mathcal{C}_M, V_M be the pure cycles and nodes (in the original graph), for which the corresponding nodes in the bipartite matching instance are in the minimum cardinality vertex cover. The dual of the subtour LP (*SUBT*) is

$$\begin{aligned} & \text{Max } 2 \sum_{S \subset V} y(S) + 2 \sum_{i \in V} y(i) - \sum_{e \in E} z(e) \\ (D) \text{ subject to: } & \sum_{S \subset V: e \in \delta(S)} y(S) + y(i) + y(j) - z(e) \leq c(e), \quad \forall e = (i, j), \\ & y(S) \geq 0, \quad \forall S \subset V, 3 \leq |S| \leq n - 3, \\ & z(e) \geq 0, \quad \forall e \in E. \end{aligned}$$

We set $z(e) = 0$ for each $e \in E$, and we set $y(i) = \frac{1}{2}$ for each $i \in V \setminus V_M$. For a pure cycle on a set of nodes C , we set $y(C) = \frac{1}{2}$, if the cycle is not in \mathcal{C}_M . The dual objective for this solution is exactly $n + r$: its value is n plus the number of pure cycles minus the size of the vertex cover, or n plus the number of pure cycles minus the size of the matching, since the vertex cover has the same size as the matching. Thus it is the same as n plus the number of pure cycles not in the matching, or $n + r$.

It remains to show that the dual constructed is feasible. Define the *load* on an edge $e = (i, j)$ of solution (y, z) to be $\sum_{S \subset V: e \in \delta(S)} y(S) + y(i) + y(j) - z(e)$. For any edge $e = (i, j)$ of cost 1 inside a cycle of the 2M, the load on the edge is at most 1, since the only potentially non-zero dual variables loading the edge are the dual variables $y(i)$ and $y(j)$. For an edge (i, j) where $i \in C$ and $j \in C' \neq C$, the load is $y(i) + y(j) + y(C) + y(C') \leq 2$. Suppose (i, j) has cost 1, and the cycles C and C' are both pure cycles. Then the edge occurs twice in the bipartite matching instance (namely, once going from i to C and once going from j to C') and hence the dual of at least two of the four objects i, j, C and C' has been reduced to 0. The total load on edge (i, j) is thus at most 1. Now, suppose C' is the non-pure cycle, then $y_{C'} = 0$, since we only increased the dual variables for the pure cycles. Moreover, at least one endpoint of the (j, C) edge in the bipartite matching instance must be in the vertex cover, so the load on edge (i, j) is again at most 1. □

Note that, combined with (4) and Theorem 1, Lemma 1 implies that the cost of the tour is at most $\frac{7}{9} \cdot \frac{10}{9} OPT(SUBT) + \frac{4}{9} OPT(SUBT) = \frac{106}{81} OPT(SUBT)$. This bound obtained on the integrality gap seems rather weak, as the best known lower bound

on the integrality gap is only $\frac{10}{9}$. Schalekamp et al. [20] have conjectured that the integrality gap (or worst-case ratio) of the subtour LP occurs when the solution to the subtour LP is a fractional 2-matching.

Conjecture 1 (Schalekamp et al. [20]) *Let α_n be the integrality gap of the subtour LP on n vertices. Then there exists an instance which has an optimal subtour LP solution that is an F2M and for which the optimal tour has cost at least α_n times the subtour LP cost.*

In the next section, we show that we can obtain better bounds on the integrality gap of the subtour LP in the case that the optimal solution is a fractional 2-matching. In Sect. 4, we then show how to combine Lemma 1 with the bounds in the next section to obtain an upper bound of $\frac{5}{4}$ on the integrality gap.

3 Better bounds if the optimal solution is an F2M

If the optimal solution to the subtour LP is a fractional 2-matching, then a natural approach to obtaining a good tour is to start with the edges with x -value 1, and add as many edges of cost 1 and x -value $\frac{1}{2}$ as possible, without creating a cycle on a subset of the nodes. More precisely, we will propose an algorithm that creates an acyclic spanning subgraph (V, T) where all nodes have degree one or two. We will refer to an acyclic spanning subgraph in which all nodes have degree one or two as a partial tour. A partial tour can be extended to a tour by adding $d/2$ edges of cost 2, where d is the number of degree 1 nodes. The cost of the tour is $c(T) + d$, where $c(T) = \sum_{e \in T} c(e)$.

We will use the following lemma.

Lemma 2 *Let $G = (V, T)$ be a partial tour. Let A be a set of edges not in T that form an odd cycle or a path on $V' \subset V$, where the nodes in V' have degree one in T . We can find $A' \subset A$ such that $(V, T \cup A')$ is a partial tour, and*

- $|A'| \geq \frac{1}{3}|A|$ if A is a cycle,
- $|A'| \geq \frac{1}{3}(|A| - 1)$ if A is a path,

We postpone the proof of the lemma and first prove the implication for the bound on the integrality gap if the optimal subtour LP solution is a fractional 2-matching.

Theorem 2 *There exists a tour of cost at most $\frac{7}{6}$ times the cost of a connected F2M solution if $c(i, j) \in \{1, 2\}$ for all i, j .*

Proof Let $P = \{e \in E : x(e) = 1\}$ (the edges in the 1-paths of x). We will start the algorithm with $T = P$. Let $R = \{e \in E : x(e) = \frac{1}{2} \text{ and } c(e) = 1\}$ (the edges of cost 1 in the cycles of x). Note that the connected components of the graph (V, R) consist of paths and odd cycles. The main idea is that we consider these components one by one, and use Lemma 2 to show that we can add a large number of the edges of each path and cycle, where we keep as an invariant that T is a partial tour. Note that by Lemma 2, the number of edges added from each path or cycle A is at least $|A|/3$, except for the paths for which $|A| \equiv 1 \pmod{3}$. Let \mathcal{P}_1 be this set of paths. We would like to claim that we add a third of the edges on average from each component,

and we therefore preprocess the paths in \mathcal{P}_1 , where we add one edge (either the first or last edge from each path in \mathcal{P}_1) to T if this is possible without creating a cycle in T , and if so, we remove this edge and its neighboring edge in R (if any) from R . After the preprocessing, we use Lemma 2 to process each of the components in (V, R) .

We call a path A in \mathcal{P}_1 “eared” if the 1-paths that are incident on the first and last node of the path are such that they go between two neighboring nodes of A . Without loss of generality we will assume for each path in \mathcal{P}_1 that is not eared, that the first edge on the path does not form a cycle with a 1-path. It is not hard to see that we can add the first edge from at least half of the paths in \mathcal{P}_1 that are not eared: Suppose we simply add the first edge from each path in \mathcal{P}_1 that is not eared to T , regardless of whether it makes a cycle in T or not. Then, each cycle in T contains at least two edges that came from a path in \mathcal{P}_1 , and hence, removing one such edge per cycle leaves a partial tour T that contains the first edge from at least half of the paths that are not eared in \mathcal{P}_1 .

After preprocessing the paths in \mathcal{P}_1 , we iterate through the connected components in (V, R) and add edges to T while maintaining that T is a partial tour. By Lemma 2, the number of edges added from each path or cycle A is at least $|A|/3$, except for the paths in \mathcal{P}_1 . We now consider two cases for the paths in \mathcal{P}_1 , depending on whether we added an edge from the path to T in the preprocessing step or not. Note that for a path A in \mathcal{P}_1 for which we added an edge to T in the preprocessing step, R contains a path of $|A| - 2$ edges after the preprocessing step, and by Lemma 2, we add at least $(|A| - 2 - 1)/3$ of these to T . Together with the edge added in the preprocessing step, we thus add at least $1 + (|A| - 2 - 1)/3 = |A|/3$ edges. For a path in \mathcal{P}_1 for which we did not add an edge to T in the preprocessing stage, we add at least $(|A| - 1)/3$ edges. Now, recall that a path A in \mathcal{P}_1 has $|A| \equiv 1 \pmod{3}$, and that the number of edges added is an integer, so in the first case, the number of edges added is at least $|A|/3 + \frac{2}{3}$ and in the second case it is $|A|/3 - \frac{1}{3}$. Let z be the number of eared paths in \mathcal{P}_1 . Then, the number of paths in \mathcal{P}_1 that are in the second case is at most z plus the number of paths in \mathcal{P}_1 that fall in the first case. Hence, the total number of edges from R that were added to T can be lower bounded by $\frac{1}{3}|R| - \frac{1}{3}z$. We now give an upper bound on the number of nodes of degree one in T .

Let k be the number of cycle nodes in x , i.e. $k = \#\{i \in V : x(i, j) = \frac{1}{2} \text{ for some } j \in V\}$, and let p be the number of cycle edges of cost 2 in x , i.e. $p = \#\{e \in E : x(e) = \frac{1}{2} \text{ and } c(e) = 2\}$. Note that (V, R) contains p paths (which may have zero edges) on the cycle nodes, and hence $p \geq z$. Initially, when T contains only the edges in the 1-paths, all k nodes have degree one, and there are $k - p$ edges in R . We argued that we added at least $\frac{1}{3}|R| - \frac{1}{3}z = \frac{1}{3}k - \frac{1}{3}p - \frac{1}{3}z$ edges to T . Each edge reduces the number of nodes of degree one by two, and hence, the number of nodes of degree one at the end of the algorithm is at most $k - 2(\frac{1}{3}k - \frac{1}{3}p - \frac{1}{3}z) = \frac{1}{3}k + \frac{2}{3}p + \frac{2}{3}z$. Recall that $c(P)$ denotes the cost of the 1-paths, and the total cost of T at the end of the algorithm is at most $c(P) + \frac{1}{3}k - \frac{1}{3}p - \frac{1}{3}z$. Since at most $\frac{1}{3}k + \frac{2}{3}p + \frac{2}{3}z$ nodes have degree one in T , we can extend T into a tour of cost at most $c(P) + \frac{2}{3}k + \frac{1}{3}p + \frac{1}{3}z$.

The cost of the solution x can be expressed as $c(P) + \frac{1}{2}k + \frac{1}{2}p$. Note that each 1-path connects two cycle nodes, hence $c(P) \geq \frac{1}{2}k$. Moreover, an eared path A is incident to one (if $|A| = 1$) or two (if $|A| > 1$) 1-paths of length two, since the support graph of x is simple. Therefore we can lower bound $c(P)$ by $\frac{1}{2}k + z$. Therefore,

$$\frac{7}{6} (c(P) + \frac{1}{2}k + \frac{1}{2}p) \geq c(P) + \frac{1}{12}k + \frac{1}{6}z + \frac{7}{12}k + \frac{7}{12}p \geq c(P) + \frac{2}{3}k + \frac{1}{3}z + \frac{1}{3}p,$$

where $p \geq z$ is used in the last inequality. \square

Proof of Lemma 2 The basic idea behind the proof of the lemma is the following: We go through the edges of A in order, and try to add them to T if this does not create a cycle or node of degree three in T . If we cannot add an edge, we simply skip the edge and continue to the next edge. Since the edges in T form a collection of disjoint paths and each node in A has degree one in T , we can always add either the first edge or the second edge of A : if the first edge cannot be added, then adding it to T must create a cycle, and since the edges in T form a collection of node disjoint paths, adding the second edge of the path or cycle to T cannot create a cycle. Similarly, we need to skip at most two edges between two edges that are successfully added to T : first, an edge is skipped because otherwise we create a node of degree three in T , and if a second edge is skipped, then this must be because adding that edge to T would create a cycle. But then, adding the next edge on the path cannot create a cycle in T .

To lower bound the number of edges from we can add from each path or cycle A , we partition the edges into groups of two or three consecutive edges. For a path A , the first group contains the first two edges, and each subsequent group contains the next three edges. The final group contains the last zero, one or two edges of the path. For each group except the last group, at least one edge is added to T . Hence, we can conclude that we can add at least $(|A| - 4)/3$ from the groups of size three, and 1 for the first group, for a total of $(|A| - 1)/3$ edges, where $|A|$ denotes the number of edges in A . For a cycle A , we need to be slightly more careful, since the argument that we can add at least one edge from the last group of size three does not hold if the very first edge was added to T (since it may be the case that the first and third edge of the group cannot be added without creating a node of degree three, and the second edge of the group cannot be added without creating a cycle). Therefore, we let the first group contain two consecutive edges, where the *second* edge is the edge that was the first to be added to T . By the same argument as for the path, we can thus conclude that we can add at least $(|A| - 1)/3$ edges.

We now show that by being a little more careful, we can in fact add $|A|/3$ edges if A is a cycle. Note that the number of nodes in A is odd, and hence there must be some j such that the path in T that starts in u_j ends in some node $v \notin A$. We claim that if we consider the edges in A starting with either edge $\{u_{j-1}, u_j\}$ or edge $\{u_j, u_{j+1}\}$, we are guaranteed that for at least one of these starting points, we can add both the first and the third edge to T .

Clearly, neither $\{u_{j-1}, u_j\}$ nor $\{u_j, u_{j+1}\}$ can create a cycle if we add it to T . So suppose that $T \cup \{u_{j-1}, u_j\} \cup \{u_{j+1}, u_{j+2}\}$ contains a cycle. This cycle does not contain the node u_j , because the path in T that starts in u_j ends in some node $v \notin C$. Hence T contains a path that starts in u_{j+1} and ends in u_{j+2} . But then $T \cup \{u_j, u_{j+1}\} \cup \{u_{j+2}, u_{j+3}\}$ does not have a cycle, since if it did, T must have a path starting in u_{j+2} and ending in u_{j+3} which is only possible if $u_{j+1} = u_{j+3}$. Since the number of nodes in A is at least three, this is not possible. \square

We remark that the ratio of $\frac{7}{6}$ in Theorem 2 is achieved if every 1-path contains just one edge of cost 1, and all cycle edges have cost 1. However, in such a case, we could find another optimal F2M solution of the same cost, which has fewer cycle edges: If

we have a 1-path of cost 1 with endpoints in two different odd cycles of edges with $x(e) = \frac{1}{2}$, we can obtain the alternative solution by removing the 1-path, and increasing the x -value on the four cycle edges incident on its endpoints to 1, and then alternating between setting the x -value to 0 and 1 around the cycles. Now, since the cycles are odd, the degree constraints are again satisfied. The objective value does not increase because we only change the x -value on edges of cost 1. For a 1-path of cost 1 with endpoints in the same odd cycle, the cycle gives us two paths between the endpoints, one of odd length and one of even length. We can alternate increasing and decreasing the x -value by $\frac{1}{2}$ on the odd-length path and finally decrease the x -value for the 1-path to $\frac{1}{2}$, to obtain a new F2M solution of the same cost with fewer cycle edges. We note that these modifications may increase the number of components of the F2M solution.

This motivates the following definition. We call an F2M solution *canonical*, if all edges in the support have cost 1 and all 1-paths contain at least two edges. If a canonical F2M solution is connected, we can improve the analysis in Theorem 2 to show the following.

Theorem 3 *There exists a tour of cost at most $\frac{10}{9}$ times the cost of a connected canonical F2M solution if $c(i, j) \in \{1, 2\}$ for all i, j .*

Proof We adapt the final paragraph of the proof of Theorem 2. As before, the cost of the tour is at most $c(P) + \frac{2}{3}k + \frac{1}{3}p + \frac{1}{3}z$. However, since all cycle edges have cost 1, $p = 0$ and $z = 0$. The cost of the tour is thus at most $c(P) + \frac{2}{3}k$.

The cost of the F2M solution is $c(P) + \frac{1}{2}k$. Since each cycle node is the endpoint of a 1-path and vice versa, the number of 1-paths is $k/2$. By the fact that x is canonical, each of these 1-paths has cost at least two, so we get that $c(P) \geq k$. The proof is concluded by noting that then $\frac{10}{9}(c(P) + \frac{1}{2}k) \geq c(P) + \frac{1}{9}k + \frac{10}{9} \cdot \frac{1}{2}k = c(P) + \frac{2}{3}k$. □

4 An upper bound of $\frac{5}{4}$ on the integrality gap

We now show how to use the results in the previous two sections to obtain an upper bound of $\frac{5}{4}$ on the integrality gap for the general case. In addition, we show that if all edges in the support of the optimal subtour LP solution have cost 1, then the integrality gap is at most $\frac{26}{21}$.

We will bound the integrality gap of the solution obtained by the Papadimitriou–Yannakakis algorithm, by (i) bounding the difference between the cost of the 2M and the subtour LP, and (ii) bounding the difference between the 2M solution and the tour constructed from it by the Papadimitriou–Yannakakis algorithm.

As in the Papadimitriou–Yannakakis algorithm described in Sect. 2, we call a cycle in a 2M a “pure” cycle if all its edges have cost 1, and a “non-pure” cycle otherwise. The idea behind this section is to show that the quantity in (i) can be “charged” to the nodes in the non-pure cycle only, and that the quantity in (ii) can be “charged” mainly to the nodes in the pure cycles.

We first state the following lemma, which formalizes the second statement.

Lemma 3 *If $OPT(SUBT) < n + 1$, then the difference between the cost of the $2M$ used and the tour constructed by the Papadimitriou–Yannakakis algorithm can be upper bounded by $\alpha n_{\text{pure}} + \beta(n_{\text{non-pure}} - \ell)$, where n_{pure} is the number of nodes in pure cycles in the $2M$, $n_{\text{non-pure}}$ is the number of nodes in the non-pure cycle, and ℓ is the number of edges of cost 2 in the non-pure cycle, for any values of α, β so that $9\alpha \geq 2$ and $3\alpha + 2\beta \geq 1$.*

Note that Lemma 1 and the assumption that $OPT(SUBT) < n + 1$ imply that the Papadimitriou–Yannakakis algorithm finds a bipartite matching that matches all the pure cycles. A careful look at the analysis of Papadimitriou and Yannakakis [18] then shows that their algorithm finds a tour which satisfies the lemma. The details basically follow the analysis of Papadimitriou and Yannakakis, and are therefore postponed to Appendix 1.

The key observation in this section is that we can indeed restrict our attention to instances with $OPT(SUBT) < n + 1$, the requirement of Lemma 3.

Lemma 4 *The worst-case integrality gap is attained on an instance with subtour LP value less than $n + 1$, where n is the number of nodes in the instance.*

Proof Consider an instance I on n nodes for which the ratio between the length of the optimal tour and the subtour LP value is γ , and suppose $OPT(SUBT) = n + k$ for some $k \geq 1$. We construct an instance I' with $n' = n + 1$ nodes, for which the ratio between the length of the optimal tour and the subtour LP value is at least γ , and the optimal value of the subtour LP is at most $n + k = n' + k - 1$. Repeatedly applying this procedure proves the lemma.

If $OPT(SUBT) = n + k$, then the subtour LP solution on I has a total x -value of k on edges of cost 2, since the objective value is equal to $\sum_{e \in E} x(e) + \sum_{e \in E: c(e)=2} x(e)$, and $\sum_{e \in E} x(e) = \frac{1}{2} \sum_{v \in V} \sum_{e \in \delta(v)} x(e) = n$. We fix an optimal subtour solution x , and we construct I' from I , by adding one node i , and adding edges (i, j) of cost 1, for every j in I that is incident on an edge e with $c(e) = 2$ and $x(e) > 0$. All other edges incident on i get cost 2. Note that the optimal tour on I' has length at least the length of the optimal tour on I , since we can take a tour on I' and shortcut i to obtain a tour on I . On the other hand, we can use x to define a feasible solution on I' , by “rerouting” one unit in total from edges $e = (j, k)$ with $c(e) = 2$ to the edges (j, i) and (i, k) . Since the cost of this solution on I' is the same as the cost of x , the ratio between the length of the optimal tour and the subtour LP value has not decreased. \square

Remark 1 We note that the proof of Lemma 4 implies that to compute integrality gaps or approximation guarantees, we may assume without loss of generality that an instance has an optimal subtour LP value of at most $n + 1$, where n is the number of nodes in the instance. If this does not hold, we may add nodes as in the proof of Lemma 4 without increasing $OPT(SUBT)$, and a tour of cost C on the extended instance can be shortcut to a tour on the original instance of cost at most C .

Theorem 4 *The integrality gap of the subtour LP is at most $\frac{5}{4}$ for the 1,2-TSP, and it is at most $\frac{26}{21}$ for 1,2-TSP instances for which $OPT(SUBT) < n + \frac{1}{2}$, where n is the number of nodes in the instance.*

Proof By Lemma 4, we can assume without loss of generality that $OPT(SUBT) < n + 1$. To compute a tour, we first drop the subtour elimination constraints and find an optimal F2M solution. Since the F2M problem is a relaxation of the subtour LP, and it is half-integral, its objective value is either $n + \frac{1}{2}$ or n . Furthermore, note that in order to bound the integrality gap, we may assume that every edge that is not in the support of the optimal F2M solution has cost 2.

We first consider the case $OPT(SUBT) < n + \frac{1}{2}$, in which case the optimal F2M solution has objective value n . Since all edges in the support of the F2M solution have cost 1, we may assume by the arguments preceding Theorem 3 that all 1-paths contain at least two edges of cost 1; in other words, we may assume the components of the F2M solution are canonical. By applying Theorem 3 we convert each fractional component of the F2M solution into a cycle on the nodes in the component.

Note that each cycle that is the result of applying Theorem 3 contains at least one edge that is not in the optimal F2M solution, and thus it has at least one edge of cost 2. By the observation of Papadimitriou and Yannakakis [18], we may merge these into a single non-pure cycle. The integer components of the F2M solution are pure cycles, since the support of the F2M solution only contains edges of cost 1. We let n_{pure} be the number of nodes in the pure cycles (or, equivalently, in the integer components of the F2M solution), and let $n_{\text{non-pure}}$ be the number of nodes in the non-pure cycle (or, equivalently, the number of nodes in the fractional components of the F2M solution). Let ℓ be the number of cost-2 edges in the computed 2-matching. Then the cost of the computed 2-matching is $n + \ell = n_{\text{pure}} + n_{\text{non-pure}} + \ell$, by the assumption that the optimal F2M solution has objective value n . By Theorem 3, $n_{\text{non-pure}} + \ell \leq \frac{10}{9}n_{\text{non-pure}}$, i.e., $\ell \leq \frac{1}{9}n_{\text{non-pure}}$.

If we apply the Papadimitriou–Yannakakis algorithm to this 2-matching, this increases the cost by at most $\alpha n_{\text{pure}} + \beta(n_{\text{non-pure}} - \ell)$, provided that $9\alpha \geq 2$ and $3\alpha + 2\beta \geq 1$ by Lemma 3. Choosing $\alpha = \frac{5}{21}$, $\beta = \frac{1}{7}$, we thus find that the total cost of the tour is at most $n + \ell + \frac{5}{21}n_{\text{pure}} + \frac{1}{7}n_{\text{non-pure}} - \frac{1}{7}\ell \leq n + \frac{5}{21}n_{\text{pure}} + (\frac{1}{7} + \frac{6}{7} \cdot \frac{1}{9})n_{\text{non-pure}} = (1 + \frac{5}{21})n$, where we used the fact that $\ell \leq \frac{1}{9}n_{\text{non-pure}}$.

If $n + \frac{1}{2} \leq OPT(SUBT) < n + 1$, the optimal F2M solution has cost at most $n + \frac{1}{2}$. We temporarily decrease the cost of the unique cost-2 edge in the F2M to 1, and follow the same procedure as above, to find a 2-matching. Let $n_{\text{non-pure}}$ be the number of nodes in the non-pure cycle, and note that $n_{\text{non-pure}}$ is at least 9, since a fractional component of a canonical F2M solution contains at least two odd cycles, containing at least six nodes, and at least three 1-paths, containing at least one additional node each.

Let the cost of the computed 2-matching (with respect to the true costs) be $n + \ell$; in other words, the procedure from Theorem 3 added $\ell - 1$ edges of cost 2 in addition to the single cost-2 edge in the F2M. By Theorem 3, $\ell - 1 \leq \frac{1}{9}n_{\text{non-pure}}$. As in the case when $OPT(SUBT) < n + \frac{1}{2}$, we apply the Papadimitriou–Yannakakis algorithm to this 2-matching, and by Lemma 3 this increases the cost by at most $\alpha n_{\text{pure}} + \beta(n_{\text{non-pure}} - \ell)$. We now choose $\alpha = \frac{1}{4}$, $\beta = \frac{1}{8}$, to get that the total cost of the tour is at most $n + \ell + \frac{1}{4}n_{\text{pure}} + \frac{1}{8}n_{\text{non-pure}} - \frac{1}{8}\ell = n + \frac{1}{4}n_{\text{pure}} + \frac{9}{8}n_{\text{non-pure}} + \frac{7}{8}(\ell - 1) + \frac{7}{8}$. Now, recall that $\ell - 1 \leq \frac{1}{9}n_{\text{non-pure}}$ and that $n_{\text{non-pure}} \geq 9$, and thus $\frac{7}{8} \leq \frac{5}{8} + \frac{1}{4} \cdot \frac{1}{9}n_{\text{non-pure}}$. Hence,

we can upper bound the cost of the tour by $n + \frac{1}{4}n_{\text{pure}} + (\frac{9}{8} + \frac{7}{8} \cdot \frac{1}{9} + \frac{1}{4} \cdot \frac{1}{9})n_{\text{non-pure}} + \frac{5}{8} = \frac{5}{4}(n + \frac{1}{2}) \leq \frac{5}{4}OPT(SUBT)$. \square

Remark 2 The bound of $\frac{5}{4}$ in Theorem 4 may be marginally improved by a more careful analysis of small instances. It appears that in order to decrease the bound to $\frac{10}{9}$, or even $\frac{11}{9}$, more substantial new ideas are needed, however.

Remark 3 The results in this section yield a polynomial time algorithm for finding a tour of cost at most $\frac{5}{4}OPT(SUBT)$: first, solve the F2M problem, and use the procedure from the proof of Lemma 4 to ensure that $OPT(SUBT) < n + 1$ by adding nodes if necessary. Then, apply the procedure from the proof of Theorem 4 to find a tour of cost at most $\frac{5}{4}OPT(SUBT)$. Finally, shortcut the nodes added by the procedure from Lemma 4 to get a tour for the original instance.

5 Computational results

In the case of the 1,2-TSP, for a fixed n we can generate all instances as follows. For each value of n , we first generate all nonisomorphic graphs on n nodes using the software package NAUTY [13]. We let the cost of edges be one for all edges in G and let the cost of all other edges be two. Then each of the generated graph G gives us an instance of 1,2-TSP problem with n nodes, and this covers all instances of the 1,2-TSP for size n up to isomorphism.

In fact, we can do slightly better by only generating biconnected graphs. We say that a graph $G = (V, E)$ is *biconnected* if it is connected and there is no vertex $v \in V$ such that removing v disconnects the graph; such a vertex v is a *cut vertex*. It is possible to show that the subtour LP value is at least $n + 1$ if G is not biconnected, hence, by Lemma 4 it suffices to consider biconnected graphs. However, the proof of Lemma 4 involves adding additional new nodes (perhaps many of them). Using a similar technique to the one in the proof of Lemma 4, one can show that given a graph on n vertices, there is a biconnected graph on at most $n + 2$ vertices that has no better ratio of optimal tour to subtour LP value. In Appendix 2 we prove two lemmas that imply the following corollary.

Corollary 1 *Let $G = (V, E)$ be the graph of cost 1 edges in a 1,2-TSP instance. Then if $G = (V, E)$ is not biconnected, there exists a biconnected $G' = (V', E')$ with $|V'| \leq |V| + 2$ such that $OPT(G)/SUBT(G) \leq OPT(G')/SUBT(G')$.*

For each instance of size n , we solve the subtour LP and the corresponding integer program using CPLEX 12.1 [12] and a Macintosh laptop computer with dual core 2GHz processor and 1GB of memory. It is known that the integrality gap is 1 for $n \leq 5$, so we only consider problems of size $n \geq 6$. The results are summarized in Table 1. For $n = 11$, the number of nonisomorphic biconnected graphs is nearly a billion and thus too large to consider, so we turn to another approach. For $n = 11$ and $n = 12$, we use the fact that we know a lower bound on the integrality gap of $\frac{n+1}{n}$, namely for the instances depicted in Fig. 2. The claimed lower bounds on the integrality gap for these instances follow readily from the integrality gap for the example in Fig. 1.

Table 1 The subtour LP integrality gap for 1,2-TSP for $6 \leq n \leq 12$, where the ratio for $6 \leq n \leq 10$ is only on biconnected graphs

n	6	7	8	9	10	11	12
Subtour IP/LP ratio	8/7.5	8/7.5	9/8.5	10/9	11/10	12/11	13/12
No. of graphs	56	468	7,123	194,066	9,743,542	900,969,091	–

The second row shows the number of nonisomorphic biconnected graphs for $6 \leq n \leq 11$

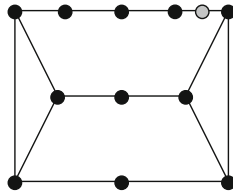


Fig. 2 Illustration of the instances with integrality gap at least $\frac{12}{11}$ for $n = 11$ (without the grey node) and $\frac{13}{12}$ for $n = 12$ (with the grey node) for the 1,2-TSP. All edges of cost 1 are shown

We then check whether this is the worst integrality gap for each vertex of subtour LP. A list of non-isomorphic vertices of the subtour LP is available for $n = 6$ to 12 at Sylvia Boyd’s website <http://www.site.uottawa.ca/~sylvia/subtourvertices>. In order to check whether the lower bound on the integrality gap is tight, we solve the following integer programming problem for each vertex x of the polytope for $n = 11$ and $n = 12$, where now the costs $c(e)$ are the decision variables, and x is fixed:

$$\begin{aligned} & \text{Max } z - \alpha_n \sum_{e \in E} c(e)x(e) \\ & \text{subject to:} \\ & \sum_{e \in T} c(e) \geq z, \quad \forall \text{ tours } T, \\ & c(e) \in \{1, 2\}, \quad \forall e \in E. \end{aligned}$$

Note that α_n is the lower bound on the integrality gap for instances of n nodes. If the objective is nonpositive for all of the vertices of the subtour LP, then we know that α_n is the integrality gap for a particular value of n .

Since the number of non-isomorphic tours of n nodes is $(n - 1)!/2$, the number of constraints is too large for CPLEX for $n = 11$ or 12. We overcome this difficulty by first solving the problem with only tours that have at least $n - 1$ edges in the support graph of the vertex x , and repeatedly adding additional violated tours. We find that the worst case integrality gap for $n = 11$ is $\frac{12}{11}$ and for $n = 12$ is $\frac{13}{12}$.

We can now observe that our overall computation leads to a bound of $\frac{10}{9}$ on the integrality gap for instances of the 1,2-TSP with $n \leq 12$. Suppose the worst-case integrality gap for these instances is attained for an instance with k vertices. If $k \leq 8$, then we know that there is a biconnected graph on at most 10 vertices with no better integrality gap, and we have determined the worst-case ratio for all biconnected graphs

on at most 10 vertices. If $k = 9$ or $k = 10$, then we know there is a biconnected graph on at most 12 vertices with no better integrality gap, and we have determined the worst-case ratio for all biconnected graphs with at most 10 vertices and all instances with 11 or 12 vertices. If $k = 11$ or $k = 12$, then we have determined the worst-case ratio for all instances with 11 or 12 vertices. Thus for any instance with $n \leq 12$, we have determined that the integrality gap is at most $\frac{10}{9}$.

6 Conjectures and conclusions

As stated in the introduction, we conjecture the following.

Conjecture 2 *The integrality gap of the subtour LP for the 1,2-TSP is $\frac{10}{9}$.*

Schalekamp et al. [20] have conjectured that to determine the integrality gap for the subtour LP, we can restrict ourselves to considering instances, which have an optimal solution that is an extreme point of the F2M polytope.

We have shown in Theorem 2 that if an analogous conjecture is true for 1,2-TSP, then the integrality gap for 1,2-TSP is at most $\frac{7}{6}$; it would be nice to show that if the analogous conjecture is true for 1,2-TSP then the integrality gap is at most $\frac{10}{9}$.

Finally, we remark that the integrality gap of the linear program obtained by adding the constraints

$$\sum_{e \in \delta(S) \setminus F} x(e) + \sum_{e \in F} (1 - x(e)) \geq 1 \quad \forall S \subset V, F \subseteq \delta(S), |F| \text{ odd},$$

to (SUBT) is at most $\frac{11}{9}$ by (4) and Lemma 1, since the 2M polytope is described by these additional constraints plus the degree constraints. It is an interesting question whether the analysis of Berman and Karpinski [3] can also be expressed in terms of the optimal value of this stronger LP.

Acknowledgments We thank Sylvia Boyd for useful and encouraging discussions. We thank two anonymous referees for helpful comments and suggestions.

Appendix 1: Proof of Lemma 3

We now prove Lemma 3 from Sect. 4.

Lemma 5 *If $OPT(SUBT) < n + 1$, then the difference between the cost of the 2M used and the tour constructed by the Papadimitriou–Yannakakis algorithm can be upper bounded by $\alpha n_{\text{pure}} + \beta(n_{\text{non-pure}} - \ell)$, where n_{pure} is the number of nodes in pure cycles in the 2M, $n_{\text{non-pure}}$ is the number of nodes in the non-pure cycle, and ℓ is the number of edges of cost 2 in the non-pure cycle, for any values of α, β so that $9\alpha \geq 2$ and $3\alpha + 2\beta \geq 1$.*

Proof Recall from Sect. 2 that the Papadimitriou–Yannakakis algorithm starts by finding a maximum cardinality bipartite matching in a graph which has a node for

each pure cycle on one side, and a node for each node in the instance on the other side. There is an edge (C, i) if $i \notin C$, and there exists some node j in C such that (i, j) is an edge of cost 1.

In Lemma 1, we show that $OPT(SUBT) \geq n + r$, where r is the number of pure cycles that are not matched in the maximum cardinality bipartite matching. Hence, the assumption that $OPT(SUBT) < n + 1$ implies that all the pure cycles are matched. In order to show that this implies the lemma, we will repeat some key parts of the algorithm and analysis of Papadimitriou and Yannakakis.

Consider the directed graph $F = (\mathcal{C}, A)$ which has a node for every cycle in the 2M, and an arc (C, C') if the maximum cardinality bipartite matching contains an edge from cycle C to a node i in cycle C' . Each node in F that corresponds to a pure cycle has outdegree 1, and the non-pure cycle (if it exists) has outdegree 0. Papadimitriou and Yannakakis show how to find a spanning subgraph of F' of F such that each nontrivial component is an in-tree of depth one or a path of length two. The only possible trivial component is the node that corresponds to the non-pure cycle. Since the non-pure cycle has outdegree 0, it can only occur in a nontrivial component as the root of an in-tree, or as the endpoint of a path of length two. It turns out that the latter does not happen in the construction described by Papadimitriou and Yannakakis, but even if it did, we could just remove the last edge in the length-two path to obtain one in-tree of depth one and one trivial component containing the non-pure cycle. Hence, we may assume the non-pure cycle only occurs in a nontrivial component as the root of an in-tree.

Papadimitriou and Yannakakis now merge the cycles in one component of F' into a single cycle containing at least one edge of cost 2 as follows: If the component is an in-tree of depth one, let C be the cycle corresponding to the root, let C_1, \dots, C_m be the remaining cycles in the component, and let v_i be the node in C such that (C_i, v_i) was in the bipartite matching. We consider the nodes in C in clockwise order, starting from a node $v \neq v_i$ for $i = 1, \dots, m$ if such a node exists, and an arbitrary node v otherwise. If we encounter two adjacent nodes v_i, v_j in $\{v_1, \dots, v_m\}$, then we merge the corresponding cycles C_i and C_j with C according to (a) in Fig. 3. Otherwise, if the current node is v_j but its clockwise neighbor is not or if its clockwise neighbor is the first node v , then we merge C_j with C as in (b) in Fig. 3. Finally, if the component is a path of length two, we merge the three cycles as in (c) in Fig. 3. Note that each cycle in the resulting graph contains at least one edge of cost 2, and hence we can find a tour of the same cost by removing the edges of cost 2, and arbitrarily connecting the resulting paths into a tour.

We now show that the number of edges of cost 2 that are added by merging cycles according to Fig. 3 can be upper bounded by $\alpha n_{\text{pure}} + \beta(n_{\text{non-pure}} - \ell)$, provided that α and β are so that $9\alpha \geq 2$ and $3\alpha + 2\beta \geq 1$.

We say a node is involved in a merging if it is either a node in one of the cycles that are fully drawn in Fig. 3, or if it is node v or v_i in subfigure (b). Note that each node is involved in at most one merging. Recall that the non-pure cycle can only occur as the root of a 1-tree of depth one or as a trivial component in F' , and hence, only the partially drawn cycle in (a) and (b) is (potentially) a non-pure cycle.

We now examine each of the cases (a), (b) and (c) in Fig. 3 in turn. In Fig. 3a one edge of cost 2 is added and we can charge this edge to the (at least) 6 nodes from pure

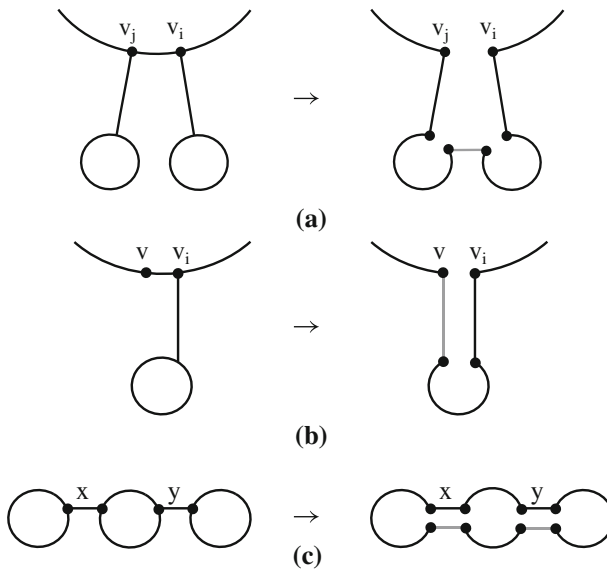


Fig. 3 The three cases for merging cycles in [18] (Figure 2 in [18]). The *black edges* indicate edges of cost 1, and the *grey edges* (potentially) have cost 2

cycles involved in this merging, as long as $6\alpha \geq 1$. This is indeed the case, because we have the stronger requirement that $9\alpha \geq 2$. In (b), again, one edge of cost 2 is added, and we can charge the edge to the (at least) three nodes of the pure cycle involved in the merging and the 2 nodes of the (potentially) non-pure cycle involved in the merging, as long as $3\alpha + 2\beta \geq 1$ (in case the 2 nodes were part of the non-pure cycle), and $5\alpha \geq 1$ (in case the 2 nodes were part of a pure cycle). Finally, in Fig. 3c, two edges of cost 2 are added; we can charge the two edges to the (at least) nine nodes from pure cycles involved in the merging as long as $9\alpha \geq 2$.

Hence, we have shown that difference in cost between the tour and the 2M can be charged to the nodes, in such a way that each node is charged at most once, and a node in a pure cycle is charged at most α and a node in a non-pure cycle is charged at most β .

Finally, we remark that a node in a non-pure cycle is charged only in case (b). Now, if (v_i, v) in Fig. 3b is an edge of cost 2, then there is no need to charge any nodes, since the cost after merging is the same as before the merge. Hence, if we direct all edges of the non-pure cycle in clockwise direction, then the head of the edges of cost 2 is never charged. The total charge to the nodes in the non-pure cycle is therefore at most $\beta(n_{\text{non-pure}} - \ell)$. \square

Appendix 2: Proof of Corollary 1

We now show that the worst-case integrality gap for the subtour LP for the 1,2-TSP can be found on graphs of cost 1 edges that are biconnected, as stated in Corollary 1 in Sect. 5. Let $OPT(G)$ and $SUBT(G)$ be the cost of the optimal tour and the value of

the subtour LP (respectively) when G is the graph of cost 1 edges. We start by proving that the worst case is obtained on a connected graph.

Lemma 6 *Let $G = (V, E)$ be the graph of cost 1 edges in a 1,2-TSP instance. Then if $G = (V, E)$ is not connected, there exists a connected graph $G' = (V', E')$ with $|V'| = |V| + 1$ such that $OPT(G)/SUBT(G) \leq OPT(G')/SUBT(G')$.*

Proof Suppose G has more than one connected component. We create $G' = (V', E')$ by adding a new vertex i^* to the graph, and adding edges from all $j \in V$ to i^* so that $V' = V \cup \{i^*\}$ and $E' = E \cup \{(i^*, j) : j \in V\}$. Given a tour of G' , we can easily produce a tour of G of no greater cost by shortcutting i^* , so that $OPT(G) \leq OPT(G')$. Let x be an optimal solution to the subtour LP for the graph G . We now define a solution x' for G' , where $x'_{ij} = x_{ij}$ if i and j are in the same connected component of G , while if i and j are in different connected components of G , then we set $x'_{ij} = 0$, $x'_{i^*i} = x_{ij}$, and $x'_{i^*j} = x_{ij}$. It is easy to see that the cost of x' is the same as that of x . We now argue that there is some solution x'' feasible for the subtour LP on G' such that its cost is no greater, so that $SUBT(G') \leq SUBT(G)$. It is clear that the bounds constraints (3) are satisfied for x' and the degree constraints (1) are satisfied for x' for all $i \in V$; however, the degree constraint for i^* may not be satisfied. Since for any component $C \subseteq V$ of G , $x(\delta(C)) \geq 2$, it is clear that $x'(\delta(i^*)) \geq 2$, but it may be the case that $x'(\delta(i^*)) > 2$. For the subtour constraints (2), consider any $S \subset V'$, $S \neq \emptyset$, such that $i^* \notin S$. Then $x'(\delta(S)) \geq x(\delta(S)) \geq 2$, and for any $S \subseteq V'$ with $i^* \in S$, $S \neq \{i^*\}$, $x'(\delta(S)) = x'(\delta(V' - S)) \geq 2$ by the previous argument. Finally, Goemans and Bertsimas [11] have shown (see also Williamson [23]) that if edge costs obey the triangle inequality, and there is some solution x' to the subtour LP in which degree constraints are exceeded but all other constraints are met, then there is another feasible solution x'' of no greater cost in which all constraints are satisfied. Hence we have that $SUBT(G') \leq SUBT(G)$. Thus we have that $OPT(G)/SUBT(G) \leq OPT(G')/SUBT(G')$. \square

Lemma 7 *Let $G = (V, E)$ be the graph of cost 1 edges in a 1,2-TSP instance. Then if $G = (V, E)$ is connected but not biconnected, there exists a biconnected $G' = (V', E')$ with $|V'| = |V| + 1$ such that $OPT(G)/SUBT(G) \leq OPT(G')/SUBT(G')$.*

Proof By hypothesis we assume that the graph $G = (V, E)$ is connected. Let i_1, \dots, i_k be all the cut vertices of G , and let C_1, \dots, C_ℓ be all the connected components formed when these vertices are removed, so that $C_1, \dots, C_\ell, \{i_1\}, \dots, \{i_k\}$ form a partition of V . We create a new graph $G' = (V', E')$ by adding a new vertex i^* , and adding edges from i^* to each vertex in $C_1 \cup \dots \cup C_\ell$, so that $V' = V \cup \{i^*\}$ and $E' = E \cup \{(i^*, j) : j \in C_p \text{ for some } p\}$. We note that G' is biconnected. As before, we have $OPT(G) \leq OPT(G')$ since given a tour of G' we can shortcut i^* to get a tour of G . Let x be an optimal subtour LP solution for graph G . We now argue, as we did in the proof of Lemma 6, that we can create an x' that costs no more than x such that all the subtour and bounds constraints are obeyed, and all degree constraints are either met or exceeded; this will imply that $SUBT(G') \leq SUBT(G)$, and complete the proof. Suppose without loss of generality that removing cut vertex i_1 creates components C_1 and $C = C_2 \cup \dots \cup C_\ell \cup \{i_2\} \cup \dots \cup \{i_k\}$, so that $C_1, \{i_1\}$,

and C partition V . We set $x'_{ij} = 0$ and $x'_{i^*i} = x'_{i^*j} = x_{ij}$ if $i \in C_1$ and $j \in C$; $x'_{ij} = x_{ij}$ otherwise. If $i \in C_1$ and $j \in C$, then $(i, j) \notin E$ since i_1 is a cut vertex, so the cost of x' is no more than that of x . The arguments that all constraints are satisfied except for the degree constraint on i^* follow as in the proof of Lemma 6. We now must argue that $x'(\delta(i^*)) \geq 2$. To do this, we show that $\sum_{i \in C_1, j \in C} x_{ij} \geq 1$. Since $x(\delta(i_1)) = 2$, it must be the case that either $\sum_{j \in C} x_{i_1j} \leq 1$ or $\sum_{j \in C_1} x_{i_1j} \leq 1$; without loss of generality we assume the former is true. Then since $x(\delta(C_1 \cup \{i_1\})) \geq 2$, and $x(\delta(C_1 \cup \{i_1\})) = \sum_{j \in C} x_{i_1j} + \sum_{i \in C_1, j \in C} x_{ij}$, it follows that $\sum_{i \in C_1, j \in C} x_{ij} \geq 1$, and the proof is complete. \square

References

1. Aggarwal, N., Garg, N., Gupta, S.: A 4/3-approximation for TSP on cubic 3-edge-connected graphs. CoRR abs/1101.5586 (2011)
2. Balinski, M.L.: Integer programming: methods, uses, computation. *Manag. Sci.* **12**, 253–313 (1965)
3. Berman, P., Karpinski, M.: 8/7-approximation algorithm for (1,2)-TSP. In: Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms, pp. 641–648 (2006)
4. Bläser, M., Shankar Ram, L.: An improved approximation algorithm for TSP with distances one and two. In: Liskiewicz, M., Reischuk, R. (eds.) *Fundamentals of Computation Theory, 15th International Symposium, FCT 2005, Lecture Notes in Computer Science*, vol. 3623, pp. 504–515. Springer (2005)
5. Boyd, S., Carr, R.: Finding low cost TSP and 2-matching solutions using certain half-integer subtour vertices. *Discrete Optim.* **8**, 525–539 (2011). Prior version available at <http://www.site.uottawa.ca/sylvia/recentpapers/halftri.pdf>. Accessed 27 June 2011
6. Boyd, S., Sitters, R., van der Ster, S., Stougie, L.: The traveling salesman problem on cubic and subcubic graphs. *Math. Program.* **144**(1–2), 227–245 (2014). A preliminary version appeared in IPCO 2011
7. Christofides, N.: Worst case analysis of a new heuristic for the traveling salesman problem. Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA (1976)
8. Dantzig, G., Fulkerson, R., Johnson, S.: Solution of a large-scale traveling-salesman problem. *Oper. Res.* **2**, 393–410 (1954)
9. Gamarnik, D., Lewenstein, M., Sviridenko, M.: An improved upper bound for the TSP in cubic 3-edge-connected graphs. *Oper. Res. Lett.* **33**(5), 467–474 (2005)
10. Goemans, M.X.: Worst-case comparison of valid inequalities for the TSP. *Math. Program.* **69**, 335–349 (1995)
11. Goemans, M.X., Bertsimas, D.J.: Survivable networks, linear programming relaxations, and the parsimonious property. *Math. Program.* **60**, 145–166 (1990)
12. IBM ILOG CPLEX 12.1 (2009)
13. McKay, B.D.: Practical graph isomorphism. *Congr. Numerantium* **30**, 45–97 (1981)
14. Mnich, M., Mömke, T.: Improved integrality gap upper bounds for TSP with distances one and two. CoRR abs/1312.2502 (2013)
15. Mömke, T., Svensson, O.: Approximating graphic TSP by matchings. In: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science, pp. 560–569 (2011)
16. Mucha, M.: $\frac{13}{9}$ -approximation for graphic TSP. *Theory Comput. Syst.*, 1–18 (2012). A preliminary version appeared in STACS 2012
17. Oveis Gharan, S., Saberi, A., Singh, M.: A randomized rounding approach to the traveling salesman problem. In: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science, pp. 550–559 (2011)
18. Papadimitriou, C.H., Yannakakis, M.: The traveling salesman problem with distances one and two. *Math. Oper. Res.* **18**, 1–11 (1993)
19. Qian, J., Schalekamp, F., Williamson, D.P., van Zuylen, A.: On the integrality gap of the subtour LP for the 1,2-TSP. In: LATIN 2012: Theoretical Informatics, 10th Latin American Symposium, Lecture Notes in Computer Science, vol. 7256, pp. 606–617 (2012)

20. Schalekamp, F., Williamson, D.P., van Zuylen, A.: 2-matchings, the traveling salesman problem, and the subtour LP: A proof of the Boyd-Carr conjecture. *Math. Oper. Res.* **39**(2), 403–417 (2014). A preliminary version appeared in SODA 2012
21. Sebő, A., Vygen, J.: Shorter tours by nicer ears: $7/5$ -approximation for graphic TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. *CoRR abs/1201.1870* (2012)
22. Shmoys, D.B., Williamson, D.P.: Analyzing the Held–Karp TSP bound: a monotonicity property with application. *Inf. Process. Lett.* **35**, 281–285 (1990)
23. Williamson, D.P.: Analysis of the Held–Karp heuristic for the traveling salesman problem. Master’s thesis, MIT, Cambridge, MA (1990). Also appears as Tech Report MIT/LCS/TR-479
24. Wolsey, L.A.: Heuristic analysis, linear programming and branch and bound. *Math. Program. Study* **13**, 121–134 (1980)