

An Asymptotic Framework for the Integrality Gap of the Traveling Salesman Problem

Janos Barta¹, Monaldo Mastrolilli^{1,3}, Eleonora Vercesi^{2,3},
Tullio Villa^{1,3*}

¹Scuola universitaria professionale della Svizzera italiana, Lugano,
Switzerland.

²Università della Svizzera italiana, Lugano, Switzerland.

³Dalle Molle Institute for Artificial Intelligence USI-SUPSI, Lugano,
Switzerland.

*Corresponding author: tullio.villa@supsi.ch;

Contributing authors: janos.barta@supsi.ch;

monaldo.mastrolilli@supsi.ch; eleonora.vercesi@usi.ch;

Abstract

The integrality gap of the subtour elimination relaxation for the Traveling Salesman Problem is a longstanding open problem, epitomized by the $4/3$ -conjecture. Understanding this gap requires a detailed analysis of the extreme points of the subtour elimination polytope. In this work, we introduce a new perspective for studying the integrality gap through what we call the *Asymptotic Framework*. Rather than focusing on individual vertices, we consider entire families of vertices and analyze their behavior in the limit. This approach allows us to capture structural properties that might otherwise remain hidden at the level of isolated instances. We develop tools for systematically constructing and analyzing such families, and we demonstrate how this framework yields new insights into the broader study of the integrality gap. As a first application of the framework, we prove the $4/3$ -conjecture for a class of instances that we call *fractional negligible*. Furthermore, we demonstrate how the framework can be used to identify infinitely many families of vertices with high integrality gap and show that such families are in some sense pervasive.

Keywords: Traveling Salesman Problem, Linear Programming Relaxation, Integrality Gap

1 Introduction

Let $K_n = (V_n, E_n)$ be the complete graph on n nodes, and let $\mathbf{c} \in \mathbb{R}^{E_n}$ be a non-negative cost vector. The *Traveling Salesman Problem* (TSP) consists of finding a Hamiltonian tour of minimum cost. In this work we only consider the *symmetric* formulation, namely, all the graphs are assumed to be undirected; accordingly, the notations ij and ji are used interchangeably to denote the edge connecting nodes i and j . We also require the costs to be *metric*, satisfying the triangle inequalities. In the remainder of the manuscript, we refer to the symmetric metric TSP simply as TSP.

The TSP is typically formulated as an Integer Linear Program (ILP). Among the most prominent formulations is the one of Dantzig, Fulkerson, and Johnson [1], whose associated relaxed Linear Problem (LP) is known in the literature as the Subtour Elimination Problem (SEP), or Held-Karp relaxation:

$$\text{minimize } \sum_{e \in E_n} c_e x_e \quad (1a)$$

$$\text{subject to: } \sum_{e \in \delta(v)} x_e = 2 \quad \forall v \in V, \quad (1b)$$

$$\sum_{e \in \delta(S)} x_e \geq 2 \quad \forall S \subseteq V, \ 3 \leq |S| \leq n-3, \quad (1c)$$

$$0 \leq x_e \leq 1 \quad \forall e \in E, \quad (1d)$$

where, for any subset $S \subseteq V$, $\delta(S)$ is the set of edges having exactly one node in S ¹. The term *degree constraints* refers to (1b) and the term *subtour elimination constraints* refers to (1c). The polytope defined by inequalities (1b)–(1d) is called *subtour elimination polytope* and denoted in this paper via P_{SEP}^n .

For a given cost vector $\mathbf{c}$, we denote with $\text{SEP}(\mathbf{c})$ the optimal solution of (1), and with $\text{TSP}(\mathbf{c})$ the optimal solution restricted to the integer feasible region obtained by replacing constraints (1d) with $x_e \in \{0, 1\}$. The *integrality gap* is the ratio of these two values; it can be evaluated for a single instance cost $\mathbf{c}$, as $\alpha(\mathbf{c})$, or considered in a worst case scenario, as α :

$$\alpha(\mathbf{c}) := \frac{\text{TSP}(\mathbf{c})}{\text{SEP}(\mathbf{c})}, \quad \alpha := \sup_{\mathbf{c} \text{ metric}} \frac{\text{TSP}(\mathbf{c})}{\text{SEP}(\mathbf{c})}.$$

An alternative granularity of the integrality gap can also be defined for a specific vertex $\mathbf{x} \in P_{\text{SEP}}^n$. Starting from the seminal work of [2], [3] introduced the notion of Gap^+ to bound the integrality gap attainable on $\mathbf{x}$:

$$\text{Gap}^+(\mathbf{x}) := \max \left\{ \frac{\text{TSP}(\mathbf{c})}{\mathbf{c}\mathbf{x}} \mid \mathbf{c} \text{ metric} \right\}.$$

This same quantity was later used in [4] to extend the exact computation of the integrality gap for the vertices of P_{SEP}^n with n up to 15.

¹Here and henceforth, when there is no risk of ambiguity, the singletons may be denoted without brackets, e.g., $v = \{v\}$.

The exact value of the integrality gap α is currently unknown, but it provably lives in the interval $4/3 \leq \alpha \leq 3/2$. Wolsey [5] first established the $3/2$ upper bound through an analysis of Christofides’ algorithm [6], which stood for nearly fifty years until Karlin, Klein, and Oveis Gharan obtained a (slightly) better-than- $3/2$ approximation algorithm [7], yielding $\alpha \leq 3/2 - \varepsilon$ for $\varepsilon \sim 10^{-36}$ [8]. The best known lower bound is $4/3$, attained by families of instances already known in 1995, when Goemans [9] formulated the celebrated *4/3-conjecture*, which remains open and states that $\alpha = 4/3$.

Beyond these general bounds, several results characterize the integrality gap for specific classes of instances (see, e.g., [2, 3, 10–19]), while other works focus their attention on instances with a high integrality gap (see, e.g., [2, 20–23]).

Our contribution

In this paper, we introduce a new perspective on the study of the integrality gap for the TSP. Rather than focusing on isolated vertices of the subtour elimination polytope, we associate with each of them an infinite family of related instances and analyze its limiting behavior. This leads to the *Asymptotic Framework*, a novel methodology that enables a unified treatment of both lower- and upper-bound techniques. More specifically, the framework simultaneously provides exact upper-bound characterizations and constructive lower-bound mechanisms, as described below.

Upper-bound characterization. We introduce a limit quantity that captures the exact asymptotic behavior of the integrality gap of a family of instances. This refines the approach of [3], which only provides upper estimates, by determining the precise limiting value. As a notable consequence, the framework proves the *4/3-conjecture* for a broad class of instances that we call *fractional negligible*.

Lower-bound constructions. We develop a simple procedure that transforms an arbitrary instance into a family whose asymptotic integrality gap is at least $4/3$. This shows that high-integrality-gap constructions are far more ubiquitous than previously recognized. Such instances have recently attracted considerable interest in the literature (see, e.g., [21, 23]) as they often provide challenging benchmarks for state-of-art TSP solvers and may serve as testing ground for new techniques.

Outline

Section 2 provides the necessary background material. In Section 3 we introduce the Asymptotic Framework and develop its main properties. In Section 4 we define the notion of fractional negligibility and prove the bound of $4/3$ for the corresponding class of instances. Section 5 presents generic constructions yielding high integrality gaps. Finally, Section 6 summarizes our findings and formulates an open question whose affirmative resolution would settle the *4/3-conjecture*.

2 Background material

We begin by clarifying some notations to avoid confusion. Given a graph $G = (V, E)$, the term *node* refers to the elements of V ; the term *vertex* is reserved exclusively for the extreme points of P_{SEP}^n . Moreover, for closed path, we use the terms *walk* and *tour*: a walk denotes the general case in which edges and nodes may be traversed multiple

times, whereas in a tour the nodes must be visited exactly once. Unless otherwise specified, both walks and tours are assumed to be Hamiltonian. $K_n = (V_n, E_n)$ always denotes the complete graph on n nodes.

When dealing with costs $\mathbf{c} \in \mathbb{R}^{E_n}$, we may often use the notation $\mathbf{c}(e) := c_e$ for edges $e \in E_n$, and $\mathbf{c}(F) := \sum_{e \in F} c_e$ for subsets of edges $F \subseteq E_n$.

2.1 Preliminaries on the Traveling Salesman Problem

Definition 2.1 (Metric completion) Let $G = (V, E)$ be a connected graph with $|V| = n$, and let $\mathbf{c} \in \mathbb{R}^E$ be a non-negative cost on E . The *metric completion* of $\mathbf{c}$ is the metric cost $\mathbf{c}^c$ on the complete graph K_n which assigns to each edge $ij \in E$ the cost of the shortest path between i and j in G with respect to $\mathbf{c}$.

Definition 2.2 (Support graph) Let $\mathbf{x} \in P_{\text{SEP}}^n$. The *support graph* of $\mathbf{x}$ is the weighted graph $G_{\mathbf{x}} = (V_{\mathbf{x}}, E_{\mathbf{x}})$, where $V_{\mathbf{x}} := V_n$, $E_{\mathbf{x}} := \{e \in E_n \mid x_e > 0\}$, and the weight on e is given by x_e . If the context does not lead to ambiguities, we may identify $G_{\mathbf{x}}$ with $\mathbf{x}$ itself.

Lemma 2.3 (Restated from [3]) Let $\mathbf{x} \in P_{\text{SEP}}^n$ and let $\mathbf{c} \in \arg \max \text{Gap}^+(\mathbf{x})$. Let $\mathbf{c}_{|E_{\mathbf{x}}}$ be the restriction of $\mathbf{c}$ on edges of the support graph $G_{\mathbf{x}}$, and consider its metric completion $\hat{\mathbf{c}} := (\mathbf{c}_{|E_{\mathbf{x}}})^c$. Then $\hat{\mathbf{c}} \in \arg \max \text{Gap}^+(\mathbf{x})$.

Remark 2.4 The previous lemma certifies that all the relevant information for the realization of the Gap^+ is found on the edges of the support graph. In virtue of this, we henceforth assume that every cost associated with a vertex $\mathbf{x}$ is of the form $(\mathbf{c}_{|E_{\mathbf{x}}})^c$, namely, it is the metric completion of a cost defined on the support graph of $\mathbf{x}$. Under this convention, it is sufficient to specify costs only on the edges in $E_{\mathbf{x}}$. Accordingly, tours may be replaced by walks on the support graph: this can be done by substituting in the tour any edge $ij \notin E_{\mathbf{x}}$ with the shortest path (on the support graph) connecting i and j .

Definition 2.5 (1-edge, 1-path, from [2]) Let $\mathbf{x} \in P_{\text{SEP}}^n$. An edge e of K_n is called *1-edge* of $\mathbf{x}$ if $x_e = 1$; e is called *fractional edge* if $0 < x_e < 1$. In general, the notation γ -edge is used for an edge e with $x_e = \gamma$ (e.g., $1/2$ -edge). When $\mathbf{x}$ is fractional (that is, it has fractional edges), we call *1-path* of $\mathbf{x}$ a maximal path of 1-edges in the support graph $G_{\mathbf{x}}$; the nodes of degree 2 in the 1-path are called *internal nodes*, the two remaining nodes are called *end nodes*.

Definition 2.6 (BB-move, from [2]) Let $\mathbf{x} \in P_{\text{SEP}}^n$ and let ab be one of its 1-edges. We call *BB-move* the construction of a new point $\mathbf{x}' \in P_{\text{SEP}}^{n+1}$ obtained by inserting a node into ab . It is defined as follows, where v is the new node added ($V_{n+1} = V \cup \{v\}$):

$$\begin{aligned} x'_{ab} &= 0, & x'_e &= 0 & \forall e \in \delta(v) \setminus \{av, vb\}, \\ x'_{av} &= x'_{vb} = 1, & x'_e &= x_e & \forall e \notin \delta(v) \cup \{ab\}. \end{aligned}$$

We denote this construction by $\text{BB}(\mathbf{x}, ab) := \mathbf{x}'$, using the initials of the authors of [2], where it was first presented with some associated properties. When it is clear from the context, we omit the edge ab in the notation: $\text{BB}(\mathbf{x})$.

Theorem 2.7 (Restated from [2]) Let $\mathbf{x} \in P_{\text{SEP}}^n$ and let ab be one of its 1-edges. Then $\text{BB}(\mathbf{x}, ab) \in P_{\text{SEP}}^{n+1}$, and $\text{BB}(\mathbf{x}, ab)$ is a vertex of P_{SEP}^{n+1} if and only if $\mathbf{x}$ is a vertex of P_{SEP}^n .

Lemma 2.8 (Restated from [3]) *The BB-move is Gap^+ -increasing: for every $\mathbf{x} \in P_{SEP}^n$, $\text{Gap}^+(\text{BB}(\mathbf{x})) \geq \text{Gap}^+(\mathbf{x})$.*

Remark 2.9 Let $\mathbf{x} \in P_{SEP}^n$. Any minimal walk $\mathbf{w}$ on the support graph $G_{\mathbf{x}}$ has only three different ways of traversing a 1-path, as represented in Figure 1:

- (i) using all the 1-edges once;
- (ii) using all the 1-edges twice;
- (iii) using all the 1-edges twice, except for a single 1-edge, not used at all.




Fig. 1 Ways to traverse a 1-path of $\mathbf{x}$ for a minimal walk $\mathbf{w}$ on $G_{\mathbf{x}}$.

We complete this section with a technical result which provides a useful simplification for the subsequent arguments: an optimal cost for Gap^+ can always be chosen to be uniform along every 1-path.

Lemma 2.10 *Let $\mathbf{x} \in P_{SEP}^n$. Then there exists $\bar{\mathbf{c}} \in \arg \max \text{Gap}^+(\mathbf{x})$ such that $\bar{\mathbf{c}} = (\bar{\mathbf{c}}|_{E_{\mathbf{x}}})^c$ and $\bar{\mathbf{c}}(e_1) = \bar{\mathbf{c}}(e_2)$ for any couple of 1-edges e_1, e_2 belonging to the same 1-path.*

Proof Consider a cost $\mathbf{c} \in \arg \max \text{Gap}^+(\mathbf{x})$, so that $\text{Gap}^+(\mathbf{x}) = \frac{\text{TSP}(\mathbf{c})}{\mathbf{c}\mathbf{x}}$. Remark 2.4 guarantees that we can restrict ourselves to consider costs and walks on the support graph $G_{\mathbf{x}}$. Let P be a 1-path of $\mathbf{x}$. We construct a cost $\bar{\mathbf{c}}$ on $E_{\mathbf{x}}$ as follows: $\bar{\mathbf{c}}(e) := \mathbf{c}(e)$ for the edges $e \in E_{\mathbf{x}} \setminus P$, and $\bar{\mathbf{c}}(e) := \frac{1}{|P|}\mathbf{c}(P)$ for the edges $e \in P$. By construction, $\bar{\mathbf{c}}$ is uniform along P , and it is easily checked that $\bar{\mathbf{c}}(P) = \mathbf{c}(P)$, $\bar{\mathbf{c}}\mathbf{x} = \mathbf{c}\mathbf{x}$, and $\bar{\mathbf{c}}$ is metric.

We argue that $\text{TSP}(\bar{\mathbf{c}}) \geq \text{TSP}(\mathbf{c})$. To this end, we consider an optimal walk $\bar{\mathbf{w}}$ for $\text{TSP}(\bar{\mathbf{c}})$, so that $\text{TSP}(\bar{\mathbf{c}}) = \bar{\mathbf{c}}\bar{\mathbf{w}}$, and we aim to prove $\bar{\mathbf{c}}\bar{\mathbf{w}} \geq \mathbf{c}\bar{\mathbf{w}}$, as it implies $\text{TSP}(\bar{\mathbf{c}}) = \bar{\mathbf{c}}\bar{\mathbf{w}} \geq \mathbf{c}\bar{\mathbf{w}} \geq \text{TSP}(\mathbf{c})$. We proceed by considering separately the cases of Remark 2.9.

If $\bar{\mathbf{w}}$ traverses P as in cases (i) or (ii), then we immediately have $\bar{\mathbf{c}}\bar{\mathbf{w}} = \mathbf{c}\bar{\mathbf{w}}$.

Otherwise, for the last case (iii), if $\bar{\mathbf{w}}$ traverses P avoiding the single edge e^* , then we have to correct the equivalence by adding the skipped edge to the count twice, thus completing the double transit on P : $\bar{\mathbf{c}}\bar{\mathbf{w}} + 2\bar{\mathbf{c}}(e^*) = \mathbf{c}\bar{\mathbf{w}} + 2\mathbf{c}(e^*)$. Since the choice of the 1-edge e^* to skip in $\bar{\mathbf{w}}$ is not relevant for the optimal value $\bar{\mathbf{c}}\bar{\mathbf{w}} = \text{TSP}(\bar{\mathbf{c}})$ (all the edges of P have the same $\bar{\mathbf{c}}$ -cost, by construction), we can assume that e^* is an edge of P of maximum $\mathbf{c}$ -cost. This consideration ensures that $\mathbf{c}(e^*) \geq \frac{1}{|P|}\mathbf{c}(P) = \bar{\mathbf{c}}(e^*)$, and the desired inequality follows: $\bar{\mathbf{c}}\bar{\mathbf{w}} = \mathbf{c}\bar{\mathbf{w}} + 2\mathbf{c}(e^*) - 2\bar{\mathbf{c}}(e^*) \geq \mathbf{c}\bar{\mathbf{w}}$.

The claim $\text{TSP}(\bar{c}) \geq \text{TSP}(c)$ implies the chain of inequalities

$$\text{Gap}^+(x) \geq \frac{\text{TSP}(\bar{c})}{\bar{c}x} \geq \frac{\text{TSP}(c)}{cx} = \text{Gap}^+(x) ,$$

finally proving that $\bar{c} \in \arg \max \text{Gap}^+(x)$. By sequentially repeating the same argument for all the 1-paths of x , we obtain, as a final construction, a cost on the edges of G_x that is optimal for the Gap^+ problem and is uniform along all the 1-paths. $\square$

2.2 Preliminaries on the Chinese Postman Problem

Consider a multigraph $\tilde{G} = (\tilde{V}, \tilde{E})$ and let $\tilde{c} \in \mathbb{R}^{\tilde{E}}$ be a non-negative cost vector. The *Chinese Postman Problem* (CPP) consists in finding a minimum-cost walk on $\tilde{G}$ which uses all the edges at least once. Equivalently, the CPP asks to find an Eulerian multi-subset of $\tilde{E}$ of minimum cost.

Given a subset of nodes $T \subseteq \tilde{V}$, a *T-join* is a set of edges $J \subseteq \tilde{E}$ such that the odd-degree nodes in the induced subgraph $(\tilde{V}, J)$ are exactly the nodes of T . The *Minimum T-Join Problem* (MTJ) asks for a *T-join* of minimum $\tilde{c}$ -cost.

A classical decomposition of the CPP breaks down the problem using the notion of *T-join*, as follows. Let $\tilde{V}_{\text{odd}}$ denote the subsets of nodes of $\tilde{G}$ of odd degree; then

$$\text{CPP}(\tilde{c}) = \tilde{c}(\tilde{E}) + \text{MTJ}(\tilde{V}_{\text{odd}}, \tilde{c}) ,$$

where the terms $\text{CPP}(\cdot)$ and $\text{MTJ}(\cdot, \cdot)$ denote the optimal value of the problems on the specified inputs. In other words, making the graph Eulerian requires adding a minimum-cost $\tilde{V}_{\text{odd}}$ -join, and the Chinese Postman Problem reduces to computing precisely this additional cost.

3 The Asymptotic Framework

We first outline the intuition at the basis of the Asymptotic Framework. Our starting point is the observation of two recurring phenomena in the literature.

- Families of instances with integrality gap approaching $4/3$ are often generated by iteratively applying the BB-move on a chosen vertex (see, e.g., [2, 21]).
- For many of these families, the costs yielding high integrality gaps exhibit a remarkable behavior: the values assigned to the fractional edges becomes negligible as the number of nodes increases.

Neither observation is a universal evidence: counterexamples are already known in the literature. The “donut instances” presented in [18] are produced not only by applying BB-moves, but also by replicating specific structures within the support graphs of the vertices; moreover, for these instances, the costs used to achieve a high integrality gap do not vanish on the fractional edges of the vertices. Nevertheless, the two phenomena described above appear frequently enough to deserve a systematic investigation. The Asymptotic Framework addresses the first observation in full generality by analyzing the limiting behavior of entire families of instances through suitable limit objects. The second observation is captured by the notion of *fractional negligibility*, which is introduced and investigated in detail in Section 4.

Before delving into the details of the formalization, we provide Tables 1 and 2 as guidance throughout the framework: they collect the different notions of integrality gap introduced in the paper, together with the main results relating them.

Structure	Description	Integrality gap notion
$\mathbf{x}$	vertex of P_{SEP}^n	$\text{Gap}^+(\mathbf{x}) := \max \left\{ \frac{\text{TSP}(\mathbf{c})}{c\mathbf{x}} \mid \mathbf{c} \text{ metric} \right\}$ $\text{Gap}_0^+(\mathbf{x}) := \max \left\{ \frac{\text{TSP}(\mathbf{c})}{c\mathbf{x}} \mid \mathbf{c} \text{ metric, } c(e) = 0 \ \forall e \text{ fractional edge of } \mathbf{x} \right\}$
$\hat{\mathbf{x}}$	weighted graph obtained from $\mathbf{x}$ replacing 1-paths with single 1-edges	$\widehat{\text{Gap}}(\hat{\mathbf{x}}) := \max \left\{ \frac{\text{TSP}_{\mathcal{P}}(\hat{\mathbf{c}})}{\hat{\mathbf{c}}\hat{\mathbf{x}}} \mid \hat{\mathbf{c}} \text{ non-negative} \right\}$ $\widehat{\text{Gap}}_0(\hat{\mathbf{x}}) := \max \left\{ \frac{\text{TSP}_{\mathcal{P}}(\hat{\mathbf{c}})}{\hat{\mathbf{c}}\hat{\mathbf{x}}} \mid \hat{\mathbf{c}} \text{ non-negative, } \hat{\mathbf{c}}(e) = 0 \ \forall e \text{ fractional edge of } \hat{\mathbf{x}} \right\}$
$\tilde{\mathbf{x}}$	multigraph obtained from $\hat{\mathbf{x}}$ contracting the fractional edges	$\widetilde{\text{Gap}}(\tilde{\mathbf{x}}) := \max \left\{ \frac{\text{CPP}(\tilde{\mathbf{c}})}{\tilde{\mathbf{c}}\tilde{\mathbf{x}}} \mid \tilde{\mathbf{c}} \text{ non-negative} \right\}$

Table 1 Structures introduced in the paper together with the corresponding notions of integrality gap. The term $\text{TSP}_{\mathcal{P}}$ denotes the TSP with the additional requirement that feasible walks must use every 1-edge at least once.

Reference	Result	Reference	Result
Lemma 2.8	$\text{Gap}^+(\text{BB}(\mathbf{x})) \geq \text{Gap}^+(\mathbf{x})$	Lemma 4.6	$\widehat{\text{Gap}}(\tilde{\mathbf{x}}) = \widehat{\text{Gap}}_0(\tilde{\mathbf{x}})$ $\widehat{\text{Gap}}(\tilde{\mathbf{x}}) \leq \widehat{\text{Gap}}(\hat{\mathbf{x}})$
Theorem 3.3	$\widehat{\text{Gap}}(\hat{\mathbf{x}}) = \lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k)$ $\text{Gap}^+(\mathbf{x}) \leq \widehat{\text{Gap}}(\hat{\mathbf{x}})$	Theorem 4.7	$\mathbf{x} \text{ FN} \Rightarrow \widehat{\text{Gap}}(\hat{\mathbf{x}}) = \lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k)$ $\text{Gap}^+(\mathbf{x}) \leq \widehat{\text{Gap}}(\tilde{\mathbf{x}})$
Lemma 4.3	$\mathbf{x} \text{ FN} \Leftrightarrow \widehat{\text{Gap}}_0(\hat{\mathbf{x}}) = \widehat{\text{Gap}}(\hat{\mathbf{x}})$	Theorem 4.9	$\mathbf{x} \text{ FN} \Rightarrow \text{Gap}^+(\mathbf{x}) \leq 4/3$

Table 2 Overview of the main results of the paper. FN stands for “fractional negligible”.

3.1 Formalization

In this section, we formalize the main elements and definitions of our framework, constructing a specific sequence of vertices that will be the object of our study.

Definition 3.1 ($\mathcal{F}(\mathbf{x})$, $\hat{\mathbf{x}}$) Let $\mathbf{x} \in P_{\text{SEP}}^n$ be a vertex. We define $\mathcal{F}(\mathbf{x}) = \{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots\}$ to be the sequence of vertices such that, for every $k \in \mathbb{N}$, $\mathbf{x}_k$ is obtained from $\mathbf{x}$ by replacing all its 1-paths with 1-paths having k internal nodes. We also define $\hat{\mathbf{x}} := \mathbf{x}_0$. See Figure 2 for an illustration of these constructions.

In the definition, the notation $\hat{\mathbf{x}}$ is introduced to emphasize a conceptual difference: even though $\hat{\mathbf{x}}$ and $\mathbf{x}_0$ coincide as weighted graphs, $\hat{\mathbf{x}}$ should be interpreted as the limit object of the sequence $\mathcal{F}(\mathbf{x})$. Intuitively, as k increases, every 1-path of $\mathbf{x}$ is populated with an ever-growing number of internal nodes. In the limit, the 1-paths are collapsed back into single 1-edges, with the understanding that they carry “infinitely many hidden nodes”. Therefore, a walk can be said to visit all of them only by traversing the corresponding edge entirely: every walk on $\hat{\mathbf{x}}$ is thus required to use each 1-edge at




Fig. 2 Example of a vertex x with the associated sequence $\mathcal{F}(x)$ and weighted graph $\hat{x}$.

least once. Following this idea, we introduce an extension of the notion of *integrality gap*:

$$\widehat{\text{Gap}}(\hat{x}) := \max \left\{ \frac{\text{TSP}_{\mathcal{P}}(\hat{c})}{\hat{c}\hat{x}} \mid \hat{c} \text{ non-negative} \right\},$$

where $\mathcal{P}$ is the set of 1-edges of $\hat{x}$, and $\text{TSP}_{\mathcal{P}}$ denotes, here and henceforth, the TSP with the additional requirement that every feasible walk must use each edge in $\mathcal{P}$ at least once.

Remark 3.2 When comparing $\hat{x}$ with x_k , there are natural correspondences between the objects defined on the two structures, both at the level of walks and of costs, as represented in Figure 3. A walk on $\hat{x}$ corresponds to a walk on x_k which uses a 1-path in place of the corresponding 1-edge. Conversely, following Remark 2.9, a minimal walk on x_k originates a walk on $\hat{x}$ by replacing each traversal of a 1-path of x_k with the corresponding 1-edge of $\hat{x}$, used once in case (i), and twice in cases (ii) or (iii). A cost $\hat{c}$ on $\hat{x}$ determines a cost c_k on x_k by uniformly distributing the cost of each 1-edge of $\hat{x}$ along the corresponding 1-path of x_k , and leaving all other edge costs unchanged, as suggested by Lemma 2.10. Conversely, a cost c_k on x_k determines a cost $\hat{c}$ on $\hat{x}$ which coincides with c_k on the fractional edges and assigns $\hat{c}(e) := c_k(P)$ for each 1-edge e of $\hat{x}$, where P is the corresponding 1-path in x_k .

We remark that the costs $\hat{c}$ appearing in the definition of $\widehat{\text{Gap}}$ do not need to satisfy metric requirements. This choice is justified by the following two observations.

For the fractional edges, metricity is enforced by optimality. Indeed, if a maximizing cost $\hat{c}$ assigns to a fractional edge $e \notin \mathcal{P}$ a value larger than the cost of a shortest path P between its endpoints, then replacing $\hat{c}(e)$ with $\hat{c}(P)$ leaves $\text{TSP}_{\mathcal{P}}$ unchanged while strictly decreasing the denominator $\hat{c}\hat{x}$, contradicting optimality.

For the 1-edges, the perspective is different. Since a 1-edge of $\hat{x}$ represents the limit of a path containing infinitely many internal nodes, its cost need not satisfy a metric condition in the collapsed graph. At the same time, every cost $\hat{c}$ on $\hat{x}$ naturally induces, as explained in Remark 3.2, a cost c_k on x_k that is uniform along the 1-paths,




Fig. 3 Natural correspondences of walks and costs between $\hat{x}$ and x_k .

consequently, for every $k \geq 1$, every 1-edge is the shortest path between its endpoints. Thus, costs $\hat{c}$ which fail to be metric on the 1-edges of $\hat{x}$ can always be interpreted as the limits of metric costs c_k on x_k .

We now state the central theorem of the framework, which clarifies the role of the $\widehat{\text{Gap}}$ quantity.

Theorem 3.3 *Let $x \in P_{SEP}^n$ and consider $\mathcal{F}(x) = \{x_0, x_1, x_2, \dots\}$ and $\hat{x}$ as in Definition 3.1. Then the corresponding sequence of Gap^+ values converges exactly to $\widehat{\text{Gap}}(\hat{x})$:*

$$\lim_{k \rightarrow \infty} \text{Gap}^+(x_k) = \widehat{\text{Gap}}(\hat{x}).$$

In particular, $\text{Gap}^+(x) \leq \widehat{\text{Gap}}(\hat{x})$.

Proof The sequence $\{\text{Gap}^+(x_k)\}_{k \in \mathbb{N}}$ is monotone non-decreasing, by Lemma 2.8, hence it converges to a limit, which is finite since all the Gap^+ values are upper bounded by $3/2$.

We first prove: $\lim_{k \rightarrow \infty} \text{Gap}^+(x_k) \leq \widehat{\text{Gap}}(\hat{x})$. Consider an arbitrary k and let c_k be an optimal cost for $\text{Gap}^+(x_k)$, so that $\text{Gap}^+(x_k) = \frac{\text{TSP}(c_k)}{c_k x_k}$. Let $\hat{c}$ be the cost on $\hat{x}$ naturally induced by c_k as in Remark 3.2. Clearly $\hat{c}\hat{x} = c_k x_k$. Moreover, $\text{TSP}(c_k) \leq \text{TSP}_{\mathcal{P}}(\hat{c})$, where $\mathcal{P}$ is the set of 1-edges of $\hat{x}$: the inequality holds because we can naturally convert any feasible walk for $\text{TSP}_{\mathcal{P}}(\hat{c})$ into a feasible walk for $\text{TSP}(c_k)$ while preserving the objective value, simply by replacing each use of a 1-edge of $\hat{x}$ with the corresponding 1-path of x_k . Therefore we obtain the chain of inequalities

$$\text{Gap}^+(x_k) = \frac{\text{TSP}(c_k)}{c_k x_k} \leq \frac{\text{TSP}_{\mathcal{P}}(\hat{c})}{\hat{c}\hat{x}} \leq \widehat{\text{Gap}}(\hat{x})$$

which propagates to the limit: $\lim_{k \rightarrow \infty} \text{Gap}^+(x_k) \leq \widehat{\text{Gap}}(\hat{x})$.

It remains to prove: $\lim_{k \rightarrow \infty} \text{Gap}^+(x_k) \geq \widehat{\text{Gap}}(\hat{x})$.

Let $\hat{c}$ be an optimal cost for $\widehat{\text{Gap}}(\hat{x})$, so that $\widehat{\text{Gap}}(\hat{x}) = \frac{\text{TSP}_{\mathcal{P}}(\hat{c})}{\hat{c}\hat{x}}$, where $\mathcal{P}$ is the set of 1-edges of $\hat{x}$. Up to rescaling, we may assume $\hat{c}\hat{x} = 1$. For every $k \in \mathbb{N}$, consider the cost c_k on x_k naturally obtained by distributing the cost uniformly along the 1-paths P of x_k (see Remark 3.2), so that each edge of P receives $1/(k+1) \hat{c}(e)$, where e is the corresponding 1-edge of $\hat{x}$. By construction, $c_k x_k = \hat{c}\hat{x} = 1$; in particular, this implies a bound of $1/k$ on

the c_k -cost of each single 1-edge².

Now consider an optimal walk w_k for $\text{TSP}(c_k)$, so that $\text{TSP}(c_k) = c_k w_k$: the associated walk $\hat{w}$ on $\hat{x}$ (see Remark 3.2) is a feasible solution for $\text{TSP}_{\mathcal{P}}(\hat{c})$. This construction may bring additional cost to the walk only in the case that w_k traverses some 1-paths of x_k avoiding individual 1-edges; however, the extra cost does not exceed $2p^{1/k}$, where p is the number of 1-paths in x . Therefore, $\text{TSP}_{\mathcal{P}}(\hat{c}) \leq \hat{c}\hat{w} \leq c_k w_k + 2p^{1/k} = \text{TSP}(c_k) + 2p^{1/k}$.

Using the normalization $c_k x_k = \hat{c}\hat{x} = 1$, we have $\widehat{\text{Gap}}(\hat{x}) = \text{TSP}_{\mathcal{P}}(\hat{c})$ and $\text{Gap}^+(x_k) \geq \text{TSP}(c_k)$. Concatenating with the previous inequality yields $\widehat{\text{Gap}}(\hat{x}) \leq \text{Gap}^+(x_k) + 2p^{1/k}$, which ultimately proves the claim in the limit, $\lim_{k \rightarrow \infty} \text{Gap}^+(x_k) \geq \widehat{\text{Gap}}(\hat{x})$, since the term $2p^{1/k}$ vanishes as $k \rightarrow \infty$.

Finally, for all $k \in \mathbb{N}$ it holds $\text{Gap}^+(x_k) \leq \widehat{\text{Gap}}(\hat{x})$, since the sequence $\{\text{Gap}^+(x_k)\}_{k \in \mathbb{N}}$ is non-decreasing; for large enough k (namely, larger than the maximum number of internal nodes in the 1-paths of x) we have $\text{Gap}^+(x) \leq \text{Gap}^+(x_k) \leq \widehat{\text{Gap}}(\hat{x})$. $\square$

Operationally, the quantity $\widehat{\text{Gap}}$ can be computed as the solution of a linear program, in a way similar to the approaches used in [2, 3]. Since in this paper we do not explicitly calculate $\widehat{\text{Gap}}$ for any instance, we omit the details of this formulation.

The computation of $\widehat{\text{Gap}}$ pursues the same objective as the GBe algorithm introduced in [3]: they both aim to estimate the largest Gap^+ that can be obtained from a given vertex x through applications of the BB-move. From this perspective, $\widehat{\text{Gap}}$ may be regarded as a refined version of the GBe. While the GBe progressively improves its result through a sequence of LP optimizations and only guarantees an upper estimate, computing $\widehat{\text{Gap}}(\hat{x})$ directly returns the exact bound in a single optimization process³. For these reasons, the Asymptotic Framework can be regarded both as a theoretical strengthening of the approach of [3] and as a computational tool providing immediate access to the tight bound associated with the family $\mathcal{F}(x)$.

4 Fractional Negligibility

In this section, we develop the intuition presented at the outset of the framework: for several important families of instances studied in the literature, costs achieving a high integrality gap are negligible on the fractional edges. We introduce a formal definition that captures this phenomenon and use the Asymptotic Framework to investigate its properties.

Definition 4.1 (Fractional negligibility) Let $x \in P_{\text{SEP}}^n$ and consider the $\mathcal{F}(x) = \{x_0, x_1, x_2, \dots\}$. We say that x is *fractional negligible* if there exists a sequence of costs $\{c_0, c_1, c_2, \dots\}$ and a sequence of small values $\{\varepsilon_0, \varepsilon_1, \varepsilon_2, \dots\}$, with $\varepsilon_k \rightarrow 0$ as $k \rightarrow \infty$, satisfying the following properties for every $k \in \mathbb{N}$:

normalization:	$c_k x_k = 1$;	
uniformity:	$c_k(e_1) = c_k(e_2)$	$\forall e_1, e_2$ 1-edges of the same 1-path;
near-optimality:	$\text{TSP}(c_k) \geq \text{Gap}^+(x_k) - \varepsilon_k$;	
negligibility:	$c_k(e) = 0$	$\forall e$ fractional edge of x .

²A sharper bound is $1/(k+1)$: we use $1/k$ for simplicity, since all we require is that it vanishes as $k \rightarrow \infty$.

³Notice that $\widehat{\text{Gap}}(\hat{x})$ bounds the Gap^+ value for all the successors of x , as defined in [3]: any successor x' of x , indeed, originates the same family $\mathcal{F}(x') = \mathcal{F}(x)$, hence their integrality gaps are bounded by the same quantity $\widehat{\text{Gap}}(\hat{x}') = \widehat{\text{Gap}}(\hat{x})$.

Accordingly, for a TSP instance $\mathbf{c}$, we say that $\mathbf{c}$ is *fractional negligible* if there exists an optimal solution for $\text{SEP}(\mathbf{c})$ that is a fractional negligible vertex.

Remark 4.2 In the definition of fractional negligibility, the first two requirements, namely normalization and uniformity, are merely technical conveniences: they are not restrictive, even with respect to optimality (see Lemma 2.10), thus they entail no loss of generality. The first essential condition is near optimality: it implies

$$\lim_{k \rightarrow \infty} \text{TSP}(\mathbf{c}_k) = \widehat{\text{Gap}}(\mathbf{x}) ,$$

where $\{\mathbf{c}_k\}_{k \in \mathbb{N}}$ is the sequence of costs appearing in Definition 4.1. Indeed, near optimality yields $\lim_{k \rightarrow \infty} \text{TSP}(\mathbf{c}_k) \geq \lim_{k \rightarrow \infty} (\text{Gap}^+(\mathbf{x}_k) - \varepsilon_k) = \lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k)$; on the other hand, under the normalization $\mathbf{c}_k \mathbf{x}_k = 1$, the inequality $\lim_{k \rightarrow \infty} \text{TSP}(\mathbf{c}_k) \leq \lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k)$ holds by definition of $\text{Gap}^+(\mathbf{x}_k)$. Therefore, $\lim_{k \rightarrow \infty} \text{TSP}(\mathbf{c}_k) = \lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k)$, and this quantity coincides with $\widehat{\text{Gap}}(\mathbf{x})$ by Theorem 3.3. Finally, the negligibility requirement is the key assumption which characterizes the properties illustrated in the remainder of the section.

Fractional negligibility concerns the interplay between each element of $\mathcal{F}(\mathbf{x})$ and a corresponding near-optimal cost; establishing whether a given vertex satisfies this property may be a non-trivial task. In the following lemma, we provide an equivalent characterization that can be computed directly in the limit constructions, thereby simplifying its verification. It is based on a restrictive variant of $\widehat{\text{Gap}}$, denoted $\widehat{\text{Gap}}_0$, in which all admissible costs must be 0 on the fractional edges:

$$\widehat{\text{Gap}}_0(\hat{\mathbf{x}}) := \max \left\{ \frac{\text{TSP}_{\mathcal{P}}(\hat{\mathbf{c}})}{\hat{\mathbf{c}} \hat{\mathbf{x}}} \mid \hat{\mathbf{c}} \text{ non-negative, } \hat{\mathbf{c}}(e) = 0 \ \forall e \text{ fractional edge of } \hat{\mathbf{x}} \right\} .$$

Lemma 4.3 *Let $\mathbf{x} \in P_{\text{SEP}}^n$, and let $\hat{\mathbf{x}}$ as in Definition 3.1. The following conditions are equivalent:*

- (1) $\mathbf{x}$ is fractional negligible;
- (2) $\widehat{\text{Gap}}_0(\hat{\mathbf{x}}) = \widehat{\text{Gap}}(\hat{\mathbf{x}})$.

Proof We first prove that condition (1) implies (2).

Consider a fractional negligible vertex $\mathbf{x}$ with the corresponding elements of Definition 4.1: $\mathbf{x}_k$, $\mathbf{c}_k$, ε_k , for $k \in \mathbb{N}$. Define Gap_0^+ as the analogous of Gap^+ under the additional constraint that the considered costs must evaluate to 0 on all fractional edges. Combining this definition with the fractional negligibility of $\mathbf{x}$ yields $\text{TSP}(\mathbf{c}_k) \leq \text{Gap}_0^+(\mathbf{x}_k)$. Furthermore, following the proof of Theorem 3.3 (with Gap_0^+ in place of Gap^+ and $\widehat{\text{Gap}}_0$ in place of $\widehat{\text{Gap}}$), we see that $\text{Gap}_0^+(\mathbf{x}_k) \leq \widehat{\text{Gap}}_0(\hat{\mathbf{x}})$ for every $k \in \mathbb{N}$, thus $\text{TSP}(\mathbf{c}_k) \leq \widehat{\text{Gap}}_0(\hat{\mathbf{x}})$. On the other hand, $\widehat{\text{Gap}}_0(\hat{\mathbf{x}}) \leq \widehat{\text{Gap}}(\hat{\mathbf{x}})$ by definition. Therefore, $\widehat{\text{Gap}}_0(\hat{\mathbf{x}})$ is trapped between $\text{TSP}(\mathbf{c}_k)$ and $\widehat{\text{Gap}}(\hat{\mathbf{x}})$. Since Remark 4.2 shows that $\lim_{k \rightarrow \infty} \text{TSP}(\mathbf{c}_k) = \widehat{\text{Gap}}(\hat{\mathbf{x}})$, the two quantities $\widehat{\text{Gap}}_0(\hat{\mathbf{x}})$ and $\widehat{\text{Gap}}(\hat{\mathbf{x}})$ must coincide.

We now prove the opposite implication, that is, (2) implies (1).

Consider an optimal cost $\hat{\mathbf{c}}$ for $\widehat{\text{Gap}}_0(\hat{\mathbf{x}})$: $\hat{\mathbf{c}}$ is optimal for $\widehat{\text{Gap}}(\hat{\mathbf{x}})$ as well, by hypothesis (2), and it is 0 on the fractional edges. Additionally, we assume it is normalized: $\hat{\mathbf{c}} \hat{\mathbf{x}} = 1$. We

construct, for every $k \in \mathbb{N}$, the corresponding cost c_k on x_k (see Remark 3.2). Following the proof of Theorem 3.3, we see that $\text{TSP}(c_k) \geq \widehat{\text{Gap}}(\hat{x}) - 2p^{1/k} \geq \text{Gap}^+(x_k) - 2p^{1/k}$, where p is the number of 1-paths in x . Thence near-optimality of Definition 4.1 holds, with $\varepsilon_k := 2p^{1/k}$. Finally, normalization, uniformity, and negligibility hold as well by construction; thus the vertex x is fractional negligible. $\square$

The result shows how $\widehat{\text{Gap}}_0$ is in fact as expressive as the original $\widehat{\text{Gap}}$ for the case of fractional negligible vertices. Motivated by this observation, we introduce an additional graph structure that makes the key features of fractional negligible vertices more evident and leads to a more efficient description of the problem.

Definition 4.4 ($\tilde{x}$) Let $x \in P_{\text{SEP}}^n$ and let $\hat{x}$ be the associated weighted graph. We define $\tilde{x}$ to be the multigraph obtained from $\hat{x}$ by contracting all fractional edges. Equivalently, each connected component of the subgraph induced by the fractional edges is contracted into a single node.




Fig. 4 A vertex x and its limit constructions $\hat{x}$ and $\tilde{x}$.

Figure 4 shows an example of this construction. Note that $\tilde{x}$ may have parallel edges (if multiple 1-paths of x end up in the same fractional connected components) as well as loops (if a 1-path of x has both end nodes in one fractional connected component). Moreover, all the edges of $\tilde{x}$ have the same weight, since they originate from 1-edges of $\hat{x}$, thus we can consider $\tilde{x}$ to be unweighted.

Remark 4.5 $\tilde{x}$ is 2-edge-connected. Indeed, any bridge in $\tilde{x}$ would correspond to a 1-path of x which in turn forms a bridge; since x is 2-edge-connected (by subtour elimination constraints), it can not have bridges, and so must be for $\tilde{x}$.

For the multigraph $\tilde{x}$, the corresponding notion of integrality gap is defined as follows,

$$\widetilde{\text{Gap}}(\tilde{x}) := \max \left\{ \frac{\text{CPP}(\tilde{c})}{\tilde{c}\tilde{x}} \mid \tilde{c} \text{ non-negative} \right\},$$

where $\tilde{c}$ may be any non-negative cost on the edges of $\tilde{x}$, without metric requirements. The CPP arises as $\tilde{x}$ is originated from $\hat{x}$ by contracting the fractional edges: only the 1-edges survive this operation, and the TSP with the additional requirement of using all of them at least once coincides with the CPP.

Lemma 4.6 *Let $\mathbf{x} \in P_{SEP}^n$ and let $\hat{\mathbf{x}}$ and $\tilde{\mathbf{x}}$ be the corresponding weighted graph and multigraph. Then $\widehat{Gap}(\tilde{\mathbf{x}}) = \widehat{Gap}_0(\hat{\mathbf{x}})$. In particular, $\widehat{Gap}(\tilde{\mathbf{x}}) \leq \widehat{Gap}(\hat{\mathbf{x}})$.*

Proof We establish natural correspondences between the objects involved in the definitions of $\widehat{Gap}_0(\hat{\mathbf{x}})$ and $\widehat{Gap}(\tilde{\mathbf{x}})$.

Since $\tilde{\mathbf{x}}$ is obtained from $\hat{\mathbf{x}}$ by contracting the components induced by fractional edges, a cost on $\hat{\mathbf{x}}$ vanishing on all fractional edges uniquely identifies a cost on $\tilde{\mathbf{x}}$, and vice-versa.

Likewise, feasible walks on the two graphs correspond naturally. Given a walk on $\hat{\mathbf{x}}$, the traversals of fractional edges may be ignored, yielding a walk on $\tilde{\mathbf{x}}$. Conversely, each walk on $\tilde{\mathbf{x}}$ can be transformed into a walk on $\hat{\mathbf{x}}$ by connecting consecutive 1-edges through suitable paths inside the fractional components. Moreover, the traversal requirements imposed by $\text{TSP}_{\mathcal{P}}$ in the definition of $\widehat{Gap}_0$ coincide exactly with those imposed by CPP in the definition of $\widehat{Gap}$. These correspondences preserve the costs of every walk, and corresponding cost assignments on $\hat{\mathbf{x}}$ and $\tilde{\mathbf{x}}$ have the same objective value. Consequently, the ratios appearing in the definitions of $\widehat{Gap}_0(\hat{\mathbf{x}})$ and $\widehat{Gap}(\tilde{\mathbf{x}})$ coincide, thence the equality $\widehat{Gap}_0(\hat{\mathbf{x}}) = \widehat{Gap}(\tilde{\mathbf{x}})$. The definition of $\widehat{Gap}_0$ trivially implies $\widehat{Gap}_0(\hat{\mathbf{x}}) \leq \widehat{Gap}(\hat{\mathbf{x}})$, and the proof is complete. $\square$

We are now in a position to combine the previous properties into a single result for fractional negligible vertices.

Theorem 4.7 *Let $\mathbf{x} \in P_{SEP}^n$ and consider $\mathcal{F}(\mathbf{x}) = \{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots\}$ and $\tilde{\mathbf{x}}$ of Definitions 3.1 and 4.4. If $\mathbf{x}$ is fractional negligible, then the corresponding sequence of Gap^+ values converges exactly to $\widehat{Gap}(\tilde{\mathbf{x}})$:*

$$\lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k) = \widehat{Gap}(\tilde{\mathbf{x}}) .$$

In particular, $\text{Gap}^+(\mathbf{x}) \leq \widehat{Gap}(\tilde{\mathbf{x}})$.

Proof Theorem 3.3 shows that $\lim_{k \rightarrow \infty} \text{Gap}^+(\mathbf{x}_k) = \widehat{Gap}(\hat{\mathbf{x}})$, and $\text{Gap}^+(\mathbf{x}) \leq \widehat{Gap}(\hat{\mathbf{x}})$. For fractional negligible vertices, the desired claim is proved by the equivalence $\widehat{Gap}(\hat{\mathbf{x}}) = \widehat{Gap}_0(\hat{\mathbf{x}}) = \widehat{Gap}(\tilde{\mathbf{x}})$, given by Lemmas 4.3 and 4.6. $\square$

This theorem provides a practical tool for constructing vertices with high integrality gaps. Given a vertex $\mathbf{x}$, the associated multigraph $\tilde{\mathbf{x}}$ can be obtained directly, and, for any fixed cost $\tilde{\mathbf{c}}$ on $\tilde{\mathbf{x}}$, both CPP($\tilde{\mathbf{c}}$) and the ratio $\frac{\text{CPP}(\tilde{\mathbf{c}})}{\tilde{\mathbf{c}} \cdot \tilde{\mathbf{x}}}$ can be computed efficiently. Since $\widehat{Gap}$ is defined as the maximum of these ratios, every choice of $\tilde{\mathbf{c}}$ immediately yields a lower bound for $\widehat{Gap}(\tilde{\mathbf{x}})$, and ultimately for $\widehat{Gap}(\hat{\mathbf{x}})$, by Lemma 4.6. In particular, we obtain a lower bound on the Gap^+ achievable by the family $\mathcal{F}(\mathbf{x})$. Explicit constructions illustrating the effectiveness of this application are given in Section 5.

Lemma 4.8 *Let $\mathbf{x} \in P_{SEP}^n$, and let $\tilde{\mathbf{x}}$ be the associated multigraph. Then $\widehat{Gap}(\tilde{\mathbf{x}}) \leq 4/3$.*

Proof Given the multigraph $\tilde{\mathbf{x}} = (\tilde{V}, \tilde{E})$, we consider the polytope of $\tilde{V}_{odd}$ -joins (where $\tilde{V}_{odd}$ is the set of odd-degree nodes of $\tilde{\mathbf{x}}$), i.e., the convex hull of the incidence vectors of all $\tilde{V}_{odd}$ -joins. The well-known result of Edmonds and Johnson [24] states that this integral polytope is described exactly by the following linear system:

$$\begin{cases} \sum_{e \in \delta(S)} t_e \geq 1 & \forall S \subseteq \tilde{V} \text{ with } |S \cap \tilde{V}_{odd}| \text{ odd} , \\ 0 \leq t_e \leq 1 & \forall e \in \tilde{E} . \end{cases}$$

Let $\mathbf{t}^* \in \mathbb{R}^{\tilde{E}}$ defined as $t_e^* = 1/3$ for all the edges $e \in \tilde{E}$. We claim that $\mathbf{t}^*$ is feasible in the $\tilde{V}_{odd}$ -join polytope. Indeed, for any subset $S \subseteq \tilde{V}$, the handshaking lemma gives

$$2|\tilde{E}(S)| + |\delta(S)| = \sum_{v \in S} \deg(v) = \sum_{v \in (S \cap \tilde{V}_{odd})} \deg(v) + \sum_{v \in (S \setminus \tilde{V}_{odd})} \deg(v) ,$$

which, considered as an equation modulo 2, yields

$$|\delta(S)| \equiv \sum_{v \in (S \cap \tilde{V}_{odd})} \deg(v) \equiv |S \cap \tilde{V}_{odd}| \pmod{2} .$$

Therefore, $|\delta(S)|$ is odd for any $S \subseteq \tilde{V}$ with $|S \cap \tilde{V}_{odd}|$ odd, and the 2-edge-connectivity of $\tilde{\mathbf{x}}$ (see Remark 4.5) forces $|\delta(S)| \geq 3$. The corresponding constraint is thus satisfied:

$$\sum_{e \in \delta(S)} t_e^* = |\delta(S)| \cdot 1/3 \geq 3 \cdot 1/3 = 1 .$$

Consider any non-negative cost vector $\tilde{\mathbf{c}} \in \mathbb{R}^{\tilde{E}}$. Then $\text{MTJ}(\tilde{V}_{odd}, \tilde{\mathbf{c}}) \leq \tilde{\mathbf{c}}\mathbf{t}^* = 1/3 \tilde{\mathbf{c}}(\tilde{E})$, since a minimum-cost solution can always be found among the vertices of the polytope, and these are exactly the $\tilde{V}_{odd}$ -joins. We apply the classical CPP decomposition to get

$$\frac{\text{CPP}(\tilde{\mathbf{c}})}{\tilde{\mathbf{c}}\tilde{\mathbf{x}}} = \frac{\tilde{\mathbf{c}}(\tilde{E}) + \text{MTJ}(\tilde{V}_{odd}, \tilde{\mathbf{c}})}{\tilde{\mathbf{c}}(\tilde{E})} \leq \frac{\tilde{\mathbf{c}}(\tilde{E}) + 1/3 \tilde{\mathbf{c}}(\tilde{E})}{\tilde{\mathbf{c}}(\tilde{E})} = \frac{4}{3} .$$

As this bound is independent of the choice of $\tilde{\mathbf{c}}$, it also holds after taking the maximum: $\text{Gap}^+(\tilde{\mathbf{x}}) \leq 4/3$. $\square$

The previous lemma holds for generic points of P_{SEP}^n ; when combined with Theorem 4.7, it finally proves the $4/3$ -conjecture for the entire class of fractional negligible vertices.

Theorem 4.9 *The integrality gap of a fractional negligible vertex does not exceed $4/3$.*

The results of this section show that the framework developed here serves a dual purpose. On the one hand, Theorem 3.3 and Lemma 4.6 apply in full generality, providing a systematic way to construct families of instances exhibiting high integrality gaps (more details in Section 5). On the other hand, Theorem 4.9 illustrates that the same framework can also be used to derive upper bounds: being fractional negligible is a sufficient condition to guarantee the $4/3$ bound.

The natural question to investigate at this stage is how broad the class of fractional negligible vertices actually is. As anticipated at the beginning of Section 3, this property is not universal. For example, whenever a vertex $\mathbf{x}$ has non-trivial integrality gap but the associated multigraph $\tilde{\mathbf{x}}$ is Eulerian, then $\text{Gap}^+(\tilde{\mathbf{x}}) = 1 < \text{Gap}^+(\mathbf{x})$,

meaning that $\tilde{x}$ can not be fractional negligible, as it would contradict Theorem 4.7. A known family of such vertices is the class of “donut instances” [18]. Nevertheless, the class appears to be rich and relevant. Explicit computations show that both the “envelope instances” [2] and the “tetrahedron instances” [21] are fractional negligible. More generally, all the “ $1/2$ -triangle instances” defined in [15] fall in the same class, as follows by combining the constructions of Section 5 with the known $4/3$ upper bound for their integrality gap. It is noteworthy that, among the well-studied instances in the literature, those achieving the highest integrality gap for a given number of nodes are indeed fractional negligible. This observation motivates the following open question.

Question 1 For every n , is fractional negligibility a necessary condition for the vertices of P_{SEP}^n attaining the maximum integrality gap?

This direction of investigation might be of particular interest: a positive answer to the question would immediately close the $4/3$ -conjecture.

5 Families with high integrality gap

This section illustrates how the framework developed so far can be used to construct infinitely many *new* families of instances whose integrality gap approaches $4/3$. One key observation is that every 2-edge-connected cubic multigraph naturally gives rise to such a family: we make this connection explicit and prove that the resulting constructions attain the $4/3$ limit. Even more generally, we show that families with high integrality gap are somehow pervasive: *any* vertex x can be modified locally to obtain a similar x' whose family $\mathcal{F}(x')$ has integrality gaps approaching $4/3$.

5.1 Canonical constructions

Consider a cubic, 2-edge-connected multigraph $\tilde{G} = (\tilde{V}, \tilde{E})$. We construct a weighted graph $G = (V, E)$ as follows:

1. for every node of $\tilde{V}$, create a triangle of three new nodes joined by $1/2$ -edges;
2. for every edge of $\tilde{E}$, add a single 1-edge between the corresponding triangles.

Since $\tilde{G}$ is 3-regular, each of the three edges incident to a node of $\tilde{G}$ can always be attached to a different node in the corresponding $1/2$ -triangle. We make such a choice throughout the construction, thereby guaranteeing that G is a graph with no parallel edges or loops. Let $n := 3|\tilde{V}|$, and denote by $x \in \mathbb{R}^{E_n}$ the vector corresponding to G .

Claim 5.1 x as described above is a vertex of P_{SEP}^n .

Proof We begin by proving the feasibility of $x \in P_{\text{SEP}}^n$.

First, we observe that all degree constraints are satisfied. Each node v of V is incident to two edges of a $1/2$ -triangle and one 1-edge corresponding to an edge of $\tilde{G}$, thus $\delta(v) = 2 \cdot 1/2 + 1 = 2$. Next, we prove that x satisfies all subtour-elimination constraints. Notice that, whenever a

cut $\delta(S)$ contains a fractional edge, then it must contain exactly two $1/2$ -edges of the same triangle. Thus every cut $\delta(S)$ is composed of 1-edges and pairs of $1/2$ -edges. Now, $\delta(S)$ cannot consist of a single 1-edge: it would correspond to a bridge of $\tilde{G}$. Analogously, $\delta(S)$ cannot consist of a single pair of $1/2$ -edges: the adjacent 1-edge would correspond to a bridge of $\tilde{G}$. This implies that $\delta(S)$ must contain at least two elements, among 1-edges and pairs of $1/2$ -edges, each contributing value 1 to $\mathbf{x}(\delta(S))$. Therefore, the corresponding subtour-elimination constraint is satisfied: $\mathbf{x}(\delta(S)) \geq 2 \cdot 1 = 2$.

It only remains to prove that $\mathbf{x}$ is an extreme point of P_{SEP}^n . Suppose that $\mathbf{x}$ can be expressed as a convex combination of a set $\{\mathbf{x}_i\}_i$ of feasible points in P_{SEP}^n . Then all the $\mathbf{x}_i$ must attain the same extreme 0–1 values of $\mathbf{x}$. Moreover, the remaining edges are compelled to the value $1/2$, since they are organized in disjoint odd cycles (specifically, triangles) and they must satisfy the degree constraints. Therefore, all the $\mathbf{x}_i$ are forced to coincide with $\mathbf{x}$: the convex combination is necessarily trivial, proving that $\mathbf{x}$ is an extreme point of P_{SEP}^n . $\square$

Now consider $\tilde{\mathbf{x}}$ of Definition 4.4: $\tilde{\mathbf{x}}$ coincides exactly with the original multigraph $\tilde{G} = (\tilde{V}, \tilde{E})$. The following claim establishes the integrality gap of the family $\mathcal{F}(\mathbf{x})$.

Claim 5.2 $\text{Gap}(\tilde{\mathbf{x}}) = 4/3$.

Proof By Lemma 4.8, $\text{Gap}(\tilde{\mathbf{x}}) \leq 4/3$, and we only need to prove $\text{Gap}(\tilde{\mathbf{x}}) \geq 4/3$. To this aim, we exhibit a specific cost achieving the desired lower bound. Let $\tilde{\mathbf{c}}$ be the cost on $\tilde{\mathbf{x}}$ which assigns 1 to every edge of $\tilde{E}$. Then, for every subset of edges $X \subseteq \tilde{E}$, the cost of X simply reduces to its cardinality: $\tilde{\mathbf{c}}(X) = |X|$; in particular, $\tilde{\mathbf{c}}\tilde{\mathbf{x}} = \tilde{\mathbf{c}}(\tilde{E}) = |\tilde{E}|$. Our goal is to prove $\text{MTJ}(\tilde{V}_{\text{odd}}, \tilde{\mathbf{c}}) \geq 1/3 \tilde{\mathbf{c}}\tilde{\mathbf{x}}$: once this is established, the claim would immediately follow, since

$$\text{Gap}(\tilde{\mathbf{x}}) \geq \frac{\text{CPP}(\tilde{\mathbf{c}})}{\tilde{\mathbf{c}}\tilde{\mathbf{x}}} = \frac{\tilde{\mathbf{c}}(\tilde{E}) + \text{MTJ}(\tilde{V}_{\text{odd}}, \tilde{\mathbf{c}})}{\tilde{\mathbf{c}}\tilde{\mathbf{x}}} \geq \frac{\tilde{\mathbf{c}}\tilde{\mathbf{x}} + 1/3 \tilde{\mathbf{c}}\tilde{\mathbf{x}}}{\tilde{\mathbf{c}}\tilde{\mathbf{x}}} = \frac{4}{3}.$$

First observe that $\tilde{V}_{\text{odd}} = \tilde{V}$, since $\tilde{G}$ is cubic. Let J be an arbitrary $\tilde{V}_{\text{odd}}$ -join. The nodes of J are exactly $\tilde{V}$, and, by definition of T-join, each of them must have odd degree in J , in particular, at least degree 1. Therefore, by the handshaking lemma, the number of edges in J must be at least $1/2 |\tilde{V}|$, thus $|J| \geq 1/2 |\tilde{V}|$. Applying the handshaking lemma again on the cubic graph $\tilde{G}$, we have $|\tilde{E}| = 1/2 \cdot 3 |\tilde{V}|$, so that $1/2 |\tilde{V}| = 1/3 |\tilde{E}|$, and ultimately $\tilde{\mathbf{c}}(J) = |J| \geq 1/2 |\tilde{V}| = 1/3 |\tilde{E}| = 1/3 \tilde{\mathbf{c}}\tilde{\mathbf{x}}$. Since this consideration applies to any $\tilde{V}$ -join, it must hold that $\text{MTJ}(\tilde{V}, \tilde{\mathbf{c}}) \geq 1/3 \tilde{\mathbf{c}}\tilde{\mathbf{x}}$, concluding the argument. $\square$

If we now consider the sequence of vertices $\mathcal{F}(\mathbf{x})$, its asymptotic gap tends to $\text{Gap}(\hat{\mathbf{x}})$ (by Theorem 3.3), which is at least as large as $\text{Gap}(\tilde{\mathbf{x}})$ (by Lemma 4.6). Therefore, the constructions developed in this section demonstrate that starting from any 2-edge-connected cubic multigraph the framework produces a non-trivial family of instances whose integrality gap converges at least to $4/3$. It is not difficult to prove that the class of such multigraphs is infinite.

Lemma 5.3 *There exist infinitely many (pairwise-non-isomorphic) 2-edge-connected cubic multigraphs.*

Proof We proceed by induction on the (even) number of nodes n .

Base case. For $n = 2$, consider the multigraph consisting of two nodes connected by three parallel edges: it is cubic and 2-edge-connected.

Inductive step. Consider a 2-edge-connected cubic multigraph with n nodes. We construct a new multigraph with $n + 2$ nodes as follows. Choose two distinct edges of the graph; subdivide each of them by inserting a new node, replacing each chosen edge with a path of two edges; finally, connect the two newly introduced nodes by an additional edge. All nodes of the resulting graph have degree 3 by construction. Moreover, 2-node-connectivity is preserved. Indeed, by the characterization of 2-edge-connected graphs via ear decomposition, the operation of subdividing an edge corresponds to extending an existing ear, while the new edge between the two added nodes can be interpreted as an additional open ear. Therefore, the resulting multigraph is cubic, 2-edge-connected, and has $n + 2$ nodes. $\square$




Fig. 5 Cubic, 2-edge-connected multigraphs $\tilde{G}$ ($= \tilde{x}$), with the corresponding vertices x , produced as in the proof of Lemma 5.3.

Beyond its intrinsic interest, this section highlights a notable feature of our framework: it generalizes several constructions from the literature. In particular, the classical “envelope instances” and “tetrahedron instances” arise naturally from the very first steps of the inductive construction described in the previous proof, corresponding to the cases $n = 2$ and $n = 4$, as illustrated in Figure 5. These instances can thus be placed within a broader and more systematic perspective, confirming the breadth and versatility of our approach.

5.2 Local modification inducing high integrality gaps

This section completes our search for families with high integrality gap. We show that, for *any* instance, a simple local modification produces an associated family whose asymptotic integrality gap is at least $4/3$.

Lemma 5.4 *Let $\mathbf{x} \in P_{SEP}^n$, and let $\tilde{\mathbf{x}}$ be the associated multigraph. If $\tilde{\mathbf{x}}$ has a node of degree 3, then $\text{Gap}(\tilde{\mathbf{x}}) = 4/3$.*

Proof Let v be a node of $\tilde{\mathbf{x}}$ of degree 3, and define the cost $\tilde{c}$ which assigns 1 to the edges incident to v and 0 otherwise. This cost achieves $\tilde{c}\tilde{\mathbf{x}} = 3$. Moreover, any feasible walk $\tilde{\mathbf{w}}$ for $\text{Gap}(\tilde{\mathbf{x}})$ must use at least 4 times the edges incident to v , yielding $\tilde{c}\tilde{\mathbf{w}} \geq 4$. This proves $\text{Gap}(\tilde{\mathbf{x}}) \geq 4/3$, thence the equality $\text{Gap}(\tilde{\mathbf{x}}) = 4/3$, by Lemma 4.8. $\square$

Given this lemma, it suffices to enforce the presence of a node of degree 3 in $\tilde{\mathbf{x}}$. A convenient way to achieve this is by inserting a $1/2$ -triangle into $\mathbf{x}$, since in the limit construction it collapses into a single node of $\tilde{\mathbf{x}}$ of degree 3. To this end, consider the structure illustrated in Figure 6, composed of two $1/2$ -triangles joined by 1-edges. Following [15], we refer to it as a *triangle gadget*. Replacing a 1-edge of $\mathbf{x}$ with a triangle gadget yields a new instance $\mathbf{x}'$ whose associated family has asymptotic integrality gap (at least) $4/3$, by Lemmas 4.6 and 5.4.




Fig. 6 The triangle gadget, used to insert two $1/2$ -triangles into a vertex.

In light of the abundance of constructions attaining this limit, the focus naturally shifts from the value $4/3$ itself to the speed at which it is reached. A compelling research direction is therefore the study of instances maximizing the rate of convergence toward $4/3$ as the number of nodes grows. From this perspective, Question 1 asks whether fractional negligible vertices occupy a distinguished position among the constructions whose asymptotic behavior attains this bound.

6 Conclusion

In this paper, we introduced an *Asymptotic Framework* for studying the integrality gap of the TSP. Rather than analyzing individual instances, our framework studies the asymptotic behavior of entire families of vertices of the subtour elimination polytope, providing a new perspective on the longstanding $4/3$ -conjecture.

This viewpoint yields progress on both sides of the conjecture. On the one hand, we introduced a limit quantity that captures the asymptotic integrality gap of a family of instances and proved that it is bounded above by $4/3$ for the class of *fractional negligible* instances. On the other hand, we considered a simple and general transformation that, starting from an arbitrary instance, produces a family whose asymptotic integrality gap is at least $4/3$.

The most intriguing direction left unresolved by this work is Question 1, asking whether the property of fractional negligibility is indeed necessary for vertices of P_{SEP}^n to attain the largest possible integrality gap. If this was the case, the corresponding upper bound proved in this manuscript would resolve the $4/3$ -conjecture.

Acknowledgments

This work was supported by the Swiss National Science Foundation project n. 200021_212929/1 “Computational methods for integrality gaps analysis”.

We also thank Professor Ahmad Abdi for insightful lectures on T-joins at the ISCO 2026 Spring School and for a precious discussion that significantly helped in shortening the proof of Lemma 4.8.

References

- [1] Dantzig, G., Fulkerson, R., Johnson, S.: Solution of a large-scale traveling salesman problem. *Journal of the operations research society of America* **2**(4), 393–410 (1954)
- [2] Benoit, G., Boyd, S.: Finding the exact integrality gap for small traveling salesman problems. *Mathematics of Operations Research* **33**, 921–931 (2008) <https://doi.org/10.1287/moor.1080.0337>
- [3] Villa, T., Vercesi, E., Barta, J., Mastrolilli, M.: The integrality gap of the traveling salesman problem is $4/3$ if the LP solution has at most $n+6$ non-zero components. In: Dey, S.S., Di Summa, M., Salvagnin, D. (eds.) *Integer Programming and Combinatorial Optimization*, pp. 128–143. Springer, Cham (2026)
- [4] Cook, W., Hougardy, S., Petrich, M.: Extending Exact Integrality Gap Computations for the Metric TSP (2026). <https://arxiv.org/abs/2603.12995>
- [5] Wolsey, L.A.: Heuristic analysis, linear programming and branch and bound. In: Rayward-Smith, V.J. (ed.) *Combinatorial Optimization II*, pp. 121–134. Springer, Berlin, Heidelberg (1980). <https://doi.org/10.1007/BFb0120913>
- [6] Christofides, N.: Worst-case analysis of a new heuristic for the travelling salesman problem. *Operations Research Forum* **3** (1976)
- [7] Karlin, A.R., Klein, N., Oveis Gharan, S.: A (slightly) improved approximation algorithm for metric TSP. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. STOC 2021*, pp. 32–45. Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3406325.3451009>
- [8] Karlin, A.R., Klein, N., Oveis Gharan, S.: A (slightly) improved bound on the integrality gap of the subtour LP for TSP. In: *63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 832–843. IEEE Computer Society, Denver, CO, USA (2022). <https://doi.org/10.1109/FOCS54457.2022.00084>
- [9] Goemans, M.X.: Worst-case comparison of valid inequalities for the TSP. *Mathematical Programming* **69**(2), 335–349 (1995)

- [10] Boyd, S., Elliott-Magwood, P.: Structure of the extreme points of the subtour elimination polytope of the STSP. In: RIMS Kokyuroku Bessatsu, pp. 33–47 (2010). <https://api.semanticscholar.org/CorpusID:15182527>
- [11] Schalekamp, F., Williamson, D.P., Zuylen, A.: 2-matchings, the traveling salesman problem, and the subtour LP: A proof of the Boyd-Carr conjecture. *Mathematics of Operations Research* **39**(2), 403–417 (2014)
- [12] Karlin, A.R., Klein, N., Oveis Gharan, S.: An improved approximation algorithm for TSP in the half integral case. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. STOC 2020, pp. 28–39. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3357713.3384273>
- [13] Gupta, A., Lee, E., Li, J., Mucha, M., Newman, H., Sarkar, S.: Matroid-based TSP rounding for half-integral solutions. *Mathematical Programming* **206**(1–2), 541–576 (2024) <https://doi.org/10.1007/s10107-024-02065-4>
- [14] Klein, N., Taziki, M.: Dual charging for half-integral TSP. In: Ene, A., Chattopadhyay, E. (eds.) *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 353, pp. 21–12122. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2025). <https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2025.21>
- [15] Boyd, S., Carr, R.: Finding low cost TSP and 2-matching solutions using certain half-integer subtour vertices. *Discrete Optimization* **8**(4), 525–539 (2011) <https://doi.org/10.1016/j.disopt.2011.05.002>
- [16] Boyd, S., Sitters, R., Ster, S., Stougie, L.: TSP on cubic and subcubic graphs. In: *Integer Programming and Combinatorial Optimization*, pp. 65–77. Springer, Berlin, Heidelberg (2011)
- [17] Mömke, T., Svensson, O.: Removing and adding edges for the traveling salesman problem. *Journal of the ACM* **63**(1), 2–1228 (2016) <https://doi.org/10.1145/2739008>
- [18] Boyd, S., Sebő, A.: The salesman’s improved tours for fundamental classes. *Mathematical Programming* **186**(1), 289–307 (2021)
- [19] Jin, B., Klein, N., Williamson, D.P.: A $4/3$ -approximation algorithm for half-integral cycle cut instances of the TSP. *Mathematical Programming* **209** (2025) <https://doi.org/10.1007/s10107-025-02193-5>
- [20] Hougardy, S.: On the integrality ratio of the subtour LP for euclidean TSP. *Operations Research Letters* **42**(8), 495–499 (2014)

- [21] Hougardy, S., Zhong, X.: Hard to solve instances of the euclidean traveling salesman problem. *Mathematical Programming Computation* **13**, 51–74 (2021)
- [22] Zhong, X.: Lower bounds on the integrality ratio of the subtour LP for the traveling salesman problem. *Discrete Applied Mathematics* **365**, 109–129 (2025)
- [23] Vercesi, E., Gualandi, S., Mastrolilli, M., Gambardella, L.M.: On the generation of metric TSP instances with a large integrality gap by branch-and-cut. *Mathematical Programming Computation* **15**(2), 389–416 (2023)
- [24] Edmonds, J., Johnson, E.L.: Matching, euler tours and the chinese postman. *Mathematical Programming* **5**(1), 88–124 (1973) <https://doi.org/10.1007/BF01580113>