

Order-2 Tightness of Block-Sparse SOS Relaxations for One-Layer ReLU Network Verification with a Matching Input-Sharing Graph

Godai Azuma Sunyoung Kim Makoto Yamashita*

Abstract: Azuma, Kim, and Yamashita formulated the verification problem for one-layer ReLU networks as a quadratically constrained quadratic program and established tight semidefinite relaxations for the edgeless case and for one-unit settings. In this work, we represent the sharing pattern of undecided ReLUs over a box input set through an input-sharing graph and focus on the case where this graph is a matching. We then derive an explicit, checkable sufficient condition for the tightness of the order-2 block-sparse SOS relaxation associated with the connected-component decomposition of this graph. Under the matching assumption, the global problem decomposes into isolated-vertex blocks and single-edge blocks. The key difficulty, which is absent from the edgeless case, is establishing tightness for a two-unit edge block. For regular rank-one edges, we show that the convex hull of each two-unit local set can be described exactly by two reduced one-unit hulls coupled through a common shared scalar. Combining the one-unit tightness result of Azuma et al. with Farkas’ lemma and affine elimination, we obtain a local order-2 certificate for each edge block. Isolated-vertex blocks reduce to one-unit problems over box input sets and are therefore handled at the same order. We prove that, when the input-sharing graph is a matching and every edge satisfies the regular rank-one condition, the order-2 block-sparse SOS relaxation is tight. This extends the tight sparse relaxation result for the edgeless case to the first sparse setting with a nontrivial two-unit interaction.

Keywords: ReLU network verification; sums-of-squares relaxation; semidefinite programming; sparse polynomial optimization; input-sharing graph.

Mathematics Subject Classification: 90C22; 90C26; 90C90; 68T07.

1 Introduction

Using neural networks in safety-critical applications requires rigorous bounds on their outputs, or on associated specification functions, over given input sets. Such bounds are essential for certifying, for instance, that admissible input perturbations do not lead to misclassification, or that a system trajectory remains within a designated safe region. Numerous verification methods have been proposed, including complete approaches based on satisfiability modulo theories (SMT) and Boolean satisfiability (SAT), abstract interpretation, branch-and-bound methods, mixed-integer optimization, semidefinite relaxations, and hybrid methods [4–6, 8–10, 19]. For surveys, see Liu et al. [14], Meng et al. [16]; for a tutorial perspective in control, see Everett [7]. A basic theoretical question is how accurately a tractable relaxation approximates the original verification problem and which structural conditions ensure tightness.

In this work, we consider the one-layer ReLU network

$$\begin{aligned}v &= Wu + b, \\x &= \text{ReLU}(v), \\y &= Ax + a\end{aligned}$$

*Corresponding author. The research of Godai Azuma was supported by JSPS KAKENHI Grant Number JP24K20738. The research of Makoto Yamashita was partially supported by JSPS KAKENHI Grant Number 24K14836 and 26K02868.

where $u \in \mathbb{R}^{n_0}$ is the input, $v \in \mathbb{R}^{n_1}$ is the pre-activation vector, $x \in \mathbb{R}^{n_1}$ is the hidden-layer output, $y \in \mathbb{R}^{n_2}$ is the final output, and $x = \text{ReLU}(v)$ denotes $x_j = \max\{v_j, 0\}$ for $j = 1, \dots, n_1$. We consider the box input set

$$\mathcal{U} := \{u \in \mathbb{R}^{n_0} \mid \underline{u} \leq u \leq \bar{u}\}.$$

We study the verification problem of maximizing a linear specification $c^\top y$ over this set:

$$\gamma^* := \max\{c^\top y \mid u \in \mathcal{U}, v = Wu + b, x = \text{ReLU}(v), y = Ax + a\}. \quad (1.1)$$

This form is a basic model that can represent, with auxiliary variables if needed, many specifications such as componentwise output bounds and margin-type inequalities between outputs.

As each ReLU constraint can be represented by linear inequalities and one quadratic equality, problem (1.1) is a quadratically constrained quadratic program (QCQP) and hence a polynomial optimization problem. One can then apply semidefinite relaxations based on sums of squares (SOS) and moment techniques [11, 13, 18]. Moreover, sparse SOS relaxations that exploit correlative or term sparsity [20–22] are relevant to structured polynomial optimization. In a related direction, Marumo et al. [15] study a T-SDP relaxation for constrained polynomial optimization and show that it can be transformed into a standard SDP relaxation with block-diagonal structure. In the setting of neural network verification, Newton and Papachristodoulou [17] apply sparse polynomial optimization and the Positivstellensatz to obtain tighter bounds for ReLU, sigmoid, and tanh networks.

Our goal is to identify graph-based conditions under which a sparse relaxation is tight for one-layer ReLU verification. The starting point for our analysis is the work of Azuma et al. [2], which established tight semidefinite relaxations for one-layer ReLU networks in three settings: the one-unit case, the ellipsoidal-input case, and the box-input case with a diagonal weight matrix. When the input set is a box and the weight matrix is diagonal, each undecided unit depends on input variables that are not shared with any other undecided unit. The problem therefore decomposes unit by unit. The diagonal-weight setting is therefore a special case of an edgeless input-sharing graph.

We study the matching case, in which each undecided ReLU shares input variables with at most one other undecided ReLU. Under this assumption, every connected component is either an isolated vertex or a single edge. Existing one-unit results handle isolated vertices, whereas a single edge requires a new tightness argument for two units coupled through a shared scalar.

For a regular rank-one edge, we describe the convex hull of the two-unit block and derive an SOS certificate at relaxation order 2. Together with the reduced-objective condition introduced later, these local certificates prove tightness of the block-sparse relaxation for matching input-sharing graphs. Isolated vertices are handled by the one-unit result of Azuma et al. [2].

The significance of this result lies in replacing a dense global certificate with local certificates supported on the graph structure of the problem. Tightness is established component-wise, through one-unit blocks and two-unit edge blocks. In this sense, the paper contributes not only to ReLU network verification, but also to graph-structured nonconvex polynomial optimization more broadly.

The main contributions of the paper are as follows.

- (i) We introduce the input-sharing graph to describe coupling among undecided ReLUs and relate the diagonal-weight case of Azuma et al. [2] to the edgeless setting.

- (ii) Under the matching assumption, we decompose the reduced verification problem into isolated-vertex and single-edge blocks and formulate the corresponding block-sparse SOS relaxation.
- (iii) For a regular rank-one edge, we obtain an exact convex-hull description and derive a local certificate at relaxation order 2.
- (iv) By assembling the one-unit and two-unit certificates, we prove global tightness under the matching and reduced-objective conditions.

The paper is organized as follows. Section 2 reviews the semialgebraic formulation of one-layer ReLU verification, the basic machinery of SOS relaxations, and the one-unit results of Azuma et al. [2]. Section 3 introduces determined units, undecided ReLUs, the input-sharing graph, matching, and regular rank-one edges. Section 4 shows that under the matching condition the problem decomposes into isolated-vertex components and single-edge components, and defines the order-2 block-sparse SOS relaxation used in this paper. Section 5 proves local tightness for a regular rank-one edge and then states the main theorem for the matching case. Section 6 concludes the paper. Appendix A gives a small example that illustrates the input-sharing graph, the component decomposition, the regular rank-one condition, and the resulting tightness statement.

2 Problem Formulation and Sparse Polynomial Relaxations via SOS

2.1 Semialgebraic formulation of the verification problem and reduction to polynomial optimization

We consider the one-layer ReLU network

$$\begin{aligned} v &= Wu + b, \\ x &= \text{ReLU}(v), \\ y &= Ax + a, \end{aligned}$$

where $u \in \mathbb{R}^{n_0}$ is the input, $v, x \in \mathbb{R}^{n_1}$ are the pre-activation variables and ReLU outputs, respectively, and $y \in \mathbb{R}^{n_2}$ is the output. The data are $W \in \mathbb{R}^{n_1 \times n_0}$, $b \in \mathbb{R}^{n_1}$, $A \in \mathbb{R}^{n_2 \times n_1}$, and $a \in \mathbb{R}^{n_2}$. We assume that the input set is the box

$$\mathcal{U} := \{u \in \mathbb{R}^{n_0} \mid \underline{u} \leq u \leq \bar{u}\}.$$

Our basic problem is the upper-bound evaluation problem for a linear specification vector $c \in \mathbb{R}^{n_2}$,

$$\gamma^* := \max\{c^\top y \mid u \in \mathcal{U}, v = Wu + b, x = \text{ReLU}(v), y = Ax + a\}. \quad (2.1)$$

This form represents not only the maximization of an output component but also, for example, componentwise output bounds and margin-type constraints defined by differences between competing classes.

For each unit $j = 1, \dots, n_1$, the ReLU constraint $x_j = \max\{v_j, 0\}$ is equivalent to

$$x_j \geq 0, \quad x_j - v_j \geq 0, \quad x_j(x_j - v_j) = 0.$$

Hence problem (2.1) can be written as a semialgebraic optimization problem consisting of affine relations, linear inequalities, and quadratic equalities. Let us collect all variables as

$$z := (u, v, x, y).$$

We write the inequality constraints as $g_1(z), \dots, g_m(z) \geq 0$ and the equality constraints as $h_1(z) = \dots = h_s(z) = 0$. Then the feasible set is

$$K := \{z \in \mathbb{R}^{n_0+2n_1+n_2} \mid g_\ell(z) \geq 0 \ (\ell = 1, \dots, m), \ h_r(z) = 0 \ (r = 1, \dots, s)\}$$

and the objective function $c^\top y$ is a linear polynomial in z .

Giving a valid upper bound γ for the verification problem is equivalent to proving that the polynomial

$$p_\gamma(z) := \gamma - c^\top y$$

is nonnegative on K . Therefore, the verification problem reduces to the question of how to certify the nonnegativity of a polynomial on a semialgebraic set by a computationally tractable certificate. This is precisely the starting point for applying Lasserre-type moment-SOS hierarchies and Putinar-type Positivstellensätze [11, 18]. Polynomial and SOS-based approaches to neural network verification are studied, for example, by Newton and Pachristodoulou [17].

2.2 SOS certificates, degree-bounded relaxations, and problem-specific sparse certificates

A polynomial $q(z)$ is said to be SOS if there exist finitely many polynomials $r_1, \dots, r_N$ such that

$$q(z) = \sum_{i=1}^N r_i(z)^2.$$

Since every SOS polynomial is globally nonnegative and SOS feasibility admits a Gram matrix representation, SOS relaxations can be formulated and solved as semidefinite programs [11, 20]. Their dual counterparts are moment relaxations, expressed in terms of moment and localizing matrices.

For the inequality constraints, define the quadratic module

$$\mathcal{M}(g) := \left\{ \sigma_0 + \sum_{\ell=1}^m \sigma_\ell g_\ell \mid \sigma_0, \sigma_1, \dots, \sigma_m \text{ are SOS polynomials} \right\},$$

and for the equality constraints, define the ideal

$$\mathcal{I}(h) := \left\{ \sum_{r=1}^s \tau_r h_r \mid \tau_1, \dots, \tau_s \text{ are polynomials} \right\}.$$

If

$$p_\gamma \in \mathcal{M}(g) + \mathcal{I}(h), \tag{2.2}$$

then $p_\gamma(z) \geq 0$ for every $z \in K$, so γ is a valid upper bound for problem (2.1). In this paper, we call an expression of the form (2.2) a certificate for p_γ .

For computational purposes, we use degree-bounded truncations. Given a relaxation order $\omega \in \mathbb{N}$, define

$$\mathcal{M}_\omega(g) := \left\{ \sigma_0 + \sum_{\ell=1}^m \sigma_\ell g_\ell \in \mathcal{M}(g) \mid \deg(\sigma_0) \leq 2\omega, \deg(\sigma_\ell g_\ell) \leq 2\omega \ (\ell = 1, \dots, m) \right\},$$

$$\mathcal{I}_\omega(h) := \left\{ \sum_{r=1}^s \tau_r h_r \in \mathcal{I}(h) \mid \deg(\tau_r h_r) \leq 2\omega \ (r = 1, \dots, s) \right\}$$

and let

$$\gamma^{(\omega)} := \inf \{ \gamma \in \mathbb{R} \mid p_\gamma \in \mathcal{M}_\omega(g) + \mathcal{I}_\omega(h) \}.$$

Then $\gamma^* \leq \gamma^{(\omega)}$ always holds. Moreover, if the feasible set is compact and the quadratic module is Archimedean, Putinar’s Positivstellensatz guarantees that $\gamma^{(\omega)} \downarrow \gamma^*$ as $\omega \rightarrow \infty$ [11, 18]. The number ω is called the relaxation order. As ω increases, higher-degree polynomials are allowed in the SOS multipliers σ_0, σ_ℓ , and in the polynomial multipliers τ_r . Consequently, the number of unknown coefficients grows rapidly, and the resulting certificate search problem becomes increasingly large. The central question of this paper is to identify conditions under which this bound is already tight at a low relaxation order, in particular at $\omega = 2$.

The relaxation used in the remainder of the paper is not the generic dense SOS relaxation itself but a block-sparse SOS certificate class tailored to the graph-structured sparsity of the ReLU verification problem. This perspective is consistent with the sparse-SOS literature exploiting correlative or term sparsity [20–22], but our aim is not to develop a general hierarchy. Instead, in Sections 3 and 4 we define a problem-specific certificate class based on the input-sharing graph and its local blocks, and then prove its tightness under the matching assumption. Thus, the sparse relaxation studied here is a problem-specific block-sparse SOS relaxation designed for the local geometry of one-layer ReLU verification.

2.3 The one-unit tightness result as a fundamental component of the analysis

Within SDP-based neural network verification, DeepSDP is a representative example [8], and Zhang [23] provides a geometric analysis of tightness for such relaxations. The present paper is motivated by the work of Azuma et al. [2], who proved tightness of semidefinite relaxations for one-layer ReLU networks in three cases: a single neuron, an ellipsoidal input set, and a diagonal weight matrix over a box input set. For our analysis, we use the following one-unit consequences of the tightness results of Azuma et al. [2].

Here, a local problem means the problem associated with one connected component, while a unit means a single ReLU, say $x_j = \text{ReLU}(v_j)$. Thus, a one-unit block may still have a multidimensional input vector.

(A1) one-unit interval-input case. The local problem contains a single unit, and the effective input reduces to a scalar s over an interval

$$U = [\underline{s}, \bar{s}] \subset \mathbb{R}.$$

In our notation, the input data consist of a 1×1 matrix $W \in \mathbb{R}^{1 \times 1}$, a scalar $b \in \mathbb{R}$, and the interval input set U , and the feasible set has the form

$$v = Ws + b, \quad x = \text{ReLU}(v), \quad s \in U.$$

(A2) ellipsoidal-input case. The local problem involves a single unit, and the input set is an ellipsoid

$$U = \{u \in \mathbb{R}^m \mid (u - u^c)^\top Q^{-1}(u - u^c) \leq 1\}$$

with $Q \succ 0$, while the pre-activation is given by the affine function

$$v = w^\top u + b.$$

(A3) one-unit rectangular-input case. The local problem contains a single unit, and the input set is a rectangle

$$U = [\underline{u}_1, \bar{u}_1] \times \cdots \times [\underline{u}_m, \bar{u}_m] \subset \mathbb{R}^m.$$

The pre-activation is affine in the input:

$$v = w^\top u + b,$$

This is the form to which isolated-vertex blocks and the reduced one-unit problems introduced later are reduced. The interval-input case **(A1)** is the special case $m = 1$ of the rectangular-input setting.

These results will be applied to isolated vertices and to the reduced one-unit problems arising from an edge block. The new difficulty is to connect the resulting one-unit certificates across the shared scalar.

3 Preprocessing, Undecided ReLUs, and the Input-Sharing Graph

In this section, we introduce the basic notation for the block-sparse SOS relaxation defined in Section 4, reflecting the structure of the one-layer ReLU verification problem studied in this paper. We first remove the determined ReLUs by interval evaluation and derive a reduced problem containing only undecided ReLUs. We then define the graph $H = (\mathcal{J}_u, E)$ that records input sharing among undecided ReLUs, and introduce the matching condition and the shared rank q_e . Finally, we define regular rank-one edges, which will be used in Section 5, and clarify what decomposes componentwise under the matching assumption.

3.1 Elimination of determined ReLUs and the reduced verification problem

Assume that for each hidden unit $j = 1, \dots, n_1$, preprocessing provides an interval bound

$$[\underline{v}_j, \bar{v}_j]$$

for the pre-activation v_j . We then partition the unit indices as

$$\begin{aligned} \mathcal{J}^- &:= \{j \in \{1, \dots, n_1\} \mid \bar{v}_j \leq 0\}, \\ \mathcal{J}^+ &:= \{j \in \{1, \dots, n_1\} \mid \underline{v}_j \geq 0\}, \\ \mathcal{J}_u &:= \{j \in \{1, \dots, n_1\} \mid \underline{v}_j < 0 < \bar{v}_j\}. \end{aligned}$$

Here $\mathcal{J}^- \cup \mathcal{J}^+$ is the set of determined ReLUs and $\mathcal{J}_u$ is the set of undecided ReLUs. In the verification literature, the same partition is often described as stable versus unstable ReLUs. In this paper, however, we consistently use the terminology “determined ReLUs” and “undecided ReLUs.” For $j \in \mathcal{J}^-$ the ReLU output is fixed at $x_j = 0$, whereas for $j \in \mathcal{J}^+$, it is given by $x_j = v_j$. Thus, all essential nonlinearity is confined to $\mathcal{J}_u$.

Let $W_{\mathcal{J}_u, \cdot}$ denote the submatrix of W consisting of the rows indexed by $\mathcal{J}_u$, and let $A_{\cdot, \mathcal{J}_u}$ denote the submatrix of A consisting of the columns indexed by $\mathcal{J}_u$. Define

$$\begin{aligned}\bar{W} &:= W_{\mathcal{J}_u, \cdot}, & \bar{b} &:= b_{\mathcal{J}_u}, \\ \bar{A} &:= A_{\cdot, \mathcal{J}_u}, & \bar{B} &:= A_{\cdot, \mathcal{J}^+} W_{\mathcal{J}^+, \cdot}, & \bar{a} &:= a + A_{\cdot, \mathcal{J}^+} b_{\mathcal{J}^+}.\end{aligned}$$

Then, the verification problem (2.1) can be rewritten equivalently as

$$\gamma^* = \max \left\{ c^\top (\bar{A}x + \bar{B}u + \bar{a}) \mid u \in \mathcal{U}, v = \bar{W}u + \bar{b}, x = \text{ReLU}(v) \right\}.$$

Here the vectors v and x are indexed only by $\mathcal{J}_u$. Therefore, in the subsequent discussion, the essential semialgebraic constraints are simply

$$v_j = \bar{W}_{j \cdot} u + \bar{b}_j, \quad x_j \geq 0, \quad x_j - v_j \geq 0, \quad x_j(x_j - v_j) = 0 \quad (j \in \mathcal{J}_u).$$

Using the notation of Section 2, let

$$p_\gamma(u, v, x) := \gamma - c^\top (\bar{A}x + \bar{B}u + \bar{a}).$$

Our aim is to understand when the nonnegativity of this p_γ can be captured tightly by SOS certificates compatible with a sparse block structure. This requires an explicit description of the input coordinates shared by the undecided ReLUs in $\mathcal{J}_u$.

3.2 Input supports, input-sharing graph, matching, and shared rank

For each undecided ReLU $j \in \mathcal{J}_u$, define the set of input coordinates on which its pre-activation depends by

$$S_j := \text{supp}(\bar{W}_{j \cdot}) = \{k \in \{1, \dots, n_0\} \mid \bar{W}_{jk} \neq 0\}.$$

Thus, the ReLU constraint for unit j depends not on the entire input vector u but only on the subvector u_{S_j} .

Based on these input supports, we introduce an undirected graph describing the interaction among undecided ReLUs.

Definition 3.1 (input-sharing graph). For distinct vertices $i, j \in \mathcal{J}_u$, define

$$\{i, j\} \in E \iff S_i \cap S_j \neq \emptyset. \quad (3.1)$$

Then, $H = (\mathcal{J}_u, E)$ is called the input-sharing graph of the undecided ReLUs.

Condition (3.1) indicates whether two undecided ReLUs are coupled through common input coordinates. Therefore, the graph H identifies where the essential interactions appear in a global block-sparse SOS certificate. In particular, when H is edgeless, no undecided ReLU shares inputs with any other undecided ReLU.

For an edge $e = \{i, j\} \in E$, we write the shared input set as

$$S_e := S_i \cap S_j.$$

We also impose the matching condition as the main graph-theoretic assumption.

Definition 3.2 (matching). The input-sharing graph $H = (\mathcal{J}_u, E)$ is called a *matching* if

$$\deg_H(j) \leq 1 \quad (\forall j \in \mathcal{J}_u)$$

holds.

Here, $\deg_H(j)$ is the number of other undecided ReLUs that share at least one input coordinate with unit j :

$$\deg_H(j) = |\{i \in \mathcal{J}_u \setminus \{j\} \mid S_i \cap S_j \neq \emptyset\}|.$$

Thus, the matching condition means that each undecided ReLU shares input variables with at most one other undecided ReLU.

Under the matching assumption, every connected component is either an isolated vertex or a single edge. This allows compatibility and assembly to be treated component by component, rather than over the entire graph at once. Accordingly, the matching assumption is a central hypothesis in the results of Sections 4 and 5.

To measure the effective dimension of the shared part of each edge, we introduce the shared rank.

Definition 3.3 (shared rank). For an edge $e = \{i, j\} \in E$, define

$$q_e := \dim \text{span}\{\bar{W}_{i,S_e}^\top, \bar{W}_{j,S_e}^\top\}.$$

We call q_e the shared rank of e .

Here $\bar{W}_{i,S_e}$ and $\bar{W}_{j,S_e}$ are the coefficient vectors obtained by restricting rows i and j to the shared input set S_e . Thus q_e measures the number of effective directions through which the two undecided units depend affinely on the shared inputs. In particular, when

$$q_e \leq 1,$$

the reduction leaves only a single common scalar to represent the shared part. This condition is the starting point for the order-2 local tightness result for two-unit blocks in Section 5.

3.3 Regular rank-one edges and the componentwise decomposition under matching

For a subset $I \subseteq \{1, \dots, n_0\}$, write the projection of the input box onto that coordinate subspace as

$$\mathcal{U}_I := \{u_I \in \mathbb{R}^I \mid \underline{u}_I \leq u_I \leq \bar{u}_I\}.$$

Fix an edge $e = \{i, j\} \in E$. For each $h \in \{i, j\}$, define the private index set by

$$P_h := S_h \setminus S_e.$$

Then the pre-activation decomposes as

$$v_h = \bar{W}_{h,P_h} u_{P_h} + \bar{W}_{h,S_e} u_{S_e} + \bar{b}_h \quad (h = i, j).$$

Based on this decomposition, define for each $h \in \{i, j\}$ the bounds of the private and shared parts by

$$\begin{aligned} \underline{p}_{h,e} &:= \min_{u_{P_h} \in \mathcal{U}_{P_h}} (\bar{W}_{h,P_h} u_{P_h} + \bar{b}_h), & \bar{p}_{h,e} &:= \max_{u_{P_h} \in \mathcal{U}_{P_h}} (\bar{W}_{h,P_h} u_{P_h} + \bar{b}_h), \\ \underline{s}_{h,e} &:= \min_{u_{S_e} \in \mathcal{U}_{S_e}} \bar{W}_{h,S_e} u_{S_e}, & \bar{s}_{h,e} &:= \max_{u_{S_e} \in \mathcal{U}_{S_e}} \bar{W}_{h,S_e} u_{S_e}. \end{aligned}$$

Definition 3.4 (regular rank-one edge). An edge $e = \{i, j\} \in E$ is called a *regular rank-one edge* if the following two conditions hold:

(M3a) **rank-one shared direction:**

$$q_e \leq 1. \quad (3.2)$$

(M3b) **edge regularity:**

$$\underline{p}_{h,e} + \bar{s}_{h,e} < 0 < \bar{p}_{h,e} + \underline{s}_{h,e} \quad (h = i, j). \quad (3.3)$$

Condition (3.2) ensures that the shared part is represented by a single scalar. Condition (3.3) states that, regardless of how the shared variable u_{S_e} is fixed within its admissible range, the private variables can still make the pre-activation of each unit take both negative and positive values. Thus a regular rank-one edge is a local condition expressing that the shared structure is effectively one-dimensional and that both units remain undecided throughout the entire edge.

Proposition 3.5 (operational meaning of edge regularity). *Fix an edge $e = \{i, j\} \in E$ and $h \in \{i, j\}$. Then the following statements are equivalent:*

- (a) $\underline{p}_{h,e} + \bar{s}_{h,e} < 0 < \bar{p}_{h,e} + \underline{s}_{h,e}$ holds;
- (b) for every $u_{S_e} \in \mathcal{U}_{S_e}$, there exist $u_{P_h}^-, u_{P_h}^+ \in \mathcal{U}_{P_h}$ such that

$$\bar{W}_{h,P_h} u_{P_h}^- + \bar{W}_{h,S_e} u_{S_e} + \bar{b}_h < 0, \quad \bar{W}_{h,P_h} u_{P_h}^+ + \bar{W}_{h,S_e} u_{S_e} + \bar{b}_h > 0.$$

In other words, even after the shared part is fixed, unit h can still attain both negative and positive pre-activation values by varying its private part.

Proof. Fix $u_{S_e} \in \mathcal{U}_{S_e}$. Then, the shared term $\bar{W}_{h,S_e} u_{S_e}$ is a point in the interval $[\underline{s}_{h,e}, \bar{s}_{h,e}]$. On the other hand, as u_{P_h} varies over $\mathcal{U}_{P_h}$, the private term $\bar{W}_{h,P_h} u_{P_h} + \bar{b}_h$ ranges over the entire interval $[\underline{p}_{h,e}, \bar{p}_{h,e}]$. Hence, for the fixed shared term, the condition that the pre-activation v_h can attain both negative and positive values is equivalent to

$$\underline{p}_{h,e} + \bar{W}_{h,S_e} u_{S_e} < 0 < \bar{p}_{h,e} + \bar{W}_{h,S_e} u_{S_e}.$$

This requirement must hold uniformly for all $u_{S_e} \in \mathcal{U}_{S_e}$, and equivalently, it is characterized by the following worst-case inequality.

$$\underline{p}_{h,e} + \bar{s}_{h,e} < 0 < \bar{p}_{h,e} + \underline{s}_{h,e}.$$

□

Proposition 3.5 provides a direct interpretation of the two requirements: condition (3.2) makes the shared dependence one-dimensional, while condition (3.3) ensures that each unit can attain both signs after the shared input is fixed. This excludes cases in which the shared coupling effectively determines one of the units.

Indeed, if $P_h = \emptyset$, so that no private variables are present, then condition (3.3) requires the shared part itself to cross zero. Conversely, if $\bar{W}_{h,S_e} = 0$, so that the shared part does not affect unit h , then (3.2) is automatically satisfied, and the analysis reduces to a one-unit problem involving only the private variables.

Finally, to connect the above discussion with the block-sparse SOS certificates defined in Section 4, for each undecided ReLU $j \in \mathcal{J}_u$, we define the one-unit local variable tuple by

$$\zeta_j := (u_{S_j}, v_j, x_j)$$

and for each edge $e = \{i, j\} \in E$ we define the two-unit local variable tuple by

$$\zeta_e := (u_{S_i \cup S_j}, v_i, x_i, v_j, x_j).$$

These local tuples serve as the blocks in the subsequent discussion. Section 4 assembles the componentwise certificates, and Section 5 proves the edge-local exactness required for regular rank-one edges.

4 Component Decomposition under the Matching Assumption

In this section, we study the case in which the input-sharing graph $H = (\mathcal{J}_u, E)$ introduced in Section 3 is a matching. Under this assumption, all nonlinear coupling is confined to isolated-vertex components and single-edge components. This structure allows us to decompose both the feasible set and the objective, and to prove an assembly result that converts local order-2 certificates into a global certificate.

4.1 Component decomposition induced by the matching assumption

We first note that, under the matching assumption, every connected component has at most two vertices.

Proposition 4.1. *Assume that the input-sharing graph $H = (\mathcal{J}_u, E)$ is a matching. Then, every connected component of H is either an isolated vertex or a single edge.*

Proof. By definition of a matching, each vertex $j \in \mathcal{J}_u$ is incident to at most one edge. Hence, if some connected component contained three or more vertices, an intermediate vertex in that component would have degree at least 2, which is impossible. Therefore every connected component is either a singleton component of degree 0 or a two-vertex, one-edge component. $\square$

Proposition 4.1 implies that all remaining nonconvex coupling in the global problem is localized within one-unit blocks and two-unit edge blocks. We next specify the input coordinates assigned to each block and use this structure to decompose both the feasible set and the objective polynomial.

Let the set of isolated vertices be

$$\mathcal{I} := \{j \in \mathcal{J}_u \mid \deg_H(j) = 0\}.$$

For each input coordinate $k \in \{1, \dots, n_0\}$, define

$$\Lambda(k) := \{j \in \mathcal{J}_u \mid k \in S_j\}.$$

Under the matching assumption, one has $|\Lambda(k)| \leq 2$ for every k . Indeed, if $|\Lambda(k)| \geq 3$, then those three vertices would be pairwise adjacent and hence form a triangle, contradicting

the matching assumption. Therefore, each input coordinate belongs to exactly one of the following sets:

$$\begin{aligned} T_j &:= \{k \mid \Lambda(k) = \{j\}\} & (j \in \mathcal{J}_u), \\ T_e &:= \{k \mid \Lambda(k) = \{i, j\}\} & (e = \{i, j\} \in E), \end{aligned}$$

and

$$T_0 := \{1, \dots, n_0\} \setminus \bigcup_{j \in \mathcal{J}_u} S_j.$$

That is,

$$\{1, \dots, n_0\} = T_0 \sqcup \left(\bigcup_{j \in \mathcal{J}_u} T_j \right) \sqcup \left(\bigcup_{e \in E} T_e \right). \quad (4.1)$$

The input coordinates in T_0 do not appear in the nonlinear constraints of the undecided ReLUs. They appear only in the box constraints

$$\underline{u}_k \leq u_k \leq \bar{u}_k \quad (k \in T_0)$$

and in the linear objective term $-\sum_{k \in T_0} \delta_k u_k$. Thus u_{T_0} forms a free linear block that shares neither nonlinear equalities nor ReLU constraints with any other local block. Accordingly, when $T_0 \neq \emptyset$, if $\mathcal{J}_u = \emptyset$, the reduced problem is simply a linear optimization problem over the input box and can be certified by the box-affine argument in Lemma 5.7. Otherwise, we fix one connected component C° of the matching and absorb u_{T_0} together with the linear term $-\sum_{k \in T_0} \delta_k u_k$ into the local block of that component. Equivalently, the variables, box generators, and local affine objective of C° are augmented to include u_{T_0} , the bounds $\underline{u}_k \leq u_k \leq \bar{u}_k$ for $k \in T_0$, and the term $-\sum_{k \in T_0} \delta_k u_k$. Since this absorption creates no new variable sharing with the other components, the direct-product structure of the reduced problem is preserved.

4.2 Componentwise decomposition of the objective polynomial and local optimal values

For the objective polynomial of the reduced problem,

$$p_\gamma(u, v, x) := \gamma - c^\top (\bar{A}x + \bar{B}u + \bar{a}),$$

define

$$\eta_\gamma := \gamma - c^\top \bar{a}, \quad \alpha_j := (\bar{A}^\top c)_j \quad (j \in \mathcal{J}_u), \quad \delta_k := (\bar{B}^\top c)_k \quad (k = 1, \dots, n_0). \quad (4.2)$$

Then

$$p_\gamma = \eta_\gamma - \sum_{j \in \mathcal{J}_u} \alpha_j x_j - \sum_{k=1}^{n_0} \delta_k u_k.$$

For each isolated vertex $j \in \mathcal{I}$, define

$$\Theta_j^{\text{lin}}(\zeta_j) := -\alpha_j x_j - \sum_{k \in T_j} \delta_k u_k,$$

where $\zeta_j := (u_{T_j}, v_j, x_j)$. For each edge component $e = \{i, j\} \in E$, define

$$\Theta_e^{\text{lin}}(\zeta_e) := -\alpha_i x_i - \alpha_j x_j - \sum_{k \in T_i \cup T_e \cup T_j} \delta_k u_k, \quad (4.3)$$

where $\zeta_e := (u_{T_i \cup T_e \cup T_j}, v_i, x_i, v_j, x_j)$. If the base component C° absorbs T_0 , then the corresponding function $\Theta_{C^\circ}^{\text{lin}}$ and local block are enlarged to include $-\sum_{k \in T_0} \delta_k u_k$ and u_{T_0} . Equivalently, the local box generators of C° are augmented by $u_k - \underline{u}_k$ and $\bar{u}_k - u_k$ for $k \in T_0$. Under this convention, distinct connected components in the matching share no variables. Therefore,

$$p_\gamma = \eta_\gamma + \sum_{e \in E} \Theta_e^{\text{lin}}(\zeta_e) + \sum_{j \in \mathcal{I}} \Theta_j^{\text{lin}}(\zeta_j). \quad (4.4)$$

We next define the local feasible sets corresponding to each component. For an isolated vertex $j \in \mathcal{I}$, let

$$\begin{aligned} g_j(\zeta_j) &:= \{u_k - \underline{u}_k, \bar{u}_k - u_k \ (k \in T_j), x_j, x_j - v_j\}, \\ h_j(\zeta_j) &:= \{v_j - \bar{W}_{j,T_j} u_{T_j} - \bar{b}_j, x_j(x_j - v_j)\}, \end{aligned}$$

and for an edge component $e = \{i, j\} \in E$, let

$$\begin{aligned} g_e(\zeta_e) &:= \{u_k - \underline{u}_k, \bar{u}_k - u_k \ (k \in T_i \cup T_e \cup T_j), x_i, x_i - v_i, x_j, x_j - v_j\}, \\ h_e(\zeta_e) &:= \left\{ v_i - \bar{W}_{i,T_i} u_{T_i} - \bar{W}_{i,T_e} u_{T_e} - \bar{b}_i, \right. \\ &\quad \left. v_j - \bar{W}_{j,T_j} u_{T_j} - \bar{W}_{j,T_e} u_{T_e} - \bar{b}_j, \right. \\ &\quad \left. x_i(x_i - v_i), x_j(x_j - v_j) \right\}. \end{aligned}$$

We write the corresponding local feasible sets as

$$K_j := \{\zeta_j \mid g_j(\zeta_j) \geq 0, h_j(\zeta_j) = 0\} \quad (j \in \mathcal{I}),$$

$$K_e := \{\zeta_e \mid g_e(\zeta_e) \geq 0, h_e(\zeta_e) = 0\} \quad (e \in E).$$

Define the corresponding local optimal values by

$$m_j := \min\{\Theta_j^{\text{lin}}(\zeta_j) \mid \zeta_j \in K_j\} \quad (j \in \mathcal{I}), \quad (4.5)$$

$$m_e := \min\{\Theta_e^{\text{lin}}(\zeta_e) \mid \zeta_e \in K_e\} \quad (e \in E). \quad (4.6)$$

Proposition 4.2 (value decomposition under the matching assumption). *Assume that the input-sharing graph $H = (\mathcal{J}_u, E)$ is a matching. Then, the feasible set of the reduced problem decomposes as a direct product of the local feasible sets of the isolated-vertex and single-edge components. Consequently, for any $\gamma \in \mathbb{R}$ one has*

$$\min p_\gamma = \eta_\gamma + \sum_{e \in E} m_e + \sum_{j \in \mathcal{I}} m_j. \quad (4.7)$$

Proof. The key point is that, under the matching assumption, both the constraints and the objective decompose according to the same component partition. Indeed, by Proposition 4.1, each connected component is either an isolated vertex or a single edge, and by (4.1), each input coordinate is assigned to exactly one component. Moreover, as mentioned above, the variables in T_0 appear only in box constraints and linear objective terms. Absorbing them into the base component C° therefore creates no coupling with the other components. Therefore, the global feasible set of the reduced problem can be identified with

$$\prod_{e \in E} K_e \times \prod_{j \in \mathcal{I}} K_j.$$

On the other hand, (4.4) shows that the objective polynomial decomposes exactly into the constant term η_γ and the sum of the local objective functions. Hence, minimizing over a direct-product set separates into componentwise minimizations:

$$\min p_\gamma = \eta_\gamma + \sum_{e \in E} \min_{\zeta_e \in K_e} \Theta_e^{\text{lin}}(\zeta_e) + \sum_{j \in \mathcal{I}} \min_{\zeta_j \in K_j} \Theta_j^{\text{lin}}(\zeta_j).$$

By the definitions (4.6) and (4.5), these minima are exactly m_e and m_j , which yields (4.7). $\square$

4.3 Assembly from local tightness to global tightness

At this stage, the global assembly argument required for the paper consists of the following three components.

- (i) the feasible set decomposes into a direct product of local feasible sets,
- (ii) the objective polynomial decomposes into the sum of local affine objective functions, and
- (iii) if each local certificate belongs to the truncated quadratic module/ideal of degree 2, then any finite sum of them yields a global certificate at the same degree.

Thus, this section isolates the assembly details needed to combine the local results, so that the remaining mathematical task is to prove order-2 tightness for a single two-unit edge component.

Define

$$\mathcal{C}_{\text{sp}}^{(2)} := \left\{ \sum_{e \in E} Q_e(\zeta_e) + \sum_{j \in \mathcal{I}} P_j(\zeta_j) \mid Q_e \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e), P_j \in \mathcal{M}_2(g_j) + \mathcal{I}_2(h_j) \right\}. \quad (4.8)$$

This defines the order-2 certificate class for the block-sparse SOS relaxation induced by the connected-component decomposition of the matching. The corresponding relaxation value is denoted by

$$\gamma_{\text{sp}}^{(2)} := \inf \{ \gamma \in \mathbb{R} \mid p_\gamma \text{ admits a certificate belonging to } \mathcal{C}_{\text{sp}}^{(2)} \}. \quad (4.9)$$

Proposition 4.3 (global tightness criterion under the matching assumption). *Assume that the input-sharing graph $H = (\mathcal{J}_u, E)$ is a matching. Suppose further that for every isolated vertex $j \in \mathcal{I}$ and every edge $e \in E$,*

$$\Theta_j^{\text{lin}} - m_j \in \mathcal{M}_2(g_j) + \mathcal{I}_2(h_j) \quad (j \in \mathcal{I}), \quad (4.10)$$

and

$$\Theta_e^{\text{lin}} - m_e \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e) \quad (e \in E) \quad (4.11)$$

hold. Then, the order-2 block-sparse SOS relaxation is tight, namely,

$$\gamma_{\text{sp}}^{(2)} = \gamma^*. \quad (4.12)$$

Proof. Applying Proposition 4.2 with $\gamma = \gamma^*$, and using the fact that the global minimum of p_{γ^*} is 0, we obtain

$$\eta_{\gamma^*} + \sum_{e \in E} m_e + \sum_{j \in \mathcal{I}} m_j = 0. \quad (4.13)$$

On the other hand, summing the assumptions (4.10) and (4.11) gives

$$\begin{aligned} & \sum_{e \in E} (\Theta_e^{\text{lin}} - m_e) + \sum_{j \in \mathcal{I}} (\Theta_j^{\text{lin}} - m_j) \\ &= \left(\sum_{e \in E} \Theta_e^{\text{lin}} + \sum_{j \in \mathcal{I}} \Theta_j^{\text{lin}} \right) - \sum_{e \in E} m_e - \sum_{j \in \mathcal{I}} m_j \\ &= p_{\gamma^*} - \eta_{\gamma^*} - \sum_{e \in E} m_e - \sum_{j \in \mathcal{I}} m_j = p_{\gamma^*}, \end{aligned}$$

where the last equality follows from (4.13). Hence, $p_{\gamma^*} \in \mathcal{C}_{\text{sp}}^{(2)}$, and therefore $\gamma_{\text{sp}}^{(2)} \leq \gamma^*$.

For the reverse inequality, each summand in (4.8) is nonnegative on the corresponding local feasible set, and therefore their sum is also nonnegative on the global feasible set. Indeed, if $p_{\gamma} \in \mathcal{C}_{\text{sp}}^{(2)}$, then p_{γ} is nonnegative on the feasible set of the reduced problem, so γ is a valid upper bound for the original verification problem. Thus $\gamma \geq \gamma^*$, and by the definition (4.9) we obtain $\gamma^* \leq \gamma_{\text{sp}}^{(2)}$. This proves (4.12). $\square$

This completes the global assembly step. It remains to prove the edge-local certificate under the regular rank-one condition. Isolated vertices are handled by the one-unit results of Azuma et al. [2].

5 Local Exactness for Regular Rank-One Edges and the Main Theorem for Matching

In this section, we establish two-unit local tightness for the only nontrivial component that remains after the component decomposition of Section 4, namely a single-edge component. Our goal is to prove that for every edge-local affine objective ℓ_e expressible in the rank-one reduced variables,

$$\ell_e - \gamma_e^* \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e)$$

holds. We then combine this edge-local certificate with the global assembly result of Section 4 to obtain the main tightness theorem for matching input-sharing graphs.

The argument proceeds in four steps: reduce the shared part to a scalar, describe the one-unit and two-unit hulls using endpoint slices, convert this hull description into an order-2 certificate using the result of Azuma et al. [2] and Farkas' lemma, and finally combine the local result with the assembly criterion of Section 4.

Recall that **(M3a)** provides a one-dimensional shared scalar, while **(M3b)** ensures that both units remain undecided after that scalar is fixed. Fix an edge $e = \{i, j\} \in E$. We first describe the geometry in affinely reduced coordinates and then translate the resulting certificate back to the original edge block.

5.1 An edge-local normal form induced by regular rank-one edges

For each edge $e = \{i, j\}$, recall the shared input set introduced in Section 3,

$$S_e := S_i \cap S_j.$$

Condition **(M3a)** allows us to reduce this shared part to a single scalar.

By the assumption $q_e \leq 1$,

$$q_e = \dim \text{span}\{\bar{W}_{i,S_e}^\top, \bar{W}_{j,S_e}^\top\} \leq 1,$$

so there exist a nonzero vector $r_e \in \mathbb{R}^{S_e}$ and scalars α_i, α_j such that

$$\bar{W}_{i,S_e} = \alpha_i r_e^\top, \quad \bar{W}_{j,S_e} = \alpha_j r_e^\top.$$

Here r_e is a reference direction representing the common input direction through which the two units depend on the shared part.

For each $h \in \{i, j\}$, split the pre-activation into its shared and nonshared parts:

$$v_h = \bar{W}_{h,S_e} u_{S_e} + \bar{W}_{h,S_h \setminus S_e} u_{S_h \setminus S_e} + \bar{b}_h.$$

Introduce the scalar obtained from the shared part by

$$t_0 := r_e^\top u_{S_e}$$

and call the remaining term

$$p_h := \bar{W}_{h,S_h \setminus S_e} u_{S_h \setminus S_e} + \bar{b}_h \quad (h = i, j)$$

the private affine term of unit h . This quantity depends only on input coordinates outside the shared part. Since it is the affine image of a box, it ranges over some interval $[A_h, B_h]$. Hence, the pre-activation can be written as

$$v_h = p_h + \alpha_h t_0, \quad p_h \in [A_h, B_h] \quad (h = i, j). \quad (5.1)$$

Write the range of the shared scalar t_0 as

$$I_e := \{r_e^\top u_{S_e} \mid \underline{u}_{S_e} \leq u_{S_e} \leq \bar{u}_{S_e}\} = [\underline{\tau}_e, \bar{\tau}_e].$$

If $\underline{\tau}_e = \bar{\tau}_e$, then the shared scalar is fixed and the term $\alpha_h t_0$ in (5.1) can be absorbed into the private term as a constant. In that case the edge block degenerates into the direct product of two one-unit blocks, hence no genuine two-unit analysis is needed. Therefore, for the nontrivial case, we assume $\underline{\tau}_e < \bar{\tau}_e$ and normalize by

$$t := \frac{2t_0 - (\underline{\tau}_e + \bar{\tau}_e)}{\bar{\tau}_e - \underline{\tau}_e}.$$

To simplify notation, we write z_h for the pre-activation in these rescaled coordinates and obtain

$$z_h = p_h + \alpha_h t, \quad p_h \in [A_h, B_h] \quad (h = i, j). \quad (5.2)$$

Moreover, by normalization,

$$t \in [-1, 1]. \quad (5.3)$$

Thus the affine change of variables maps $t_0 = \underline{\tau}_e$ to $t = -1$ and $t_0 = \bar{\tau}_e$ to $t = 1$. In passing from (5.1) to (5.2), the constant part $\alpha_h(\underline{\tau}_e + \bar{\tau}_e)/2$ is absorbed into the private term and the coefficient of t is renamed α_h . The degenerate case is treated separately in Theorem 5.12.

Let

$$L_h(t) := A_h + \alpha_h t, \quad U_h(t) := B_h + \alpha_h t \quad (h = i, j). \quad (5.4)$$

Then, the statement that both units remain undecided at the endpoints $t = \pm 1$ can be written as

$$L_h(-1) < 0, \quad L_h(1) < 0, \quad U_h(-1) > 0, \quad U_h(1) > 0 \quad (h = i, j). \quad (5.5)$$

Lemma 5.1 (equivalence between edge regularity and endpoint conditions). *Assume that $e = \{i, j\}$ is a regular rank-one edge in the sense of Section 3. Then one of the following holds.*

- (a) *The range $I_e = [\underline{\tau}_e, \bar{\tau}_e]$ of the shared scalar is degenerate, that is, $\underline{\tau}_e = \bar{\tau}_e$, and the edge block degenerates into the direct product of two one-unit blocks.*
- (b) *One has $\underline{\tau}_e < \bar{\tau}_e$, and after choosing the normalization in (5.2)–(5.3) appropriately, condition (5.5) holds.*

Conversely, in the nondegenerate case, if (5.5) holds under some normalization, then the edge-regularity condition (3.3) holds.

Proof. If $I_e = [\underline{\tau}_e, \bar{\tau}_e]$ is degenerate with $\underline{\tau}_e = \bar{\tau}_e$, then the shared scalar is fixed, and by absorbing $\alpha_h t_0$ into the private term, the edge block degenerates into the direct product of two one-unit blocks. Thus (a) follows.

We now consider the nondegenerate case $\underline{\tau}_e < \bar{\tau}_e$. By (3.2), the shared part reduces to a one-dimensional scalar, and the pre-activations take the form (5.2). If we normalize so that the center of the affine image of the shared part is absorbed into the private term and only its signed length remains in $\alpha_h t$, then (3.3) is equivalent to

$$A_h + |\alpha_h| < 0 < B_h - |\alpha_h| \quad (h = i, j).$$

On the other hand, by (5.4),

$$\max\{L_h(-1), L_h(1)\} = A_h + |\alpha_h|, \quad \min\{U_h(-1), U_h(1)\} = B_h - |\alpha_h|,$$

which is equivalent to (5.5). Thus (b) follows. The converse implication is identical. $\square$

Therefore, in the nondegenerate case, we may replace the regular rank-one edge condition of Section 3 by the endpoint condition (5.5) via Lemma 5.1. Under these conditions, the convex hull of the slice of each unit at $t = \pm 1$ is the three-vertex polytope

$$\text{conv}\{(L_h(\pm 1), 0), (0, 0), (U_h(\pm 1), U_h(\pm 1))\}.$$

In degenerate boundary cases where equality occurs, one of the branches simply disappears, and the argument below becomes even simpler.

5.2 Description of the one-unit hull by endpoint slices

For each $h \in \{i, j\}$, define the reduced one-unit graph by

$$\mathcal{S}_h := \{(t, z_h, x_h) \mid t \in [-1, 1], L_h(t) \leq z_h \leq U_h(t), x_h = \text{ReLU}(z_h)\}. \quad (5.6)$$

Also write the convex hulls of the endpoint slices at $t = -1$ and $t = 1$ as

$$\begin{aligned} C_h^- &:= \text{conv}\{(L_h(-1), 0), (0, 0), (U_h(-1), U_h(-1))\}, \\ C_h^+ &:= \text{conv}\{(L_h(1), 0), (0, 0), (U_h(1), U_h(1))\}. \end{aligned}$$

Here C_h^- and C_h^+ are subsets of the (z_h, x_h) -plane, whose first coordinate is the pre-activation z_h and second coordinate is the ReLU output x_h . At the endpoints $t = \pm 1$, the one-unit graph consists of the negative branch $(z_h, 0)$ and the positive branch (z_h, z_h) ; thus, its convex hull is the convex hull of the three points $(L_h(\pm 1), 0)$, $(0, 0)$, and $(U_h(\pm 1), U_h(\pm 1))$. We use these endpoint slices to describe the exact hulls of both the one-unit and the two-unit problems.

Lemma 5.2 (endpoint decomposition of the one-unit hull). *For each $h \in \{i, j\}$,*

$$\text{conv}(\mathcal{S}_h) = \text{conv}(\{-1\} \times C_h^- \cup \{1\} \times C_h^+). \quad (5.7)$$

Proof. Write the right-hand side as H_h . First, from the definition it is clear that $\{-1\} \times C_h^- \subset \text{conv}(\mathcal{S}_h)$ and $\{1\} \times C_h^+ \subset \text{conv}(\mathcal{S}_h)$. Hence,

$$H_h \subset \text{conv}(\mathcal{S}_h).$$

Conversely, take any $(t, z_h, x_h) \in \mathcal{S}_h$ and set

$$\lambda := \frac{1-t}{2} \in [0, 1].$$

Then $t = \lambda(-1) + (1-\lambda)(1)$.

If $z_h \leq 0$, then $x_h = \text{ReLU}(z_h) = 0$. By the affine dependence in (5.4),

$$[L_h(t), 0] = \lambda[L_h(-1), 0] + (1-\lambda)[L_h(1), 0],$$

thus, there exist $z_h^- \in [L_h(-1), 0]$ and $z_h^+ \in [L_h(1), 0]$ such that

$$z_h = \lambda z_h^- + (1-\lambda)z_h^+.$$

Therefore,

$$(t, z_h, x_h) = \lambda(-1, z_h^-, 0) + (1-\lambda)(1, z_h^+, 0) \in H_h.$$

If instead $z_h \geq 0$, then $x_h = z_h$. Similarly,

$$[0, U_h(t)] = \lambda[0, U_h(-1)] + (1-\lambda)[0, U_h(1)],$$

hence, there exist $z_h^- \in [0, U_h(-1)]$ and $z_h^+ \in [0, U_h(1)]$ such that

$$z_h = \lambda z_h^- + (1-\lambda)z_h^+.$$

Hence,

$$(t, z_h, x_h) = \lambda(-1, z_h^-, z_h^-) + (1-\lambda)(1, z_h^+, z_h^+) \in H_h.$$

Thus $\mathcal{S}_h \subset H_h$, and since the right-hand side is convex, we obtain

$$\text{conv}(\mathcal{S}_h) \subset H_h.$$

The proof of (5.7) is complete. $\square$

By Lemma 5.2, the one-unit convex hull can be described without treating the entire interval directly: it is the convex combination of the two three-vertex convex hulls at the endpoints $t = \pm 1$. We now lift this fact to the two-unit edge block.

5.3 Exact decomposition of the two-unit edge hull

Define the reduced two-unit edge graph by

$$\mathcal{S}_e := \{(t, z_i, z_j, x_i, x_j) \mid (t, z_i, x_i) \in \mathcal{S}_i, (t, z_j, x_j) \in \mathcal{S}_j\}.$$

Using the convex hull of each unit, define also

$$G_e := \{(t, z_i, z_j, x_i, x_j) \mid (t, z_i, x_i) \in \text{conv}(\mathcal{S}_i), (t, z_j, x_j) \in \text{conv}(\mathcal{S}_j)\}.$$

Theorem 5.3 (exact decomposition of the two-unit edge hull). *Under $q_e \leq 1$ and (5.5),*

$$\text{conv}(\mathcal{S}_e) = G_e = \text{conv}(\{-1\} \times C_i^- \times C_j^- \cup \{1\} \times C_i^+ \times C_j^+). \quad (5.8)$$

Proof. First, $\mathcal{S}_e \subset G_e$ and G_e is convex, so

$$\text{conv}(\mathcal{S}_e) \subset G_e.$$

Conversely, take any

$$y = (t, z_i, z_j, x_i, x_j) \in G_e.$$

By Lemma 5.2, with $\lambda = (1 - t)/2 \in [0, 1]$, there exist $a_i^- \in C_i^-$, $a_i^+ \in C_i^+$, $a_j^- \in C_j^-$, and $a_j^+ \in C_j^+$ such that

$$\begin{aligned} (t, z_i, x_i) &= \lambda(-1, a_i^-) + (1 - \lambda)(1, a_i^+), \\ (t, z_j, x_j) &= \lambda(-1, a_j^-) + (1 - \lambda)(1, a_j^+). \end{aligned}$$

Hence,

$$y = \lambda(-1, a_i^-, a_j^-) + (1 - \lambda)(1, a_i^+, a_j^+),$$

and therefore,

$$G_e \subset \text{conv}(\{-1\} \times C_i^- \times C_j^- \cup \{1\} \times C_i^+ \times C_j^+).$$

On the other hand, both endpoint product slices on the right-hand side are contained in G_e . By convexity of G_e , the reverse inclusion also holds. Thus

$$G_e = \text{conv}(\{-1\} \times C_i^- \times C_j^- \cup \{1\} \times C_i^+ \times C_j^+).$$

Finally, when $t = \pm 1$ is fixed, the degrees of freedom on the i -side and the j -side are independent. Therefore,

$$\begin{aligned} \{-1\} \times C_i^- \times C_j^- &= \text{conv}(\mathcal{S}_e \cap \{t = -1\}), \\ \{1\} \times C_i^+ \times C_j^+ &= \text{conv}(\mathcal{S}_e \cap \{t = 1\}), \end{aligned}$$

and hence the right-hand side is contained in $\text{conv}(\mathcal{S}_e)$. Together with the first inclusion, it proves (5.8). $\square$

Theorem 5.3 shows that the convex hull of the entire edge block is described exactly as a convex combination of the two endpoint product slices at the two ends of the shared scalar. Consequently, the required two-unit analysis is reduced not to a direct description of mixed branches, but to an exact coupling of one-unit interval hulls at the endpoints. This representation immediately yields a linear extended formulation.

Corollary 5.4 (linear extended formulation). *The exact hull of Theorem 5.3 can be represented linearly by using nonnegative variables λ_- , λ_+ and scaled endpoint variables*

$$(\hat{z}_h^-, \hat{x}_h^-), (\hat{z}_h^+, \hat{x}_h^+) \quad (h = i, j)$$

as follows:

$$\lambda_- + \lambda_+ = 1, \quad t = -\lambda_- + \lambda_+, \quad (5.9)$$

$$z_h = \hat{z}_h^- + \hat{z}_h^+, \quad x_h = \hat{x}_h^- + \hat{x}_h^+ \quad (h = i, j), \quad (5.10)$$

$$(\hat{z}_h^-, \hat{x}_h^-) \in \lambda_- C_h^-, \quad (\hat{z}_h^+, \hat{x}_h^+) \in \lambda_+ C_h^+ \quad (h = i, j). \quad (5.11)$$

This is the two-term specialization of the convex-hull formulation for a finite union of polyhedra in disjunctive programming. Here, $\lambda C_h^\pm$ denotes the usual perspective scaling of the polytope $C_h^\pm$, which is described by linear inequalities [3].

Proof. By Theorem 5.3, every point in the exact hull can be written as a convex combination of one point from the endpoint slice at $t = -1$ and one point from the endpoint slice at $t = 1$. Thus, there exist $\lambda_-, \lambda_+ \geq 0$ with $\lambda_- + \lambda_+ = 1$, and points

$$(z_h^-, x_h^-) \in C_h^-, \quad (z_h^+, x_h^+) \in C_h^+ \quad (h = i, j),$$

such that

$$(t, z_i, z_j, x_i, x_j) = \lambda_-(-1, z_i^-, z_j^-, x_i^-, x_j^-) + \lambda_+(1, z_i^+, z_j^+, x_i^+, x_j^+).$$

Equating coordinates gives $t = -\lambda_- + \lambda_+$ and, for $h = i, j$,

$$z_h = \lambda_- z_h^- + \lambda_+ z_h^+, \quad x_h = \lambda_- x_h^- + \lambda_+ x_h^+.$$

Introduce the scaled endpoint variables

$$(\hat{z}_h^-, \hat{x}_h^-) := \lambda_- (z_h^-, x_h^-), \quad (\hat{z}_h^+, \hat{x}_h^+) := \lambda_+ (z_h^+, x_h^+).$$

Then the preceding coordinate equations become (5.9)–(5.10), while the endpoint membership conditions become (5.11).

Conversely, suppose that variables satisfying (5.9)–(5.11) are given. If $\lambda_\sigma > 0$ for $\sigma \in \{-, +\}$, then

$$(z_h^\sigma, x_h^\sigma) := \lambda_\sigma^{-1} (\hat{z}_h^\sigma, \hat{x}_h^\sigma) \in C_h^\sigma.$$

If $\lambda_\sigma = 0$, the condition $(\hat{z}_h^\sigma, \hat{x}_h^\sigma) \in \lambda_\sigma C_h^\sigma$ means $(\hat{z}_h^\sigma, \hat{x}_h^\sigma) = (0, 0)$, and the corresponding endpoint term has zero weight. Therefore, any point of C_h^σ may be chosen for (z_h^σ, x_h^σ) . Using the equalities in the formulation, the resulting point (t, z_i, z_j, x_i, x_j) belongs to

$$\text{conv}(\{-1\} \times C_i^- \times C_j^- \cup \{1\} \times C_i^+ \times C_j^+),$$

which is $\text{conv}(\mathcal{S}_e)$ by Theorem 5.3. Since each C_h^σ is a polytope, the scaled membership condition $(\hat{z}_h^\sigma, \hat{x}_h^\sigma) \in \lambda_\sigma C_h^\sigma$ is described by linear inequalities. Hence, the formulation is linear and exact. $\square$

5.4 Reduction to an order-2 local certificate for the SOS relaxation

We have now obtained a geometric description of the two-unit edge hull in reduced coordinates. It remains to convert the polyhedral description into an order-2 local certificate. In this subsection, we show that the conversion can be performed at relaxation order 2. The proof has three steps: (i) applying the one-unit result of Azuma et al. [2] to the reduced one-unit problem, (ii) using a Farkas representation for affine optimization over the reduced edge hull, and (iii) lifting the resulting certificate back to the original edge block by affine elimination.

For each $h \in \{i, j\}$, define the reduced one-unit set, with the private scalar included explicitly, by

$$\hat{\mathcal{S}}_h := \{(t, p_h, z_h, x_h) \mid t \in [-1, 1], p_h \in [A_h, B_h], z_h = p_h + \alpha_h t, x_h = \text{ReLU}(z_h)\}. \quad (5.12)$$

Its projection onto (t, z_h, x_h) is exactly $\mathcal{S}_h$ in (5.6). For $\sigma \in \{-, +\}$, write

$$L_h^\sigma := L_h(\sigma 1), \quad U_h^\sigma := U_h(\sigma 1).$$

In the nondegenerate case of Lemma 5.1, one has

$$L_h^\sigma < 0 < U_h^\sigma \quad (h = i, j, \sigma \in \{-, +\}).$$

Lemma 5.5 (explicit linear description of the endpoint triangles). *For each $h \in \{i, j\}$ and each $\sigma \in \{-, +\}$, the endpoint triangle C_h^σ can be written as*

$$C_h^\sigma = \left\{ (z_h, x_h) \in \mathbb{R}^2 \mid x_h \geq 0, x_h - z_h \geq 0, x_h \leq \frac{U_h^\sigma}{U_h^\sigma - L_h^\sigma}(z_h - L_h^\sigma) \right\}. \quad (5.13)$$

Proof. Since $C_h^\sigma = \text{conv}\{(L_h^\sigma, 0), (0, 0), (U_h^\sigma, U_h^\sigma)\}$ is a triangle, it suffices to write the half-spaces corresponding to its three edges. The line through $(L_h^\sigma, 0)$ and $(0, 0)$ results in $x_h = 0$, and the line through $(0, 0)$ and (U_h^σ, U_h^σ) leads to $x_h = z_h$. The remaining edge through $(L_h^\sigma, 0)$ and (U_h^σ, U_h^σ) has equation

$$x_h = \frac{U_h^\sigma}{U_h^\sigma - L_h^\sigma}(z_h - L_h^\sigma).$$

Since the interior of the triangle is the intersection of these half-spaces, we obtain (5.13). $\square$

Thus, at each endpoint $t = \pm 1$, the last inequality in (5.13) is the nontrivial facet of the interval ReLU hull, as used in convex ReLU formulations [1, 6]. We use affine inequalities of this type on the one-unit side to represent the optimality gap over the reduced two-unit edge hull.

Lemma 5.6 (the reduced one-unit problem falls within the scope of the one-unit result of Azuma et al. [2]). *For each $h \in \{i, j\}$, the set $\hat{\mathcal{S}}_h$ in (5.12) corresponds to a single-neuron ReLU problem with input variables (t, p_h) ranging over the rectangle*

$$[-1, 1] \times [A_h, B_h].$$

Therefore, the one-unit tightness result of Azuma et al. [2] described in Section 2 applies. In particular, for any affine inequality

$$\varphi_h(t, z_h, x_h) \geq 0$$

valid on $\text{conv}(\mathcal{S}_h)$, we have

$$\varphi_h(t, z_h, x_h) \in \mathcal{M}_2(g_h^{\text{red}}) + \mathcal{I}_2(h_h^{\text{red}}), \quad (5.14)$$

where $(g_h^{\text{red}}, h_h^{\text{red}})$ is the family of box constraints and linear equalities that define the reduced one-unit block. Equivalently, the one-unit result imported in this section provides order-2 certificates without requiring higher-order relaxations.

Proof. The set $\hat{\mathcal{S}}_h$ is a one-unit problem with input (t, p_h) , pre-activation z_h , and output x_h , and the input set is the rectangle $[-1, 1] \times [A_h, B_h]$. Thus, it is a special case of the rectangular-input one-unit setting of Azuma et al. [2].

Assume that $\varphi_h(t, z_h, x_h) \geq 0$ is valid on $\text{conv}(\mathcal{S}_h)$. Its pullback to $\hat{\mathcal{S}}_h$ is the affine function

$$\hat{\varphi}_h(t, p_h, z_h, x_h) := \varphi_h(t, z_h, x_h),$$

which is valid on $\hat{\mathcal{S}}_h$. Let

$$\hat{m}_h := \min\{\hat{\varphi}_h(t, p_h, z_h, x_h) \mid (t, p_h, z_h, x_h) \in \hat{\mathcal{S}}_h\}.$$

Then $\hat{m}_h \geq 0$. By the one-unit tightness result of Azuma et al. [2] applied to the affine objective $\hat{\varphi}_h$, we have

$$\hat{\varphi}_h - \hat{m}_h \in \mathcal{M}_2(\hat{g}_h) + \mathcal{I}_2(\hat{h}_h).$$

Since $\widehat{m}_h$ is a nonnegative constant, it belongs to the quadratic module as an SOS polynomial. Therefore,

$$\widehat{\varphi}_h \in \mathcal{M}_2(\widehat{g}_h) + \mathcal{I}_2(\widehat{h}_h).$$

Finally, applying Lemma 5.8(i) to the affine substitution $p_h = z_h - \alpha_h t$ transfers this certificate to (5.14) for the reduced block $(g_h^{\text{red}}, h_h^{\text{red}})$. $\square$

Pullbacks also play a role in positivity certificates for polynomial optimization [12]. Here, however, the required maps are affine, and therefore preserve SOS structure without increasing the degree.

Lemma 5.7 (affine inequalities valid on a box have order-2 box certificates). *Let*

$$B = [\underline{u}_1, \bar{u}_1] \times \cdots \times [\underline{u}_m, \bar{u}_m]$$

and let

$$g_B := \{u_k - \underline{u}_k, \bar{u}_k - u_k \mid k = 1, \dots, m\}$$

be its box generators. If an affine polynomial

$$q(u) = q_0 + \sum_{k=1}^m a_k u_k$$

is nonnegative on B , then

$$q \in \mathcal{M}_2(g_B).$$

Proof. The minimum of q on the box is attained at the vertex

$$u_k^* = \begin{cases} \underline{u}_k, & a_k \geq 0, \\ \bar{u}_k, & a_k < 0. \end{cases}$$

Set $m_q := q(u^*) \geq 0$. Then

$$q(u) = m_q + \sum_{a_k \geq 0} a_k (u_k - \underline{u}_k) + \sum_{a_k < 0} (-a_k) (\bar{u}_k - u_k).$$

Since m_q is a nonnegative constant, it is an SOS polynomial. Moreover, all coefficients multiplying the box generators are nonnegative constants. Hence, this representation belongs to $\mathcal{M}_2(g_B)$. $\square$

Lemma 5.8 (affine pullback preserves order-2 certificates). *The following statements hold.*

(i) *For each $h \in \{i, j\}$, consider the affine substitution*

$$\Phi_h : \mathbb{R}[t, p_h, z_h, x_h] \rightarrow \mathbb{R}[t, z_h, x_h], \quad (\Phi_h q)(t, z_h, x_h) := q(t, z_h - \alpha_h t, z_h, x_h).$$

If $\widehat{\psi}(t, p_h, z_h, x_h) \in \mathcal{M}_2(\widehat{g}_h) + \mathcal{I}_2(\widehat{h}_h)$, then

$$\psi := \Phi_h(\widehat{\psi}) \in \mathcal{M}_2(g_h^{\text{red}}) + \mathcal{I}_2(h_h^{\text{red}})$$

holds.

(ii) Let π_e be the affine map from the original edge-local variables ζ_e to the reduced variables $\xi = (t, z_i, x_i, z_j, x_j)$, and let its pullback be

$$\Pi_e : \mathbb{R}[\xi] \rightarrow \mathbb{R}[\zeta_e], \quad (\Pi_e q)(\zeta_e) := q(\pi_e(\zeta_e)).$$

If a reduced polynomial $\tilde{\psi}(\xi)$ satisfies

$$\tilde{\psi} \in \mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}}),$$

then its pullback satisfies

$$\Pi_e(\tilde{\psi}) = \tilde{\psi}(\pi_e(\zeta_e)) \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e).$$

Moreover, in both (i) and (ii), the polynomial degree is preserved under the pullback. Hence, the degree conditions associated with relaxation order 2 remain valid.

Proof. (i) An affine substitution is a ring homomorphism. Hence if $\hat{\sigma} = \sum_m r_m^2$ is SOS, then

$$\Phi_h(\hat{\sigma}) = \sum_m (r_m(t, z_h - \alpha_h t, z_h, x_h))^2$$

is again SOS. Since the map $p_h \mapsto z_h - \alpha_h t$ is affine, it does not increase degree. Hence,

$$\deg \Phi_h(q) \leq \deg q$$

for every polynomial q .

Now, write

$$\hat{\psi} = \hat{\sigma}_0 + \sum_{\ell} \hat{\sigma}_{\ell} \hat{g}_{h,\ell} + \sum_r \hat{\tau}_r \hat{h}_{h,r} \in \mathcal{M}_2(\hat{g}_h) + \mathcal{I}_2(\hat{h}_h).$$

Each $\hat{\sigma}_{\ell}$ is SOS, and each term satisfies the degree constraints associated with relaxation order 2. After substituting $p_h = z_h - \alpha_h t$, we obtain

$$\Phi_h(p_h - A_h) = z_h - L_h(t), \quad \Phi_h(B_h - p_h) = U_h(t) - z_h.$$

Thus, the box constraints associated with $t \in [-1, 1]$ are preserved. Moreover, the linear equality

$$z_h - p_h - \alpha_h t = 0$$

becomes identically zero after pullback. Therefore, $\Phi_h(\hat{\psi})$ belongs to the sum of the truncated quadratic module and ideal associated with the reduced one-unit block $(g_h^{\text{red}}, h_h^{\text{red}})$. Since no term increases in degree, we obtain

$$\Phi_h(\hat{\psi}) \in \mathcal{M}_2(g_h^{\text{red}}) + \mathcal{I}_2(h_h^{\text{red}}).$$

(ii) Similarly, Π_e is the ring homomorphism induced by the affine map π_e , and therefore sends SOS polynomials to SOS polynomials without increasing degree. It remains only to track the reduced generators under the pullback. The ReLU generators x_h and $x_h - z_h$ become the original generators x_h and $x_h - v_h$, respectively, while the reduced complementarity equality $x_h(x_h - z_h) = 0$ becomes the original equality $x_h(x_h - v_h) = 0$.

The remaining reduced generators are affine box-type bounds. The pullbacks of $1+t$ and $1-t$ are affine functions that are nonnegative on the shared input box by the normalization of t . Similarly, the pullbacks of $z_h - L_h(t)$ and $U_h(t) - z_h$ are $p_h - A_h$ and $B_h - p_h$, respectively,

and these affine functions are nonnegative on the private input box by the definitions of A_h and B_h . By Lemma 5.7, each of these affine pullbacks belongs to the order-2 quadratic module generated by the corresponding original box constraints. Therefore, every reduced inequality generator pulls back either to an original ReLU generator or to an element of the original box quadratic module, while every reduced equality generator pulls back to the original ideal.

Thus if

$$\tilde{\psi} \in \mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}}),$$

then applying Π_e term by term maps it into $\mathcal{M}_2(g_e) + \mathcal{I}_2(h_e)$. The degree bounds are preserved since the pullback is affine and Lemma 5.7 represents each pulled-back affine box-type generator as a nonnegative constant combination of affine original box generators. Consequently,

$$\Pi_e(\tilde{\psi}) \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e).$$

□

Lemma 5.9 (adding unused variables preserves order-2 certificates). *For each $h \in \{i, j\}$, suppose*

$$\psi_h(t, z_h, x_h) \in \mathcal{M}_2(g_h^{\text{red}}) + \mathcal{I}_2(h_h^{\text{red}}).$$

Then, viewing ψ_h as a polynomial in the reduced edge variables (t, z_i, x_i, z_j, x_j) , we also have

$$\psi_h \in \mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}}).$$

Proof. Each generator of the reduced one-unit block appears verbatim as part of the generators of the reduced edge block. Therefore, any representation

$$\psi_h = \sigma_0 + \sum_{\ell} \sigma_{\ell} g_{h,\ell}^{\text{red}} + \sum_r \tau_r h_{h,r}^{\text{red}}$$

in $\mathcal{M}_2(g_h^{\text{red}}) + \mathcal{I}_2(h_h^{\text{red}})$ remains a valid representation in $\mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}})$, after viewing the coefficient polynomials as independent of the unused variables $(z_{h'}, x_{h'})$. The SOS structure and the degree bounds are unchanged, which proves the claim. □

Lemma 5.10 (affine optimization over the reduced edge graph reduces to linear programming). *Assume the nondegenerate case $\underline{\tau}_e < \bar{\tau}_e$, and let the reduced edge hull be*

$$P_e := \text{conv}(\mathcal{S}_e).$$

Then, for every reduced affine objective $\tilde{\ell}_e$,

$$\min\{\tilde{\ell}_e(\xi) \mid \xi \in \mathcal{S}_e\} = \min\{\tilde{\ell}_e(\xi) \mid \xi \in P_e\} \quad (5.15)$$

holds. Moreover, using the linear extended formulation of Corollary 5.4, the right-hand side can be solved as a linear program in lifted variables.

Proof. Since $\tilde{\ell}_e$ is affine, its value at a convex combination of points in $\mathcal{S}_e$ is the same convex combination of the corresponding values. Hence, no point in $\text{conv}(\mathcal{S}_e)$ can have a smaller value than the minimum over $\mathcal{S}_e$, and the reverse inequality is immediate from $\mathcal{S}_e \subset \text{conv}(\mathcal{S}_e)$. Therefore,

$$\min\{\tilde{\ell}_e(\xi) \mid \xi \in \mathcal{S}_e\} = \min\{\tilde{\ell}_e(\xi) \mid \xi \in \text{conv}(\mathcal{S}_e)\}.$$

Since $P_e = \text{conv}(\mathcal{S}_e)$, this proves (5.15). The second claim follows from Corollary 5.4, which provides a linear extended formulation of P_e . □

Lemma 5.11 (Farkas representation of the optimality gap on the reduced edge hull). *Assume the nondegenerate case $\underline{\tau}_e < \bar{\tau}_e$. For each $h \in \{i, j\}$, let*

$$P_h := \text{conv}(\mathcal{S}_h).$$

Then, for any reduced affine objective $\tilde{\ell}_e$, if its reduced optimal value is

$$\tilde{\gamma}_e^* := \min\{\tilde{\ell}_e(\xi) \mid \xi \in P_e\},$$

there exist finitely many affine inequalities $\varphi_{i,1}, \dots, \varphi_{i,N_i}$ valid on P_i , affine inequalities $\varphi_{j,1}, \dots, \varphi_{j,N_j}$ valid on P_j , and nonnegative coefficients $\lambda_{i,1}, \dots, \lambda_{i,N_i}, \lambda_{j,1}, \dots, \lambda_{j,N_j}$ such that

$$\tilde{\ell}_e - \tilde{\gamma}_e^* = \sum_{r=1}^{N_i} \lambda_{i,r} \varphi_{i,r}(t, z_i, x_i) + \sum_{s=1}^{N_j} \lambda_{j,s} \varphi_{j,s}(t, z_j, x_j) \quad (5.16)$$

holds on the reduced edge variables.

Proof. By Theorem 5.3, the reduced edge hull can be written as

$$P_e = \{(t, z_i, z_j, x_i, x_j) \mid (t, z_i, x_i) \in P_i, (t, z_j, x_j) \in P_j\}.$$

Moreover, by Lemma 5.2 and Corollary 5.4, each P_h is a polytope. Hence, we can write

$$P_h = \{(t, z_h, x_h) \mid \varphi_{h,r}(t, z_h, x_h) \geq 0 \text{ for } r = 1, \dots, N_h\}$$

using finitely many valid affine inequalities.

Let $M_i(t, z_i, x_i, z_j, x_j) = (t, z_i, x_i)$ and $M_j(t, z_i, x_i, z_j, x_j) = (t, z_j, x_j)$. Then P_e admits the polyhedral description

$$P_e = \{\xi \mid \varphi_{i,r}(M_i\xi) \geq 0 \ (r = 1, \dots, N_i), \varphi_{j,s}(M_j\xi) \geq 0 \ (s = 1, \dots, N_j)\}.$$

Each row is simply an affine inequality valid on P_i or P_j , embedded into the reduced edge variables.

Now, $\tilde{\gamma}_e^*$ is the minimum of the affine function $\tilde{\ell}_e$ over the polytope P_e . Therefore, by the polyhedral form of Farkas' lemma, there exist nonnegative vectors λ_i, λ_j such that

$$\tilde{\ell}_e - \tilde{\gamma}_e^* = \sum_{r=1}^{N_i} \lambda_{i,r} \varphi_{i,r}(M_i\xi) + \sum_{s=1}^{N_j} \lambda_{j,s} \varphi_{j,s}(M_j\xi).$$

Renaming the rows as $\varphi_{i,r}$ and $\varphi_{j,s}$ yields (5.16). $\square$

Before applying the reduced edge hull description to the original edge block, we specify the class of affine objectives covered by the reduced-hull argument. Let ℓ_e be an affine edge-local objective whose dependence on the original input variables is given by

$$d_i^\top u_{T_i} + d_e^\top u_{T_e} + d_j^\top u_{T_j},$$

where T_i , T_e , and T_j denote, respectively, the private input set of unit i , the shared input set, and the private input set of unit j . We say that ℓ_e is *expressible in the rank-one reduced variables* if

$$d_i \in \text{span}\{\bar{W}_{i,T_i}^\top\}, \quad d_j \in \text{span}\{\bar{W}_{j,T_j}^\top\}, \quad d_e \in \text{span}\{r_e\}. \quad (5.17)$$

We call (5.17) the reduced-objective condition. Under this condition, the input-variable dependence of ℓ_e is affine in the reduced scalars p_i , p_j , and $t_0 = r_e^\top u_{T_e}$. After normalizing t_0 and using the substitution $p_h = z_h - \alpha_h t$, the objective restricted to the feasible edge block can be expressed as a reduced affine objective,

$$\ell_e(\zeta_e) = \tilde{\ell}_e(t, z_i, x_i, z_j, x_j).$$

Theorem 5.12 (edge-local tightness for a regular rank-one edge at relaxation order 2). *Assume that the edge $e = \{i, j\}$ is a regular rank-one edge in the sense of Section 3. Let ℓ_e be an affine objective on the edge block satisfying the reduced-objective condition (5.17), and let the local optimal value be denoted by*

$$\gamma_e^* := \min\{\ell_e(\zeta_e) \mid \zeta_e \in K_e\}.$$

Then,

$$\ell_e - \gamma_e^* \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e) \quad (5.18)$$

holds.

Proof. We first consider the degenerate case $\underline{\tau}_e = \bar{\tau}_e$. By Lemma 5.1, the edge block then reduces to the direct product of two one-unit blocks. Applying the one-unit tightness result of Azuma et al. [2] to each block and summing the resulting certificates yields an order-2 certificate for this direct-product block. Pulling this certificate back through the affine elimination in Lemma 5.8(ii), we obtain

$$\ell_e - \gamma_e^* \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e).$$

We now turn to the nondegenerate case $\underline{\tau}_e < \bar{\tau}_e$.

Using the affine elimination introduced at the beginning of this section, let π_e map the original edge variables ζ_e to

$$\xi := (t, z_i, x_i, z_j, x_j).$$

By the reduced-objective condition, there exists a reduced affine objective $\tilde{\ell}_e$ such that, on the feasible edge block,

$$\ell_e(\zeta_e) = \tilde{\ell}_e(\pi_e(\zeta_e)).$$

Set

$$\tilde{\gamma}_e^* := \min\{\tilde{\ell}_e(\xi) \mid \xi \in P_e\}.$$

Moreover, the image of the original feasible edge block under π_e is the reduced graph $\mathcal{S}_e$, and $P_e = \text{conv}(\mathcal{S}_e)$. Since $\tilde{\ell}_e$ is affine, its minimum over $\mathcal{S}_e$ coincides with its minimum over P_e . Hence,

$$\tilde{\gamma}_e^* = \gamma_e^*.$$

Next, apply Lemma 5.11 to $\tilde{\ell}_e$. Then, there exist affine inequalities $\varphi_{i,r}$ valid on P_i , affine inequalities $\varphi_{j,s}$ valid on P_j , and nonnegative coefficients $\lambda_{i,r}, \lambda_{j,s}$ such that

$$\tilde{\ell}_e - \tilde{\gamma}_e^* = \sum_r \lambda_{i,r} \varphi_{i,r}(t, z_i, x_i) + \sum_s \lambda_{j,s} \varphi_{j,s}(t, z_j, x_j).$$

Each $\varphi_{i,r}$ and $\varphi_{j,s}$ is valid on the corresponding reduced one-unit hull. Thus, by Lemma 5.6, we have

$$\varphi_{i,r} \in \mathcal{M}_2(g_i^{\text{red}}) + \mathcal{I}_2(h_i^{\text{red}}), \quad \varphi_{j,s} \in \mathcal{M}_2(g_j^{\text{red}}) + \mathcal{I}_2(h_j^{\text{red}}).$$

Applying Lemma 5.9 to each term, we may regard these terms as elements of

$$\mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}}).$$

Since the coefficients are nonnegative, their sum belongs to the same truncated quadratic module plus ideal. Therefore,

$$\tilde{\ell}_e - \tilde{\gamma}_e^* \in \mathcal{M}_2(g_e^{\text{red}}) + \mathcal{I}_2(h_e^{\text{red}}).$$

Finally, applying Lemma 5.8(ii) to $\tilde{\psi} = \tilde{\ell}_e - \tilde{\gamma}_e^*$ gives a pullback certificate satisfying

$$\ell_e - \gamma_e^* \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e),$$

which is exactly (5.18). $\square$

Thus, a regular rank-one edge is handled by transferring the one-unit result of Azuma et al. [2] through the reduced hull and the affine pullback.

5.5 Main theorem for matching

We use the notation of Section 4 and write

$$\mathcal{I} = \{j \in \mathcal{J}_u \mid \deg_H(j) = 0\}$$

for the set of isolated vertices. The final global theorem is proved under the following three structural assumptions.

(M1) Each isolated vertex block (g_j, h_j) , $j \in \mathcal{I}$, satisfies at least one of the one-unit conditions of Azuma et al. [2].

(M2) The input-sharing graph $H = (\mathcal{J}_u, E)$ is a matching.

(M3) Every edge $e \in E$ is a regular rank-one edge as defined in Section 3.

Theorem 5.13 (main theorem for matching). *Assume that **(M1)**–**(M3)** hold. Suppose further that, for the given output objective, every edge-local affine objective Θ_e^{lin} defined in (4.3) satisfies the reduced-objective condition (5.17). Then, the order-2 block-sparse SOS relaxation induced by the connected components of the matching is tight. In particular,*

$$\gamma_{\text{sp}}^{(2)} = \gamma^*.$$

Proof. First, let $j \in \mathcal{I}$ be any isolated vertex. By assumption **(M1)**, the one-unit tightness result of Azuma et al. [2] from Section 2 applies to the local problem (g_j, h_j) . Therefore, for the local affine objective Θ_j^{lin} defined in Section 4, we have

$$\Theta_j^{\text{lin}} - m_j \in \mathcal{M}_2(g_j) + \mathcal{I}_2(h_j),$$

which is exactly (4.10).

Next, for any edge $e \in E$, assumption **(M3)** implies that e is a regular rank-one edge. The reduced-objective condition in the theorem statement ensures that Theorem 5.12 applies to the local affine objective $\ell_e = \Theta_e^{\text{lin}}$. Therefore,

$$\Theta_e^{\text{lin}} - m_e \in \mathcal{M}_2(g_e) + \mathcal{I}_2(h_e),$$

which is (4.11).

Finally, by assumption **(M2)**, the graph H is a matching. Hence Proposition 4.3 from Section 4 applies to the component decomposition. Summing the local certificates established above for the isolated vertices and single-edge components yields a global certificate. Therefore

$$\gamma_{\text{sp}}^{(2)} = \gamma^*.$$

$\square$

Remark 5.14. The reduced-objective condition (5.17) is objective-dependent and is not an additional graph assumption. With **(M1)**–**(M3)**, it is directly checkable from the reduced data and the input box. Extending Theorem 5.13 to objectives outside this class would require retaining the input directions eliminated by the rank-one reduction.

6 Conclusion and Future Work

We have established tightness at relaxation order 2 for a block-sparse SOS relaxation of one-layer ReLU verification under the matching, regular rank-one, and reduced-objective conditions. The matching structure decomposes the problem into isolated-vertex and single-edge blocks. Existing one-unit results handle the former, while the endpoint-slice description developed here yields local certificates for the latter. Summing these certificates provides the global tightness result.

Extending the analysis to general forests remains open since vertices of degree at least 2 introduce compatibility conditions that must be represented by low-degree sparse SOS certificates. Other directions include edges with shared rank $q_e \geq 2$, graphs with cycles, and higher-degree certificates for structures beyond rank-one regularity.

Acknowledgements

We used AI-based language tools to improve the wording, readability, and presentation of some parts of the manuscript. We also used such tools at an early stage of writing to help generate possible examples. All mathematical statements, proofs, formulations, and numerical results were checked independently by the authors, without relying on AI-generated claims. The authors take full responsibility for the accuracy, validity, and integrity of the contents of this paper.

References

- [1] R. Anderson, J. Huchette, W. Ma, C. Tjandraatmadja, and J. P. Vielma. Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming*, 183:3–39, 2020. doi: 10.1007/s10107-020-01474-5.
- [2] G. Azuma, S. Kim, and M. Yamashita. Verifying robustness of neural networks with tight semidefinite relaxations. *Computational Optimization and Applications*, 94(1): 41–72, 2026. doi: 10.1007/s10589-026-00765-5.
- [3] E. Balas. Disjunctive programming. In *Discrete Optimization II*, volume 5 of *Annals of Discrete Mathematics*, pages 3–51. North-Holland, 1979. doi: 10.1016/S0167-5060(08)70342-X.
- [4] R. Bunel, I. Turkaslan, P. H. S. Torr, P. Kohli, and P. K. Mudigonda. A unified view of piecewise linear neural network verification. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, 2018.
- [5] R. Bunel, J. Lu, I. Turkaslan, P. H. S. Torr, P. Kohli, and M. P. Kumar. Branch and bound for piecewise linear neural network verification. *Journal of Machine Learning Research*, 21(42):1–39, 2020.
- [6] R. Ehlers. Formal verification of piece-wise linear feed-forward neural networks. *arXiv preprint arXiv:1705.01320*, 2017. doi: 10.48550/arXiv.1705.01320.
- [7] M. Everett. Neural network verification in control. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021. doi: 10.1109/CDC45484.2021.9683154.

- [8] M. Fazlyab, M. Morari, and G. J. Pappas. Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming. *IEEE Transactions on Automatic Control*, 67(1):1–15, 2022. doi: 10.1109/TAC.2020.3046193.
- [9] T. Gehr, M. Mirman, D. Drachler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev. Ai2: Safety and robustness certification of neural networks with abstract interpretation. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 3–18, 2018. doi: 10.1109/SP.2018.00058.
- [10] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *Computer Aided Verification*, volume 10426 of *Lecture Notes in Computer Science*, pages 97–117, 2017. doi: 10.1007/978-3-319-63387-9_5.
- [11] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. doi: 10.1137/S1052623400366802.
- [12] J. B. Lasserre and M. Putinar. Positivity and optimization for semi-algebraic functions. *SIAM Journal on Optimization*, 20(6):3364–3383, 2010. doi: 10.1137/090775221.
- [13] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, volume 149 of *IMA Volumes in Mathematics and its Applications*, pages 157–270. Springer, 2009.
- [14] C. Liu, T. Arnon, C. Lazarus, C. Strong, C. Barrett, and M. J. Kochenderfer. Algorithms for verifying deep neural networks. *Foundations and Trends in Optimization*, 4(3–4):244–404, 2021. doi: 10.1561/24000000035.
- [15] H. Marumo, S. Kim, and M. Yamashita. T-semidefinite programming relaxation with third-order tensors for constrained polynomial optimization. *Computational Optimization and Applications*, 89(1):183–218, 2024. doi: 10.1007/s10589-024-00582-8.
- [16] M. H. Meng, G. Bai, S. G. Teo, Z. Hou, Y. Xiao, Y. Lin, and J. S. Dong. Adversarial robustness of deep neural networks: A survey from a formal verification perspective. *IEEE Transactions on Dependable and Secure Computing*, 2022. doi: 10.1109/TDSC.2022.3179131.
- [17] M. Newton and A. Papachristodoulou. Sparse polynomial optimisation for neural network verification. *Automatica*, 157:111233, 2023. doi: 10.1016/j.automatica.2023.111233.
- [18] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.
- [19] V. Tjeng, K. Y. Xiao, and R. Tedrake. Evaluating robustness of neural networks with mixed integer programming. In *International Conference on Learning Representations (ICLR)*, 2019.
- [20] H. Waki, S. Kim, M. Kojima, and M. Muramatsu. Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM Journal on Optimization*, 17(1):218–242, 2006. doi: 10.1137/050623802.
- [21] J. Wang, V. Magron, and J.-B. Lasserre. TSSOS: A moment-SOS hierarchy that exploits term sparsity. *SIAM Journal on Optimization*, 31(1):30–58, 2021. doi: 10.1137/19M1307871.

- [22] J. Wang, V. Magron, J.-B. Lasserre, and N. H. A. Mai. CS-TSSOS: Correlative and term sparsity for large-scale polynomial optimization. *ACM Transactions on Mathematical Software*, 48(4), 2022. doi: 10.1145/3569709.
- [23] R. Y. Zhang. On the tightness of semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2020)*, 2020.

A A Concrete Example Illustrating the Theory of Sections 3–5

In this appendix, we present a small concrete instance that illustrates the framework developed in Sections 3–5. We first show how the input-sharing graph, the matching structure, and the regular rank-one edge condition are obtained from the input data $(W, b, \mathcal{U})$. We then compute the local optimal values m_e and m_3 through the component decomposition and derive the optimal value $\gamma^* = 13/2$ of the original ReLU verification problem. Finally, we give a reduced certificate of relaxation order 2 for the edge block and confirm tightness of the block-sparse SOS relaxation for this instance.

Appendix A.1 describes the original verification problem and the reduced data. Appendix A.2 constructs the input-sharing graph and verifies that its unique edge is a regular rank-one edge. Appendix A.3 computes the optimal value $\gamma^* = 13/2$ from the component decomposition. Appendix A.4 introduces the reduced polynomial optimization problem for the edge block and presents a relaxation-order-2 certificate.

A.1 The original ReLU verification problem and the reduced data

Let the input dimension be $n_0 = 4$, the number of hidden units be $n_1 = 4$, and the output dimension be $n_2 = 1$. Let the input set be

$$\mathcal{U} = [-1, 1]^4. \quad (\text{A.1})$$

We define the weight matrix and bias by

$$W = \begin{pmatrix} \frac{1}{2} & 1 & 0 & 0 \\ -\frac{1}{2} & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad b = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \end{pmatrix},$$

and the output map by

$$A = \begin{pmatrix} 1 & 2 & -1 & 1 \end{pmatrix}, \quad a = 0.$$

Let the linear specification vector be $c = 1$. Then the original ReLU verification problem is

$$\gamma^* = \max\{Ax \mid u \in \mathcal{U}, v = Wu + b, x = \text{ReLU}(v)\}. \quad (\text{A.2})$$

The pre-activations are given by

$$v_1 = \frac{1}{2}u_1 + u_2, \quad v_2 = -\frac{1}{2}u_1 + u_3, \quad v_3 = u_4, \quad v_4 = u_2 + 2.$$

Therefore, from (A.1) one immediately obtains

$$v_1 \in [-\frac{3}{2}, \frac{3}{2}], \quad v_2 \in [-\frac{3}{2}, \frac{3}{2}], \quad v_3 \in [-1, 1], \quad v_4 \in [1, 3].$$

These interval bounds show that units 1, 2, 3 are undecided, while unit 4 is a determined positive unit. Hence,

$$\mathcal{J}_u = \{1, 2, 3\}, \quad \mathcal{J}^+ = \{4\}, \quad \mathcal{J}^- = \emptyset.$$

Following the reduced notation of Section 3, the data corresponding to the undecided units are

$$\bar{W} = \begin{pmatrix} \frac{1}{2} & 1 & 0 & 0 \\ -\frac{1}{2} & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \bar{b} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \quad \bar{A} = \begin{pmatrix} 1 & 2 & -1 \end{pmatrix}.$$

On the other hand, the contribution of the determined positive unit 4 is

$$A_{:, \mathcal{J}^+} W_{\mathcal{J}^+, :} = \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix}, \quad A_{:, \mathcal{J}^+} b_{\mathcal{J}^+} = 2,$$

and after incorporating this into the objective, the reduced verification problem becomes

$$\gamma^* = \max\{x_1 + 2x_2 - x_3 + u_2 + 2 \mid u \in [-1, 1]^4, \ v = \bar{W}u, \ x = \text{ReLU}(v)\}.$$

A.2 Input-sharing graph, matching structure, and regular rank-one edge

The input supports of the undecided units are

$$S_1 = \{1, 2\}, \quad S_2 = \{1, 3\}, \quad S_3 = \{4\}.$$

Hence the edge set of the input-sharing graph $H = (\mathcal{J}_u, E)$ is

$$E = \{\{1, 2\}\}.$$

That is, units 1 and 2 share the input coordinate u_1 , whereas unit 3 is an isolated vertex. Thus the input-sharing graph is a matching, and the global decomposition in the paper appears here in its simplest nontrivial form.

For the edge $e = \{1, 2\}$, the shared input set and private sets are

$$S_e = \{1\}, \quad P_1 = \{2\}, \quad P_2 = \{3\}.$$

The shared rank is

$$q_e = \dim \text{span} \left\{ \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} -1 \\ 2 \end{pmatrix} \right\} = 1,$$

so condition (3.2) holds.

Furthermore,

$$v_1 = u_2 + \frac{1}{2}u_1, \quad v_2 = u_3 - \frac{1}{2}u_1,$$

so the lower and upper bounds of the private and shared parts are

$$\begin{aligned} \underline{p}_{1,e} &= -1, & \bar{p}_{1,e} &= 1, & \underline{s}_{1,e} &= -\frac{1}{2}, & \bar{s}_{1,e} &= \frac{1}{2}, \\ \underline{p}_{2,e} &= -1, & \bar{p}_{2,e} &= 1, & \underline{s}_{2,e} &= -\frac{1}{2}, & \bar{s}_{2,e} &= \frac{1}{2}. \end{aligned}$$

Therefore the edge-regularity condition (3.3) becomes

$$-1 + \frac{1}{2} < 0 < 1 - \frac{1}{2},$$

which holds for both units. Thus the unique edge $e = \{1, 2\}$ is a regular rank-one edge.

A.3 Component decomposition and the optimal value of the original ReLU problem

In this example, the set of isolated vertices is

$$\mathcal{I} = \{3\}.$$

Using the notation of Section 4, we then have

$$T_1 = \{2\}, \quad T_e = \{1\}, \quad T_2 = \{3\}, \quad T_3 = \{4\}, \quad T_0 = \emptyset.$$

Moreover, the coefficients defined in (4.2) are

$$\alpha_1 = 1, \quad \alpha_2 = 2, \quad \alpha_3 = -1, \quad (\delta_1, \delta_2, \delta_3, \delta_4) = (0, -1, 0, 0).$$

Hence the local affine objectives are given by

$$\Theta_e^{\text{lin}} = -x_1 - 2x_2 - u_2,$$

and

$$\Theta_3^{\text{lin}} = x_3.$$

Therefore the gap polynomial decomposes as

$$p_\gamma = \gamma - 2 + \Theta_e^{\text{lin}} + \Theta_3^{\text{lin}}. \tag{A.3}$$

On the isolated-vertex side, the local optimal value is immediately

$$m_3 = 0.$$

For the edge component, introduce the reduced coordinates

$$t := u_1, \quad p_1 := u_2, \quad p_2 := u_3, \quad z_1 = p_1 + \frac{1}{2}t, \quad z_2 = p_2 - \frac{1}{2}t.$$

The edge-local objective is expressible in the rank-one reduced variables introduced in Section 5: its only input-linear term is $-u_2 = -p_1$, and it contains no input-linear component orthogonal to the reduced shared scalar $t = u_1$. Then the local objective becomes

$$\Theta_e^{\text{lin}} = -p_1 - x_1 - 2x_2. \tag{A.4}$$

For fixed t , minimizing Θ_e^{lin} is equivalent to maximizing $p_1 + x_1 + 2x_2$. Since $x_1 = \text{ReLU}(p_1 + \frac{1}{2}t)$ is nondecreasing in p_1 and $x_2 = \text{ReLU}(p_2 - \frac{1}{2}t)$ is nondecreasing in p_2 , the optimum is attained at $p_1 = p_2 = 1$. In that case,

$$z_1 = 1 + \frac{1}{2}t, \quad z_2 = 1 - \frac{1}{2}t,$$

so the quantity to be maximized reduces to

$$1 + \frac{1}{2}t + 2\left(1 - \frac{1}{2}t\right) + 1 = 4 - \frac{1}{2}t.$$

Therefore the maximum value is $9/2$ at $t = -1$, and hence

$$m_e = -\frac{9}{2}.$$

Consequently,

$$\eta_{\gamma^*} + m_e + m_3 = 0,$$

so the optimal value of the reduced problem is

$$\gamma^* = \frac{13}{2}. \quad (\text{A.5})$$

Moreover, in the original ReLU problem (A.2), if we take

$$u^* = (-1, 1, 1, -1), \quad (\text{A.6})$$

then

$$v = (\frac{1}{2}, \frac{3}{2}, -1, 3), \quad x = (\frac{1}{2}, \frac{3}{2}, 0, 3),$$

and

$$Ax = \frac{1}{2} + 2 \cdot \frac{3}{2} + 3 = \frac{13}{2}.$$

Thus (A.6) gives an explicit optimal solution of the original verification problem.

A.4 A reduced-edge certificate of relaxation order 2

We next give an order-2 local certificate for the unique edge block $e = \{1, 2\}$. In this example, let the variables of the reduced edge block be

$$(t, z_1, x_1, z_2, x_2),$$

and define the target polynomial by

$$\tilde{p}_e(t, z_1, x_1, z_2, x_2) := \frac{9}{2} + \frac{1}{2}t - z_1 - x_1 - 2x_2.$$

This is obtained from (A.4) by substituting $p_1 = z_1 - \frac{1}{2}t$, and it is precisely the gap polynomial corresponding to the local optimum $m_e = -9/2$.

For this reduced block, the box-type generators are

$$\begin{aligned} g_1 &= 1 + t, & g_2 &= 1 - t, \\ g_3 &= z_1 + 1 - \frac{1}{2}t, & g_4 &= 1 + \frac{1}{2}t - z_1, \\ g_5 &= x_1, & g_6 &= x_1 - z_1, \\ g_7 &= z_2 + 1 + \frac{1}{2}t, & g_8 &= 1 - \frac{1}{2}t - z_2, \\ g_9 &= x_2, & g_{10} &= x_2 - z_2, \end{aligned}$$

and the ideal generators arising from complementarity are

$$h_1 = x_1(x_1 - z_1), \quad h_2 = x_2(x_2 - z_2).$$

Here g_1, g_2 encode the shared scalar constraint $t \in [-1, 1]$. The generators g_3, g_4 encode $z_1 \in [-1 + \frac{1}{2}t, 1 + \frac{1}{2}t]$, while g_7, g_8 encode $z_2 \in [-1 - \frac{1}{2}t, 1 - \frac{1}{2}t]$. Finally, g_5, g_6 and g_9, g_{10} correspond to the linear ReLU inequalities for $x_1 = \text{ReLU}(z_1)$ and $x_2 = \text{ReLU}(z_2)$, respectively.

The desired membership is

$$\tilde{p}_e = \sigma_0 + \sum_{r=1}^{10} \sigma_r g_r + \tau_1 h_1 + \tau_2 h_2, \quad (\text{A.7})$$

where

$$\sigma_0 \text{ is SOS with } \deg \sigma_0 \leq 4, \quad \sigma_r \text{ is SOS with } \deg \sigma_r \leq 2 \ (r = 1, \dots, 10), \quad \deg \tau_1, \deg \tau_2 \leq 2.$$

Equivalently, σ_0 admits a Gram representation on a basis of polynomials of degree at most 2, while each σ_r admits a Gram representation on a basis of polynomials of degree at most 1. This is the concrete meaning of “relaxation order 2” in the present example.

Solving this SDP with MOSEK yields the termination status **OPTIMAL**, objective value 0, and a maximum coefficient residual in the polynomial identity equal to

$$1.49 \times 10^{-8}.$$

Hence (A.7) is satisfied numerically to high precision. This confirms that a local certificate of relaxation order 2 can be constructed for the edge block in this instance.

For the isolated vertex 3, one has

$$\Theta_3^{\text{lin}} - m_3 = x_3,$$

and since x_3 itself is one of the generators, it is handled trivially at the same degree. Therefore, from the additive decomposition along the matching in (A.3) and (A.5),

$$p_{\gamma^*} = (\Theta_e^{\text{lin}} - m_e) + (\Theta_3^{\text{lin}} - m_3)$$

holds, and by adding the certificates of the edge block and the isolated-vertex block, we obtain

$$p_{\gamma^*} \in \mathcal{C}_{\text{sp}}^{(2)}.$$

Therefore, this example shows that

- (i) the input-sharing graph is a matching and its unique edge is a regular rank-one edge,
- (ii) the component decomposition along the matching yields $\gamma^* = 13/2$, and
- (iii) the same optimal value is reproduced by the order-2 block-sparse SOS relaxation.

This example therefore provides a concrete illustration of the argument developed in Sections 3–5.

GODAI AZUMA

Department of Mathematical and Computing Science, Institute of Science Tokyo, 2-12-1-W8-29 Oh-Okayama, Meguro-ku, Tokyo 152-8550, Japan

Department of Industrial and Systems Engineering, Aoyama Gakuin University, 5-10-1 Fuchinobe, Chuo-ku, Sagamihara-shi, Kanagawa 252-5258, Japan

E-mail address: `azuma@comp.isct.ac.jp`

SUNYOUNG KIM

Department of Mathematics, Ewha W. University, 52 Ewhayeodae-gil, Seodaemun-gu, Seoul 03760, Korea

E-mail address: `skim@ewha.ac.kr`

MAKOTO YAMASHITA

Department of Mathematical and Computing Science, Institute of Science Tokyo, 2-12-1-W8-29 Oh-Okayama, Meguro-ku, Tokyo 152-8550, Japan

E-mail address: `Makoto.Yamashita@comp.isct.ac.jp`