Issue 196: The SDP Primal Constraint Is a Submeasurement Constraint

2026-04-30

1 At a glance

- **Difficulty: easy.** This issue is about the feasible set of one displayed SDP, not about proving SDP strong duality. The strong-duality and complementary-slackness producer is documented separately.[1]
- **Estimated weight:** already resolved; replaying the repair is about **2–4 interface changes** and roughly **100–250 lines** of project-local definitions and adapters.
- **Mathlib/project split:** about **20% Mathlib** and **80% project-local**. Mathlib supplies finite sums and matrix order; the project supplies the submeasurement, measurement, SDP-wrapper, and completion interfaces.
- **Key Mathlib inputs:** finite sums over the polynomial set, matrix positive-semidefinite/order facts, and scalar multiplication of positive operators.

2 Key theorem forms

The resolved formal interface should isolate four statements.

1. **Submeasurement primal form.** The primal variable is a family $T = \{T_g\}$ with $T_g \geq 0$ and $\sum_g T_g \leq I$.

[1] The full strong-duality producer is discussed in `docs/paper-gaps/issue-1230-self-improvement-sdp-usage.tex`.

2. **Strict feasible witness.** For $M = |\mathcal{P}|$, the family $T_g = (2M)^{-1}I$ satisfies $\sum_g T_g = \frac{1}{2}I$ and is strictly feasible.

3. **Saturation as output.** The identity $\sum_g T_g = I$ belongs to the selected optimal witness, together with $T_g Z = T_g A_g$; it is not a feasibility condition.

4. **Completion only downstream.** If a later wrapper needs a full measurement, complete the chosen submeasurement explicitly, and do not identify that completion with the paper's strict Slater witness.

3 The primal feasible set in the paper

This note concerns the semidefinite program used in the self-improvement argument of [JNV+20, Subsection *A semidefinite program*]. [2] Let $\mathcal{P}$ be the finite set of degree-bounded polynomials, and set

$$A_g = \mathbb{E}_{u\in\mathbb{F}_q^m} A^u_{g(u)} \qquad (g \in \mathcal{P}).$$

The primal SDP in the paper is

$$\begin{aligned} \sup_{\{T_g\}_{g\in\mathcal{P}}} \quad & \sum_{g\in\mathcal{P}} \mathrm{Tr}(T_g A_g) && (\texttt{eq:primal-objective}) \\ \text{subject to} \quad & T_g \geq 0 \qquad (g \in \mathcal{P}), \\ & \sum_{g\in\mathcal{P}} T_g \leq I. \end{aligned}$$

Thus a feasible primal variable is a submeasurement: a family of positive semidefinite effects whose total mass is bounded above by the identity. It is not, at this point, required to be a measurement.

The same lemma later asserts the existence of an optimal primal-dual pair $(\{T_g\}, Z)$ for which

$$\sum_{g\in\mathcal{P}} T_g = I, \qquad T_g Z = T_g A_g \quad (g \in \mathcal{P}). \qquad (\texttt{lem:sdp})$$

This equality is a conclusion about a chosen optimal solution. It is not the feasibility constraint of the primal SDP.

[2] The mismatch was recorded in GitHub issue #196 and repaired in GitHub pull request #251. The documentation issue for this note is GitHub issue #935. In the TeX source consulted here, the primal SDP is in `references/ldt-paper/self_improvement.tex`, lines 64–74; the SDP lemma is in lines 82–88; and the Slater witness is in lines 168–176.

4 Why equality is a stronger formulation

Replacing $\sum_g T_g \leq I$ by $\sum_g T_g = I$ changes the feasible set. The change is proper. If $M = |\mathcal{P}|$, the paper's strict primal witness is

$$T_g = \frac{1}{2M} I \qquad (g \in \mathcal{P}).$$

Its total mass is

$$\sum_{g \in \mathcal{P}} T_g = \frac{1}{2} I,$$

so it is feasible for the paper's SDP and satisfies the inequality strictly, but it is not feasible for the equality-constrained SDP.

This point matters for the proof. The family above is the strict-feasibility witness used to invoke Slater's condition. It is deliberately not saturated: the missing mass is represented in the canonical SDP by the additional slack block $X_{M+1,M+1}$. Complementary slackness then forces this slack block to vanish at an optimal solution, which yields $\sum_g T_g = I$ for the selected optimum. If equality is imposed at the start, this derivation is replaced by an assumption and the strict feasible point used in the paper is excluded.

Thus the equality formulation is not a harmless restatement of the displayed primal program. It is a smaller SDP. It may still contain saturated optima, but that saturation is exactly part of what the paper proves from the submeasurement program.

5 The current formal encoding

The current development has a public paper-local type for submeasurements: a family $T = \{T_g\}$ with $T_g \geq 0$, total $\sum_g T_g$, and proof that this total is at most I. A measurement is the special case where the total is exactly I. [3] The self-improvement SDP definitions now use the submeasurement type for the primal variable. In particular, the primal contribution, primal objective, and complementary-slackness equation all take T to be a submeasurement, not a measurement. [4] The proof that the sandwiched family is a submeasurement

[3] These are the declarations MIPStarRE.LDT.SubMeas and MIPStarRE.LDT.Measurement in `MIPStarRE/LDT/Basic/SubMeasurementCore.lean`, lines 13–28.

[4] See MIPStarRE.LDT.SelfImprovement.sdpPrimalContributionOperator, MIPStarRE.LDT.SelfImprovement.sdpPrimalObjectiveOperator, MIPStarRE.LDT.SelfImprovement.sdpPrimalObjective, and MIP-

uses the inequality $\sum_g T_g \leq I$: after grouping the summands by the value of $h(u)$, the filtered sum is bounded by the total mass of T, and hence by I.

The strict primal witness from the paper is represented before completion as a submeasurement. The declaration `sdpStrictPrimalSubMeas` has outcomes $T_g = (2|\mathcal{P}|)^{-1}I$, and the theorem `sdpStrictPrimalSubMeas_total` proves that its total is $\frac{1}{2}I$. [5] This is the formal counterpart of the paper's Slater witness. Earlier Lean versions also exposed a declaration sdpPrimalWitness, which completed this submeasurement at the zero polynomial for a reduced theorem named sdp. That completion wrapper and the reduced theorem it served were retired in PR #2430. The active SDP route keeps the strict witness as a submeasurement and obtains a saturated measurement only from the optimal-pair conclusion. Thus completion is not part of the primal feasible set of the SDP.

The informal blueprint already agrees with the paper on this point: its SDP lemma states the primal with $\sum_g T_g \leq I$, and only afterwards states the existence of an optimal pair whose total is I. The resolved mismatch was therefore confined to the formal encoding, not a three-way disagreement with the blueprint. [6]

At the statement layer, the active source-facing route exposes the primal witness as a measurement inside SdpStatementWithSlackness.witness. The associated witness record is indexed by the underlying submeasurement and still contains a field `primalTotalOperator` asserting that this total is I; this is the paper's saturated-optimum conclusion, not a definition of primal feasibility. [7] The matrix-level realization follows the same convention: primal variables are matrix submeasurements, while the slackness statement stores the completed matrix measurement obtained from the identity-total conclusion. [8]

StarRE.LDT.SelfImprovement.sdpComplementarySlacknessEquation in `MIPStarRE/LDT/SelfImprovement/Defs.lean`, lines 139–184. The same file also defines the sandwiched family `sandwichedPolynomialSubMeasAt` and the averaged family `averagedSandwichedPolynomialSubMeas` with T a submeasurement, in lines 198–333.

[5]These are in `MIPStarRE/LDT/SelfImprovement/Defs.lean`, lines 37–97.

[6]See `blueprint/src/chapter/ch07_self_improvement.tex`, lines 281–304.

[7]See MIPStarRE.LDT.SelfImprovement.SdpOptimalPairWithSlackness and MIPStarRE.LDT.SelfImprovement.SdpStatementWithSlackness in `MIPStarRE/LDT/SelfImprovement/Theorems/Statements.lean`, lines 61–134. The source-facing theorem is MIPStarRE.LDT.SelfImprovement.sdp_statement_with_slackness in `MIPStarRE/LDT/SelfImprovement/Theorems/Results/HelperCompleteness/Bracketed.lean`, lines 493–510, using the canonical block-SDP transport from `MIPStarRE/LDT/SelfImprovement/Theorems/Results/SdpMatrixBridge.lean`.

[8]The matrix-level primal-side declarations `matrixSdpPrimalContributionOperator`, `matrixSdpPrimalObjective`, `matrixSdpComplementarySlacknessDefect`, and `MatrixSd`

6 Conclusion

The verdict is that issue #196 recorded a genuine mismatch. The paper's primal SDP is over submeasurements satisfying $\sum_g T_g \leq I$. The earlier equality-based formal encoding treated the primal variable as a measurement and thereby excluded the paper's strict-feasible Slater witness $T_g = (2|\mathcal{P}|)^{-1} I$.

The mismatch is now resolved at the SDP-interface level. The formal objective, complementary-slackness relation, strict primal witness, and sandwiched-family construction all expose the feasible primal variable as a submeasurement. Saturation is recorded only for the selected optimal witness, and the slackness-carrying SDP statement stores that saturated witness as a measurement.

The former reduced interface has been retired: PR #2430 removed sdpPrimalWitness, the reduced SdpStatement structure, the theorem named sdp, and the wrapper selfImprovementHelperConstruction. The current LEAN route for Lemma `lem:sdp` is sdp_statement_with_slackness, whose output is SdpStatementWithSlackness. The canonical block-SDP strong-duality route is documented in `docs/paper-gaps/issue-1230-self-improvement-sdp-usage.tex`. Future refactors should preserve the submeasurement feasible set until the point where the paper derives the saturated optimal witness $\sum_g T_g = I$.

References

[JNV+20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test, 2020. arXiv:2009.12982.

`pOptimalWitness` are in `MIPStarRE/LDT/SelfImprovement/MatrixRealization.lean`, lines 42–96; the witness structure itself is in lines 81–96.