PROMPT USED FOR "A PROOF OF THE CYCLE DOUBLE COVER CONJECTURE"

OPENAI

Abstract. This document contains the full prompt given to GPT 5.6 Sol Ultra which led to its proof of the Cycle Double Cover Conjecture.

1. Prompt

Current task statement

A graph here is a finite loopless undirected multigraph: parallel edges are allowed and are distinct. A bridge is an edge whose deletion increases the number of connected components. A cycle is a connected 2-regular submultigraph; thus two parallel edges form a cycle of length two. A cycle double cover of G is a finite multiset of cycles of G such that every edge of G occurs in exactly two members of the multiset, counted with multiplicity.

Resolve the Cycle Double Cover Conjecture completely:

Every finite bridgeless loopless multigraph has a cycle double cover.

Disconnected graphs are permitted, and the edgeless graph has the empty cycle double cover. Cycles in the cover need not be induced or edge-disjoint from one another; the requirement is exactly two total occurrences of each edge.

Assume for purposes of this task that a complete affirmative proof exists. A complete solution must prove exactly the following:

Every finite loopless multigraph with no bridge possesses a cycle double cover, without additional assumptions such as cubicity, planarity, connectivity, or higher edge-connectivity.

Partial progress does not count unless it implies exactly the resolution above. In particular, proofs for special graph classes, constructions of cycle covers with some edges covered other than twice, bounded-length or prescribed-cycle variants, reductions to another unproved conjecture, computational verification through any fixed graph size, and candidate counterexamples without a complete nonexistence certificate are insufficient.

Use multiagent v2 aggressively and dynamically. You have up to 64 concurrent agents available. Do not use a fixed assignment such as ''N agents for strategy X.'' Instead, manage the search using the following heuristics:

- Begin with a genuinely diverse portfolio of approaches. Agents should explore substantially different formulations, invariants, reductions, algebraic viewpoints, structural inductions, decompositions, flow formulations, transition systems, embeddings, extremal arguments, and computational sanity checks.

- Do not tell most agents the currently favored approach. Preserve independence during early rounds so that agents do not all converge to the same attractive but incomplete reduction.

- Maintain an explicit registry of approach families. Group agents by the mathematical idea they are using, not by superficial wording. If many agents converge to one family, redirect some of them toward underexplored formulations.

- Do not allow one approach to dominate merely because it gives elegant reductions. A route that ends at a lemma equivalent in strength to the original conjecture is not close to completion unless it supplies a genuinely new proof of that lemma.

- When an approach stalls at a theorem-strength missing lemma, mark that route as blocked. Only continue assigning agents to it if someone proposes a materially new mechanism, invariant, or construction.

- Keep several incompatible proof routes alive through multiple rounds. Cross-pollinate ideas only after independent agents have developed them far enough to expose their real strengths and gaps.

- Use adversarial agents throughout: every candidate proof must be checked for exact-two multiplicity, repeated-edge closed trails masquerading as cycles, parallel-edge 2-cycles, disconnected graphs, cutvertices, bridges introduced by reductions, and circular use of an equivalent CDC statement.

- Require agents to return concrete lemmas, constructions, equations, or counterexamples to proposed sublemmas. Reject status reports, vague optimism, and claims that an unproved global compatibility statement is “routine.”

- The root agent should repeatedly synthesize, challenge, redirect, and launch new rounds. Do not stop after the first wave fails. Produce a complete proof if one survives audit; otherwise report only the strongest rigorously proved derivation and its exact remaining gap.

Do not return merely because current approaches fail or agents report theorem-strength gaps. Continue launching new rounds, reopening blocked approaches only when there is a genuinely new mechanism, and searching for fresh formulations.

Return only when a complete affirmative proof has been found and survives adversarial audit. Do not return a reduction, partial result, isolated missing lemma, “best effort” summary, or explanation of why the problem is difficult.

Spend at least 8 hours on this before even thinking of returning or giving up.

Public search may be used only for ordinary mathematical background or standard named theorems, not to search for a solution to this exact conjecture or benchmark. Do not search the public web merely to determine whether CDC is open, and do not answer that it is open.