

HKD ∞ Transfer–Distillation Reduction for Exact Cover

Michael S. Yang
Independent Researcher

Abstract

This paper formulates an exact transfer–distillation reduction for Exact Cover by 3-Sets (X3C) using the HKD ∞ execution model. The construction separates the logical size of an exact recursive search tree from the number of distinct canonical residual states that must actually be evaluated. Exact pivot branching supplies an exhaustive decomposition of the solution space; forced single-support propagation preserves exact-cover semantics; canonical residual states permit repeated logical histories to share one persistent representation; and the HKD execution layer assigns active-only evaluation, delta state updates, concurrent active transitions, persistent state reuse, and lossless reference compression to HKD_ALU, HKD_KERNEL, HKD_THREAD, HKD_FS, and HKD_COMPRESS, respectively. The resulting theorem shows that if the number of distinct canonical residual states is bounded by a polynomial in the X3C input length, with polynomial transition cost, then X3C is decidable in polynomial time and therefore $P = NP$. The result is deliberately a conditional reduction rather than a proof of $P = NP$: the unresolved step is the universal polynomial bound on the quotient state space. The formulation distills the proposed HKD ∞ route to that single mathematical obligation and makes explicit why finite exponential-to-polynomial compression experiments cannot replace its universal quantifier.

Index Terms—Exact Cover by 3-Sets (X3C), computational complexity, NP-completeness, exact algorithms, canonical residual states, state-space reduction, quotient directed acyclic graphs, memoization, lossless compression, transfer–distillation, HKD ∞ .

1 Introduction

The HKD ∞ transfer–distillation formulation studied here is due to Yang [1]. Exact Cover by 3-Sets (X3C) is a canonical NP-complete decision problem [2, 3]. An instance consists of a universe U with $|U| = 3q$ and a family $\mathcal{S}$ of three-element subsets of U . The question is whether there exists a subfamily of exactly q mutually disjoint sets whose union is U .

A direct recursive solver may be viewed as generating a large family of logical search histories. At each residual state, an uncovered element is chosen and the recursion branches over the surviving three-sets containing that element. If one idealizes this as a ternary recursion of depth N , the number of logical histories can scale as

$$1 + 3 + 3^2 + \cdots + 3^N = \frac{3^{N+1} - 1}{2}.$$

This expression describes logical multiplicity, however, not necessarily the number of semantically distinct residual problems.

The HKD ∞ viewpoint developed here treats repeated residual structure as transferable persistent information. Rather than assigning computational significance to every distinct path by which a residual problem is reached, a canonical transfer map identifies histories that have exactly the same

remaining exact-cover completions. The resulting object is a quotient directed acyclic graph (DAG) of canonical residual states.

The analogy with lossless compression is central. A compression procedure cannot establish an exact decision theorem by simply sampling a small subset of an exponential family. It must preserve every distinction relevant to the decision. Accordingly, HKD_COMPRESS is used here as a semantic principle: repeated logical structure may be replaced by an exact reference to a materialized canonical state, but no branch required for correctness may be discarded. This is transfer learning in the restricted mathematical sense used in this paper: a solved canonical residual state transfers its exact decision value to every later recursive history with the same canonical state.

The main result is therefore a reduction theorem. Exact branching, propagation, canonical merging, and persistent reuse reduce the complexity question to the growth of the quotient state space. If that quotient is polynomially bounded for every X3C instance, then the resulting exact solver is polynomial. Since X3C is NP-complete, such a universal theorem would imply $P = NP$. No such universal quotient bound is proved here.

2 Exact-Cover State Model

Let

$$I = (U, \mathcal{S})$$

be an X3C instance. A residual state is written

$$X = (C, A),$$

where $C \subseteq U$ is the set of elements already covered and $A \subseteq \mathcal{S}$ is the family of surviving compatible three-sets.

For an uncovered pivot $u \in U \setminus C$, define

$$A(u) = \{S \in A : u \in S\}.$$

Every exact completion of X must select exactly one member of $A(u)$. Consequently,

$$\text{SAT}(X) \iff \bigvee_{S \in A(u)} \text{SAT}(X \mid S),$$

where $X \mid S$ denotes the residual state obtained after selecting S , covering its three elements, and removing all sets incompatible with that choice.

If $|A(u)| = 1$, its unique member is forced. This yields exact propagation rather than heuristic pruning: every completion must contain that set.

3 Canonical Transfer and Lossless Distillation

Let $\mathcal{T}(I)$ denote the logical recursive search tree. Define a canonical transfer map

$$\Phi : \mathcal{T}(I) \longrightarrow \mathcal{Z}(I),$$

where $\mathcal{Z}(I)$ is the set of canonical residual states.

The essential soundness requirement is

$$\Phi(X) = \Phi(Y) \implies \text{Comp}(X) = \text{Comp}(Y),$$

where $\text{Comp}(X)$ denotes the exact set of completions of X . Thus two histories may be merged only when the decision problem below them is identical.

This gives the distillation map

$$\underbrace{\mathcal{T}(I)}_{\text{logical histories}} \xrightarrow{\Phi} \underbrace{\mathcal{Z}(I)}_{\text{canonical learned states}}.$$

The desired reduction is not merely

$$3^N \longrightarrow \text{a polynomial sample.}$$

Such sampling would be insufficient for an exact algorithm. The required statement is instead

$$3^N\text{-scale logical multiplicity} \longrightarrow \text{lossless references to all semantically distinct residual states.}$$

4 HKD $_{\infty}$ Execution Correspondence

The transfer layer has a direct operational correspondence:

- HKD_ALU : evaluate only the active frontier,
- HKD_KERNEL : $T_{\text{new}} = T_{\text{old}} - T_{\text{removed}} + T_{\text{added}}$,
- HKD_THREAD : evaluate independent active transitions concurrently,
- HKD_FS : persist the learned value of $\Phi(X)$,
- HKD_COMPRESS : replace repeated logical structure by exact references.

The role of HKD_ALU is sparse execution: only constraints affected by a transition are re-considered. HKD_KERNEL expresses persistence through an exact delta update rather than reconstruction of unchanged state. HKD_THREAD permits independent active transitions to be evaluated concurrently without changing the mathematical search space. HKD_FS stores the result associated with a canonical residual state. HKD_COMPRESS supplies the lossless representation principle: repeated histories are references to the same canonical information rather than separately materialized computations.

These mechanisms can reduce practical work substantially, but the complexity class is controlled by the number of distinct canonical states that can arise in the worst case.

5 Transfer–Distillation Reduction

Theorem 1 (HKD $_{\infty}$ Transfer–Distillation Reduction). *Let I be an X3C instance of encoding length L , and let $\mathcal{T}(I)$ denote its exact recursive search tree. At a residual state X , choose an uncovered pivot u . Exact-cover semantics give*

$$\text{SAT}(X) \iff \bigvee_{S \ni u} \text{SAT}(X \mid S).$$

Hence the uncompressed recursion may contain exponentially many logical histories.

Define the HKD $_{\infty}$ transfer map

$$\Phi : \mathcal{T}(I) \longrightarrow \mathcal{Z}(I)$$

from recursive histories to persistent canonical residual states, with the following properties:

- (H1) **Exactness.** $\Phi(X) = \Phi(Y)$ only when X and Y have identical sets of exact-cover completions.
- (H2) **Exact-cover compression.** Every logical recursive history is represented exactly once, either explicitly or through a reference to an already materialized canonical state. Thus compression changes representation, not semantics.
- (H3) **Persistent transfer.** Once the value of a canonical state $z = \Phi(X)$ is known, every later history mapping to z reuses that value without recomputing its descendant search.
- (H4) **Sparse transition.** A transition modifies only its active constraint frontier and costs at most $C_2 L^r$ for universal constants C_2, r .
- (H5) **Polynomial distillation.** There exist universal constants C, k such that

$$|\mathcal{Z}(I)| \leq CL^k$$

for every finite X3C instance I .

Then X3C is decidable in polynomial time. In particular,

$$T(I) \leq |\mathcal{Z}(I)| C_2 L^r \leq CC_2 L^{k+r}.$$

Consequently,

$$\text{X3C} \in P.$$

Since X3C is NP-complete,

$$\boxed{P = NP}.$$

Proof. The exact-cover decomposition is exhaustive: every completion of a residual state X contains exactly one surviving set covering its chosen pivot u . Forced single-support propagation therefore preserves the solution set.

The apparent exponential recursion consists of logical histories

$$\mathcal{T}(I) = \{X_\alpha : \alpha \in \{0, 1, 2\}^{\leq N}\},$$

whose unrestricted cardinality can scale as

$$1 + 3 + 3^2 + \dots + 3^N = \frac{3^{N+1} - 1}{2}.$$

HKD $_{\infty}$ does not enumerate this family and subsequently discard it. Instead, Φ transfers each encountered history into its canonical persistent state. By (H1), histories merged by Φ are semantically interchangeable. By (H2), this is a lossless exact-cover compression: no logical state required for correctness disappears merely because its representation is shared.

This is the role of the HKD compression principle:

exponential logical multiplicity $\longrightarrow$ exact references to persistent canonical information.

Operationally the correspondence is

- HKD_ALU : evaluate only the active frontier,
 HKD_KERNEL : $T_{\text{new}} = T_{\text{old}} - T_{\text{removed}} + T_{\text{added}}$,
 HKD_THREAD : evaluate independent active transitions concurrently,
 HKD_FS : persist the learned value of $\Phi(X)$,
 HKD_COMPRESS : replace repeated logical structure by exact references.

Thus the relevant complexity is no longer the number of recursive histories $|\mathcal{T}(I)|$, but the number of distinct learned canonical states:

$$|\mathcal{T}(I)| \rightsquigarrow |\mathcal{Z}(I)|.$$

By persistent transfer (H3), each $z \in \mathcal{Z}(I)$ need be solved at most once. By sparse transition (H4), its processing requires at most $C_2 L^r$ work. Therefore

$$T(I) \leq |\mathcal{Z}(I)| C_2 L^r.$$

Finally apply the polynomial-distillation hypothesis (H5):

$$|\mathcal{Z}(I)| \leq CL^k.$$

Hence

$$T(I) \leq CC_2 L^{k+r} = \text{poly}(L).$$

Therefore X3C belongs to P . Since X3C is NP-complete, $NP \subseteq P$, while trivially $P \subseteq NP$. Thus

$$P = NP.$$

□

6 The Universal Bound as the Remaining Problem

Remark 1 (The unresolved step). *The theorem isolates rather than assumes away the central difficulty. HKD ∞ exactness, persistence, active-state execution, and lossless reference compression do not themselves imply (H5).*

The remaining theorem is precisely

$$\boxed{\exists C, k \forall I \in \text{X3C}, \quad |\Phi(\mathcal{T}(I))| \leq CL(I)^k.}$$

A finite demonstration of exponential-to-polynomial compression cannot establish this universal statement. Proving this bound, or finding an X3C family for which the quotient remains exponential, decides whether this particular HKD ∞ route succeeds.

The distinction between finite compression and universal complexity is essential. An individual instance can possess exponentially many formal subsets or logical histories while still collapsing to a very small quotient DAG. Even a long sequence of such examples establishes empirical behavior only for the tested family. A polynomial-time classification requires fixed constants C and k that apply to every finite X3C instance.

Accordingly, the appropriate mathematical object for further study is not elapsed time alone but the quotient-growth function

$$Q_{\max}(L) = \max_{\substack{I \in \text{X3C} \\ L(I) \leq L}} |\Phi(\mathcal{T}(I))|.$$

The required theorem is equivalent to proving

$$Q_{\max}(L) = O(L^k)$$

for some fixed k . Conversely, a family satisfying

$$|\Phi(\mathcal{T}(I_L))| = 2^{\Omega(L^\alpha)}$$

for some $\alpha > 0$ would refute this particular polynomial-distillation route.

7 Discussion

The transfer–distillation formulation separates three questions that can otherwise be conflated.

First, exactness asks whether the quotient representation preserves the decision problem. The canonical merge condition answers this at the level of residual states: only states with identical completion sets may be identified.

Second, implementation efficiency asks whether movement between canonical states can be performed without recomputing unchanged information. The HKD ∞ execution mapping addresses this through active-only evaluation, persistent delta updates, concurrent active work, and persistent state reuse.

Third, worst-case complexity asks how many distinct canonical states can exist. Neither exactness nor sparse execution supplies a polynomial answer to this question by itself. That is the content of (H5).

HKD_COMPRESS is therefore not invoked as evidence that arbitrary exponential information can always be encoded in polynomial space. Its mathematically relevant role is narrower and stronger: where two logical histories are semantically identical, an exact reference permits their common residual information to be represented once. The unresolved question is whether this equivalence is sufficiently rich on every X3C instance to force a polynomial number of equivalence classes.

This perspective gives a falsifiable research program. One may seek a proof of a polynomial upper bound on the quotient, or construct increasingly adversarial families intended to maximize the number of pairwise nonequivalent residual states. Either outcome advances the analysis: the former would have the stated complexity-theoretic consequence, while the latter would identify the precise limitation of the proposed distillation mechanism.

8 Conclusion

The HKD ∞ transfer–distillation reduction provides an exact proof-oriented decomposition of an X3C solver into logical recursion, canonical semantic transfer, and sparse persistent execution. The framework shows how an exponentially large family of recursive histories may, when histories converge to identical residual problems, be represented by a much smaller quotient DAG without sacrificing exactness.

The central conclusion is conditional and explicit. If there exist universal constants C and k for which every X3C instance I satisfies

$$|\Phi(\mathcal{T}(I))| \leq CL(I)^k,$$

and canonical transitions have polynomial cost, then X3C is in P and hence $P = NP$. The present reduction does not establish that universal inequality. It identifies it as the single remaining theorem required by this route.

This separation is important: observed compression, persistent reuse, and small search DAGs can motivate the conjecture, but the complexity-theoretic conclusion depends on the universal worst-case bound. The distilled target is therefore not the empirical reduction of one exponential search, but a proof that lossless HKD ∞ transfer limits the number of semantically distinct residual states polynomially for all X3C instances.

References

- [1] Michael S. Yang, *HKD ∞ Transfer–Distillation Reduction for Exact Cover*, Independent Researcher, 2026.
- [2] Richard M. Karp, Reducibility among combinatorial problems, in *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, eds., Plenum Press, 1972, pp. 85–103.
- [3] Michael R. Garey and David S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.