The Clean Paper

WHAT THE PAPER SAYS · WHAT IT DOESN'T · WHY IT MATTERS

Cryptographic proof kan lè hide secret rẹ̀ nípa jíjẹ́ kí missing simulator nira láti prove

Zero-knowledge proofs jẹ́ kí ẹnikan prove statement kan lái reveal witness. Classical theory sọ pé o kò lè ní full zero-knowledge pẹ̀lú one message, no trusted setup àti perfect soundness. Paper Rahul Ilango kò make impossibility yẹn vanish. Ó change target: dípò requiring pé simulator kan wà gidi, ó ask pé proof system tí a yan kò lè efficiently prove pé simulator kò exist. Under major proof-complexity àti cryptographic assumptions, èyí tó láti recover falsifiable, game-based consequences of zero-knowledge property by property. Point náà kì í ṣe plug-in internet primitive. Ó jẹ́ proof-theoretic way láti turn mathematical unprovability into cryptographic cover.

Written by Cinzia Vaglio · Cross-checked by Vera Rubin · Edited by Michele Renda · Figures by Nova Vela · AI-assisted, human-reviewed

Paper: *Gödel in Cryptography: Effectively Zero-Knowledge Proofs for NP with No Interaction, No Setup, and Perfect Soundness*, Rahul Ilango, FOCS 2025 / IACR ePrint 2025/1296 · DOI: 10.1109/FOCS63196.2025.00058

Ètan náà kì í ṣe láti fi èrí múlẹ̀ pé secret náà farapamọ́

Jẹ́ ká bẹ̀rẹ̀ pẹ̀lú èdà zero-knowledge tó rọrùn jù.

Alice fẹ́ convince Bob pé Sudoku puzzle kan ní solution. Bí ó bá send solution náà, Bob máa convinced, ṣùgbọ́n puzzle náà ti bàjẹ́. Ohun tí Alice fẹ́ yàtọ̀ díẹ̀: proof pé solution kan wà, láì reveal solution náà.

Ìlérí zero-knowledge proof nìyẹn. Prover (Alice) convince verifier (Bob) pé statement kan jẹ́ true, nígbà tí kò reveal ohunkóhun ju truth statement náà lọ.

ìṣòro náà ni pé ìlérí yìí ní cost. Ordinary mathematical proof ní àbùdá méjì tó convenient. Ó jẹ́ **ọ̀kan message**: o kọ ọ́ sílẹ̀, o fún ẹlòmíràn, o sì lọ. Ó tún jẹ́ **perfectly sound**: false statement kò ní valid proof rárá. Classical impossibility èsì sọ pé zero-knowledge gbọ́dọ̀ give up àbùdá méjèèjì — kì í ṣe pé méjèèjì pọ̀ nìkan; ọkọọkan wọn fúnra rẹ̀ off-limits.

Àkọ́kọ́, zero-knowledge proof nílò conversation. Bí Alice bá send message kan ṣoṣo, láìsí trusted setup tí a arrange tẹ́lẹ̀, zero-knowledge guarantee collapse — èyí sì hold bó tilẹ̀ jẹ́ pé o willing láti trade soundness púpọ̀ away.

Èkejì, zero-knowledge proof nílò tolerance kékeré fún àṣìṣe. Demanding perfect soundness turns out láti quietly destroy interaction náà: verifier tí a kò lè fool rárá, whatever random choices tó ṣe, lè kan fix choices yẹn tẹ́lẹ̀ — nígbà tí verifier bá predictable, Alice lè answer gbogbo nǹkan nínú message kan, èyí gan-an ni ọ̀ràn tí a ti mọ̀ pé ó break.

ìwé ìwádìí Rahul Ilango jẹ́ nípa ọ̀nà kan láti yí double wall yẹn ká. Kì í ṣe nípa pretending pé wall kò sí, kì í sì ṣe nípa producing classical zero-knowledge nínú setting tí kò ṣeé ṣe. Move náà subtle ju: weaken ohun tí “reveals nothing” túmọ̀ sí, ṣùgbọ́n weaken rẹ̀ ní ọ̀nà tó preserve security properties tí cryptographers lè ìdánwò gidi.

èsì náà ni a pe ní **effectively zero-knowledge**.

A proof without leaking the witness

The paper keeps the ordinary-proof shape by weakening what privacy must prove.

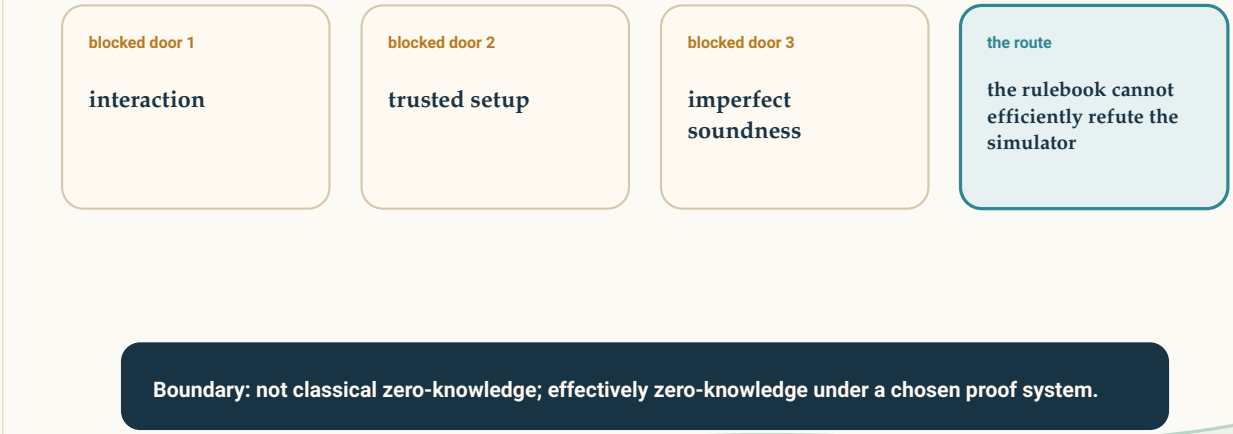

Zero-knowledge ti blocked ní doors méta — interaction, trusted setup, àti imperfect soundness. Construction llango yọ kọjá ní ẹnu-ọ̀nà mìíràn: rulebook náà kò lè efficiently refute simulator náà.

Original diagram — The Clean Paper CC BY 4.0

Classical privacy vs effective privacy

The relaxation is the point: the paper changes what the privacy claim has to establish.

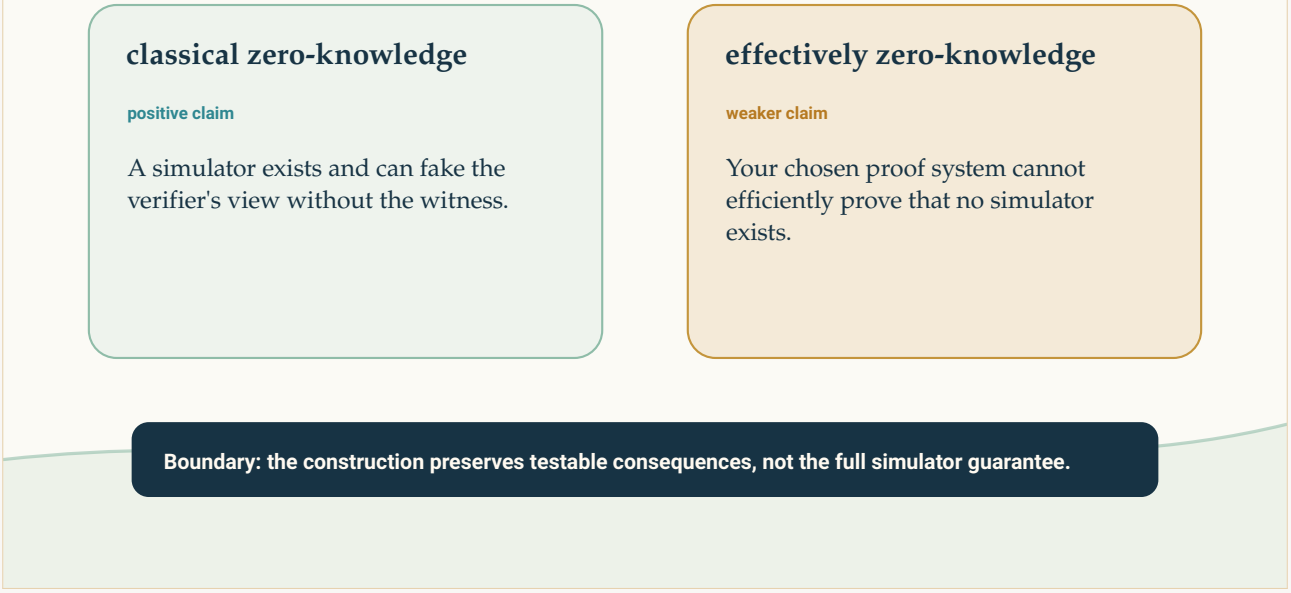

Classical zero-knowledge béèrè bóyá simulator kan wà; “effectively zero-knowledge” béèrè nìkan bóyá rulebook tí o yan lè efficiently fi ẹ̀rí múlẹ̀ pé simulator kò lè wà. Ìbéèrè tó weaker yìí ni ohun tó jẹ́ kí construction

náà keep ọ̀kan message, no setup àti perfect soundness.

ìdánwò atijọ́: simulator kan wà

Ọ̀nà classical láti formalize zero-knowledge lo fictional helper kan tí a ń pè ní **simulator**.

Idea náà ni èyí: imagine Jane, tí **kò** mọ secret Alice. Bí Jane bá lè generate, fúnra rẹ̀ patapata, proofs tí ó look gẹ́gẹ́ bí proofs tí Bob bá ti gba láti Alice, nígbà náà proofs Alice kò kọ́ Bob ohunkóhun tuntun. Jane ti lè fake experience náà láìsí secret Alice.

Torí náà classical zero-knowledge demands actual simulator kan. Efficient algorithm kan gbọ́dọ̀ wà tó lè produce fake-looking proofs láì mọ secret — *witness*, ní jargon; fún Sudoku, witness náà ni solved grid gan-an.

Definition náà powerful, ṣùgbọ́n ibẹ̀ gan-an ni old impossibility ti bite. Intuition ni èyí. Truly non-interactive proof jẹ́ string kan nìkan. Nígbà tí Bob bá ní string yẹn, ó lè fi hàn ẹlòmíràn: ó ti gain ability láti fi ẹ̀rí múlẹ̀ statement náà fún àwọn mìíràn, èyí sì ti sound bí sí i than “nothing.” Classical theorems sharpen intuition yẹn sí impossibilities tó wà lókè.

> **Properties mẹ́ta tí paper yìí insist lórí**
>
> Title ìwé ìwádìí náà name constraints mẹ́ta:
> **No interaction:** Alice send proof string kan. Kò sí back-and-forth protocol.
> **No setup:** Alice àti Bob kò rely lórí trusted common reference string tàbí pre-arranged public randomness mìíràn. Ọ̀pọ̀ ètò tí a pe ní “non-interactive zero-knowledge” ṣì rely lórí setup; ìwé ìwádìí yìí túmọ̀ sí zero setup.
> **Perfect soundness:** false statement kò ní valid proof. Kì í ṣe “almost never accepted”; valid proof kankan kò exist.
> Properties mẹ́ta yìí gan-an ni ordinary written mathematics ní — àti, gẹ́gẹ́ bí a ṣe explain lókè, classical zero-knowledge kò lè keep wọn.

Mega-Sudoku ẹ̀dà kan ti ìyàtọ̀ náà

Ẹ jẹ́ ká lo deliberately simplified way láti feel difference náà.

Má ṣe lo ordinary 9-by-9 Sudoku fún serious part of analogy. Ó kere ju, ó sì finite ju: computer lè kan solve rẹ̀, tàbí fi ẹ̀rí múlẹ̀ pé kò ní solution. Dípò bẹ́ẹ̀ imagine family of **MegaSudoku(n)** puzzles. Scale usual òfin up: yan block ìwọ̀n *n*, jẹ́ kí $N = n^2$, kí o sì build N by N grid tí a divide sí n by n blocks, pẹ̀lú N symbols. Ordinary Sudoku jẹ́ tiny $n = 3$, $N = 9$ ọ̀ràn nìkan: 9-by-9 grid, 3-by-3 blocks àti symbols mẹ́sàn-án. Proof-complexity ìtàn bẹ̀rẹ̀ gidi nígbà tí n lè grow, àti nígbà tí grid lè carry extra gadgets tí yóò jẹ́ kó behave bí SAT formula tí a wọ̀ ní aṣọ Sudoku. SAT formula jẹ́ list of yes/no constraints nìkan: ṣé o lè assign true/false values sí variables kí gbogbo constraints satisfy?

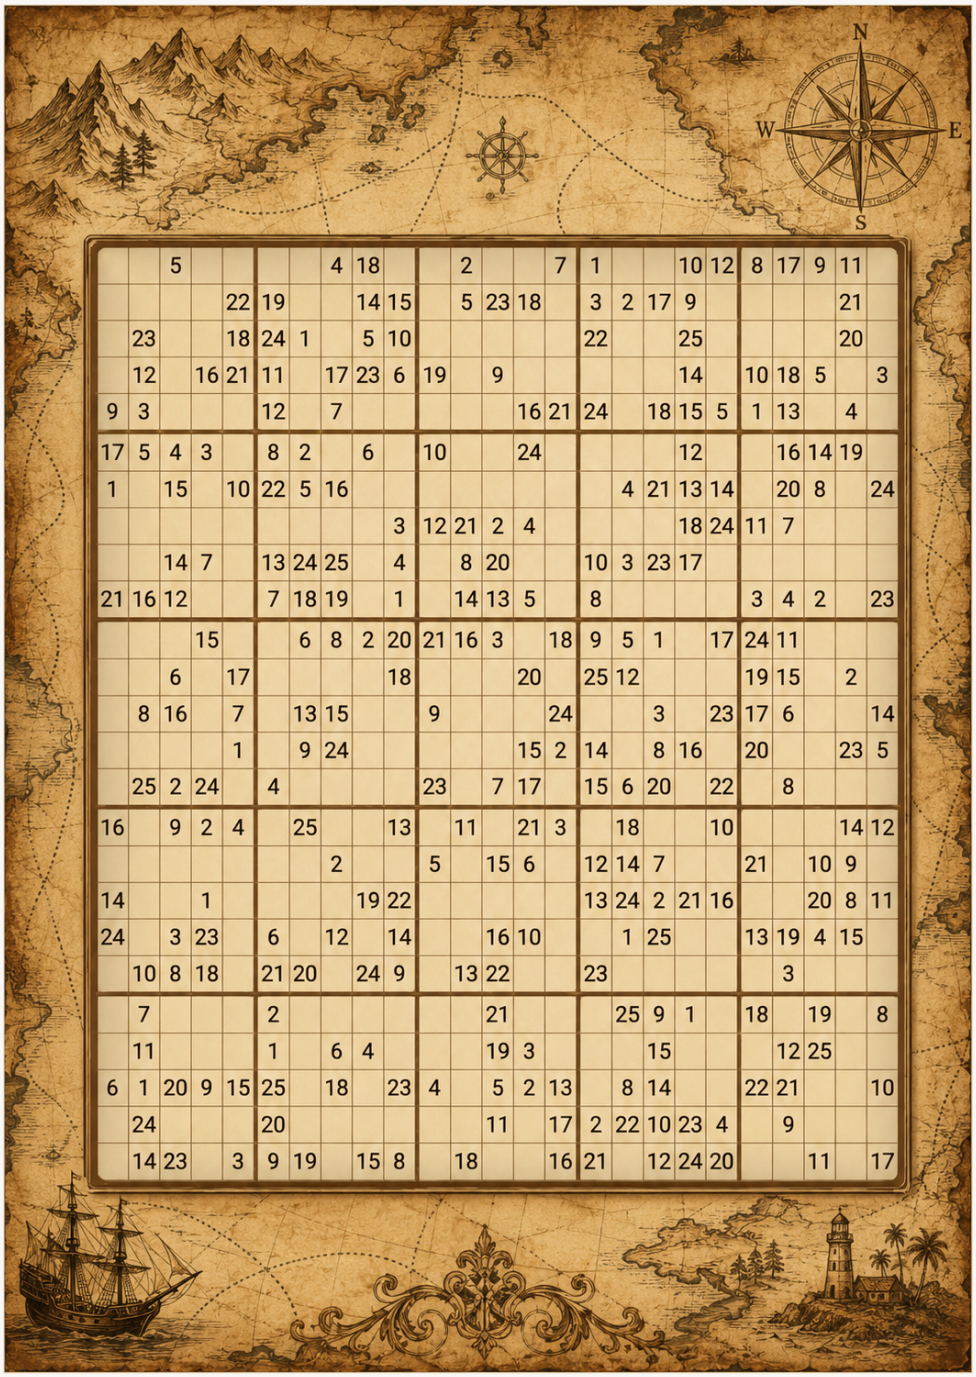

		5					4	18			2			7	1			10	12	8	17	9	11	
				22	19			14	15		5	23	18		3	2	17	9					21	
	23			18	24	1		5	10						22			25					20	
	12		16	21	11		17	23	6	19		9						14		10	18	5		3
9	3				12		7						16	21	24		18	15	5	1	13		4	
17	5	4	3		8	2		6		10			24					12			16	14	19	
1		15		10	22	5	16									4	21	13	14		20	8		24
									3	12	21	2	4					18	24	11	7			
		14	7		13	24	25		4		8	20			10	3	23	17						
21	16	12			7	18	19		1		14	13	5		8					3	4	2		23
			15			6	8	2	20	21	16	3		18	9	5	1		17	24	11			
		6		17					18				20		25	12				19	15		2	
	8	16		7		13	15			9				24			3		23	17	6			14
				1		9	24						15	2	14		8	16		20			23	5
	25	2	24		4					23		7	17		15	6	20		22		8			
16		9	2	4		25			13		11		21	3		18			10				14	12
							2			5		15	6		12	14	7			21		10	9	
14			1					19	22						13	24	2	21	16			20	8	11
24		3	23		6		12		14			16	10			1	25			13	19	4	15	
	10	8	18		21	20		24	9		13	22			23						3			
	7				2							21				25	9	1		18		19		8
	11				1		6	4				19	3				15				12	25		
6	1	20	9	15	25		18		23	4		5	2	13		8	14			22	21			10
	24				20							11		17	2	22	10	23	4		9			
	14	23		3	9	19		15	8		18			16	21		12	24	20			11		17

25x25 Sudoku kan: a lè check òfin rẹ lái reveal finished grid — visual stand-in fún proof kan tó verify hidden solution, witness náà.

AI-generated editorial thumbnail — The Clean Paper CC BY 4.0

Sudoku àti SAT: puzzle kan náà nínú aṣọ méjì

Claim pé Sudoku kan lè “behave like a SAT formula” kì í ṣe metaphor. Translation náà ń lọ ní ìtọ́ṣọ́nà méjèèjì, easy ìtọ́ṣọ́nà sì lè kọ sílẹ̀ patapata.

Láti Sudoku sí SAT. SAT sọ true/false nìkan, torí náà fún un ní boolean variable kan fún every (row, column, value) triple: *x(r,c,v)* túmọ̀ sí “ṣẹ̀ṣẹ̀li ní row *r*, column *c* contains value *v*.” 4-by-4 Sudoku (2-by-2

blocks, values 1–4) nílò 4·4·4 = 64 variables; classical 9-by-9 nílò 729. Every Sudoku òfin náà di batch of clauses. (Clause jẹ́ OR of variables tàbí negations wọn; whole formula ni AND of gbogbo clauses.)
Every sẹẹli holds at least ọ̀kan value — ọ̀kan clause per sẹẹli:

$$x(1,1,1) \vee x(1,1,2) \vee x(1,1,3) \vee x(1,1,4)$$

Every sẹẹli holds at most ọ̀kan value — "not both" clause kan fún each pair of values:

$$\neg x(1,1,1) \vee \neg x(1,1,2) \quad \neg x(1,1,1) \vee \neg x(1,1,3) \quad \ldots$$ àti bẹ́ẹ̀ lọ fún gbogbo six pairs.

Every row contains every value — fún row 1 àti value 3: at least once,

$$x(1,1,3) \vee x(1,2,3) \vee x(1,3,3) \vee x(1,4,3)$$

àti at most once: $\neg x(1,1,3) \vee \neg x(1,2,3)$, àti bẹ́ẹ̀ lọ fún each pair of sẹẹli nínú row.
Columns àti blocks — identical batches; ẹgbẹ́ of sẹẹli nìkan ló yàtọ̀. Fún top-left block àti value 2:

$$x(1,1,2) \vee x(1,2,2) \vee x(2,1,2) \vee x(2,2,2)$$

plus pairwise "not both" clauses.
Printed clues — easiest part: clue kọọkan jẹ́ clause pẹ̀lú single variable. Printed 3 ní top-left corner di clause

$$x(1,1,3)$$

AND of gbogbo èyí satisfiable exactly when Sudoku ní solution — satisfying assignment náà sì *jẹ́* solution: ka x(r,c,v) wo ni true, kí o fill grid. Fún 9-by-9, èyí jẹ́ variables 729 àti a few thousand clauses, tí modern SAT solver lè dispatch ní milliseconds. Ṣe akiyesi clue clause x(1,1,3): ó sọ "sẹẹli yìí equals exactly 3," kì í ṣe "sẹẹli wòyí are gbogbo yàtọ̀" — asymmetry kan náà tó máa force extra trick fún clue sẹẹli nínú protocol note tó wà ní isalẹ.
Láti SAT sí Sudoku. ìwé ìwádìí nílò opposite, harder ìtọ́sọ́nà: fún *arbitrary* SAT formula, build mega-Sudoku tó ní solution exactly when formula náà does. Native òfin Sudoku lè sọ "sẹẹli wòyí gbogbo yàtọ̀" nìkan, torí náà arbitrary logical constraints gbọ́dọ̀ *built* — èyí gan-an ni gadgets. Gadget jẹ́ small pre-fabricated cluster of sẹẹli, ọ̀kan per clause of formula, níbi tí designated sẹẹli ṣe role of variables (symbol tí wọ́n hold encode true tàbí false), internal constraints cluster náà sì engineered kí nìkan legal fillings correspond sí assignments tí satisfy clause. Èyí jẹ́ standard craftsmanship láti NP-completeness proofs; fún generalized Sudoku, Yato àti Seta ṣe é ní 2003.
Together, ìtọ́sọ́nà méjèèjì sọ pé N-by-N Sudoku àti SAT jẹ́ ìṣòro kan náà tí wọ́n wọ aṣọ yàtọ̀. Èyí ni ohun tó license article yìí — àti ìwé ìwádìí náà — láti tell ìtàn nípa gbogbo NP pẹ̀lú grids àti symbols.

Witness náà ṣì rọrùn láti picture. Alice mọ complete valid filling ti mega-Sudoku. Bob fẹ́ convinced pé filling bẹ́ẹ̀ wà, ṣùgbọ́n Alice kò fẹ́ reveal rẹ̀. Bí ó bá send whole filling, Bob convinced, ṣùgbọ́n secret ti gone.

Nínú classical zero-knowledge ẹ̀dà, Alice àti Bob interact. Old-style mental àwòṣe kan lo covered tiles. Alice hide solved grid, secretly rename symbols ṣáájú round kọọkan, ó sì jẹ́ kí Bob inspect ọ̀kan randomly chosen agbègbè constraint: row, column, box, tàbí gadget. Bí opened sẹẹli bá fi hàn all-different symbols, Bob gain ìgbẹ́kẹ̀lé. Lẹ́yìn náà gbogbo nǹkan ni a cover lẹ́ẹ̀kan síi, symbols sì freshly renamed. (Wrinkle kan: given clues puzzle náà nílò extra trick, torí renaming symbols hide wọn náà. Note tó wà ní isalẹ explain bí classical protocols ṣe solve èyí; toy

picture náà tó fún ohun tó tẹ̀lé.)

> **Bí classical protocols ṣe handle clue cells gidi**
>
> Renaming trick ní blind spot. Row, column àti box òfin gbogbo sọ “these sẹẹli are gbogbo yàtọ̀,” *gbogbo yàtọ̀* sì survive renaming of symbols èyíkéyìí. Ṣùgbọ́n clue kan sọ “sẹẹli yìí contains exactly 5,” lẹ́yìn renaming Bob rí $\sigma(5)$ nìkan — masked symbol kan — láì mọ renaming σ. Kò lè check ohunkóhun. Bí a kò bá fix èyí, Alice lè fi ẹ̀rí múlẹ̀ pé *some* valid grid exists nígbà tó ignore printed clues patapata, èyí kò fi ẹ̀rí múlẹ̀ ohunkóhun nípa *this* puzzle. Classical literature ní standard repairs méjì.
> **Palette náà.** Fi extra row kan ti N sẹẹli kun hidden grid — palette tí Alice fill pẹ̀lú symbols 1…N ní fixed public order, lẹ́yìn náà ó rename rẹ̀ pẹ̀lú everything else, torí náà ó contains $\sigma(1)\ldots\sigma(N)$. Random challenge Bob ní extra option kan báyìí. Yàtọ̀ sí picking row, column, box tàbí gadget láti open, ó lè pick **palette plus ọ̀kan clue sẹẹli**. Alice uncover méjèèjì; palette reveal renaming round yẹn, Bob sì check pé clue sẹẹli fi hàn exactly renamed ẹ̀dà of printed clue. Èyí stay zero-knowledge torí Bob learn σ nìkan — freshly drawn every round, useless on its own — àti value of sẹẹli tí ó ti mọ láti puzzle. Kò sí secret sẹẹli information tó leak, simulator sì lè fake view nípa drawing random σ. Ó sound torí cheating Alice máa caught pẹ̀lú fixed probability per round, rounds sì repeated títí doubt fi negligible.
> **Compiling clues away.** Structural variant mîì remove special challenge dípò kó add rẹ̀. Dípò *verifying* clue value, force rẹ̀ pẹ̀lú difference constraints: link clue sẹẹli sí every palette sẹẹli except ọ̀kan carrying own value — “yàtọ̀ láti $\sigma(1)$, yàtọ̀ láti $\sigma(2)$, …, yàtọ̀ láti everything ṣùgbọ́n $\sigma(5)$.” Only symbol sẹẹli náà lè legally hold ni clue náà. Every constraint báyìí jẹ́ “these méjì differ” lẹ́ẹ̀kan síi — invariant lábẹ́ renaming, checkable exactly like row. Maneuver kan náà ni a lo fún pre-colored vertices nínú classical graph-coloring protocol, ó sì jẹ́ spirit of word *gadgets* lókè: nínú MegaSudoku-as-SAT picture, clues are compiled sínú inequality gadgets bí every other constraint.
> **Physical protocol náà.** ayé gidi card protocol fún Sudoku (Gradwohl, Naor, Pinkas àti Rothblum, 2007) kò lo renaming rárá, ó sì settle clues ṣáájú kí hiding tó bẹ̀rẹ̀. Fún every sẹẹli, Alice lay down identical cards méta pẹ̀lú value sẹẹli náà — face-down fún secret sẹẹli, ṣùgbọ́n **face-up fún clue sẹẹli**, kí Bob fi ojú ara rẹ̀ rí pé clues respected ṣáájú kí cards flipped. Léyìn náà card kan láti each sẹẹli wọ packet row rẹ̀, ọ̀kan sínú column, ọ̀kan sínú box; packet kọ̀ọ̀kan shuffled àti revealed, Bob sì check pé ó contains gbogbo N symbols. Shuffling destroy position information (ìyẹn ni zero-knowledge), ṣùgbọ́n clues ti nailed down ní dealing àkókò.
> Either way, lesson náà jẹ́ ohun kan náà tí article yìí ń padà sí: zero-knowledge protocol jẹ́ careful bookkeeping ti *which* facts survive hiding. Renaming preserve “gbogbo yàtọ̀” ó sì erase “equals 5” — torí náà “equals 5” gbọ́dọ̀ smuggled back in ní other means.

Èyí kì í ṣe protocol inú ìwé ìwádìí. Ó jẹ́ mental àwòṣe fún classical zero-knowledge:

- Alice àti Bob go back àti forth.
- Bob choose random checks.
- Alice reveal agbègbè consistency nìkan, kì í ṣe whole solution.
- Privacy proof ṣiṣẹ́ nípa showing pé Bob’s view lè have been generated láìsí secret solution Alice.

Torí náà classical zero-knowledge built around positive fact kan:

> Simulator kan wà gidi.

Báyìí remove comfortable parts. Alice send proof string kan, ó sì walk away. Kò sí trusted setup, kò sí shared random string prepared in advance, Bob sì gbọ́dọ̀ never accept false puzzle. Èyí ni setting tí classical zero-knowledge kò lè survive.

Character kan síi nílò ṣáájú trick nàà. Fix **rulebook** kan: formal proof ètò, ní sense ti logician — fixed set of axioms plus mechanical òfin fún checking written mathematical proofs. ZFC, standard axioms of mathematics, ni canonical example. Gbogbo nǹkan láti ibí lọ ni stated relative sí rulebook tí a choose in advance, choice sì flexible: construction ṣíṣẹ́ fún rulebook èyíkéyìí tí o fix, ZFC included.

(Note kan lórí words, borrowed láti ìwé ìwádìí fúnra rẹ̀: “proof ètò” níbí nigbagbogbo túmọ̀ sí rulebook yìí — formal ètò tó check mathematical proofs — kò túmọ̀ sí messages Alice send. Machinery Alice àti Bob ni a pe ní “prover àti verifier.”)

Gödel-style ẹ̀dà keep mega-Sudoku ìtàn, ṣùgbọ́n ó change proof nàà.

Yan second constraint ètò kan ti displayed ìwọ̀n kan nàà, pè é ní **D**. Fún ìtàn nàà, S àti D jẹ́ MegaSudoku(n) puzzles méjì ní format kan nàà. Behind scenes, D lè have started as hard logical formula of yàtọ̀ ìwọ̀n; bí needed, a lè pad rẹ̀ pẹ̀lú harmless dummy constraints kí ó fit grid kan nàà. D built láti logical formula tó jẹ́ **unsatisfiable** gidi: kò sí ṣeé ṣe assignment of values tó make gbogbo constraints true, bí broken puzzle tí kò ní legal completed grid. Toy example ni formula tó demand “X is true” àti “X is false” méjèèjì. Torí nàà D kò ní valid filling.

Ṣùgbọ́n D kò gbọ́dọ̀ jẹ́ broken puzzle tí ó *easy to expose*. Toy example lókè fail: rulebook èyíkéyìí refute “X àti not-X” ní line kan. D gbọ́dọ̀ false ní way tí chosen rulebook kò lè certify pẹ̀lú short argument. Bí rulebook bá lè refute D pẹ̀lú short proof, ìtàn tó wà ní isalẹ collapse: alternative route tó lè have produced proofs láìsí secret Alice lè formally ruled out, àṣírí guarantee nàà sì parí. Torí nàà a yan D láti family tí fixed rulebook kò lè efficiently refute: kò sí short proof, nínú rulebook yẹn, pé D has no solution.

One-message proof Alice jẹ́ nípa either/tàbí statement:

> either gidi mega-Sudoku S has a solution, tàbí decoy D has a solution.

Èyí ni logical link. D **kò** generated ní magical way kan tó make S true. Proof nàà kò argue “D has no solution, therefore S has a solution.” Ó ń fi ẹ̀rí múlẹ̀ disjunction **S tàbí D**. Perfect soundness sọ pé false disjunction kò lè ní valid proof. Níwọ̀n bí D ti false gidi — kò ní solution — nìkan way disjunction lè true ni pé S true. Torí nàà bí proof accepted, S gbọ́dọ̀ ní solution. Decoy kò lè make false S true.

Ṣùgbọ́n fún zero-knowledge-style part, béèrè ohun tó máa ṣẹlẹ̀ bí D *bá ní* solution. Decoy solution yẹn máa act as alternative witness. Ó máa jẹ́ kí ẹnikan produce proofs láì mọ gidi mega-Sudoku solution Alice — simulator, ní other words. Ní reality D kò ní solution, torí nàà simulator route yìí closed. Point nàà ni pé rulebook kò lè efficiently fi ẹ̀rí múlẹ̀ pé route nàà closed.

Torí nàà D ní jobs méjì. Fún **soundness**, D false, torí valid proof of “S tàbí D” force S. Fún **effective zero-knowledge**, D hard to refute, torí rulebook kò lè quickly òfin out decoy route tó bá ti make simulation ṣeé ṣe.

Torí nàà security ìdánwò kò longer jẹ́:

> Ṣé a lè fi ẹ̀rí múlẹ̀ pé simulator kan wà gidi?

Ó di:

> Ṣé rulebook rẹ lè efficiently fi ẹ̀rí múlẹ̀ pé simulator impossible?

Bí answer bá jẹ́ no, nǹkan surprisingly lágbára follow: every security guarantee tí (a) a lè observe nípa running ìdánwò, àti (b) provably follows — inside rulebook yẹn — láti existence of simulator, actually holds. Successful ìkọlù lórí eyikeyi wọn yóò fúnra rẹ̀ amount sí missing short refutation, missing short refutation yẹn kò sì exist. Ìyẹn ni “effective” part of effectively zero-knowledge.

Torí náà classroom contrast ni:

Classical zero-knowledge: proofs láìléwu torí simulator exists.

Gödel-style effective zero-knowledge: proofs treated as láìléwu fún observable security ìdánwò torí rulebook kò lè efficiently fi ẹ̀rí múlẹ̀ pé simulator impossible.

Claim kejì weaker. Ó tún jẹ́ reason tí ìwé ìwádìí lè keep àbùdá méta tí broke classical ẹ̀dà: ọ̀kan message, no setup àti perfect soundness.

ìdánwò tuntun: o kò lè fi ẹ̀rí múlẹ̀ pé simulator absent

Relaxation Ilango change ìbéèrè náà.

Classical zero-knowledge asks:

> Does a simulator exist?

Effectively zero-knowledge asks something weaker:

> Can your chosen rulebook efficiently fi ẹ̀rí múlẹ̀ that no simulator exists?

Èyí sound bí technical dodge, ṣùgbọ́n core idea ni. Construction náà wà nínú strange ipò: simulator gidi kò exist — ìwé ìwádìí explicit nípa èyí — ṣùgbọ́n rulebook tí o fix kò lè efficiently fi ẹ̀rí múlẹ̀ pé kò exist. Bí every bad consequence tí o care nípa bá require such refutation, èyí náà ṣì behave like zero-knowledge fún consequences wọ̀nyẹn.

Ibí ni Gödel enter. Kì í ṣe decoration, kì í sì ṣe “Gödel makes crypto secure.” Connection náà proof-theoretic. Rulebook kan ni a pe ní **optimal** bí, ní precise sense, ó jẹ́ best ṣeé ṣe ọ̀kan: whenever any rulebook lè refute formula of relevant kind pẹ̀lú short proof, optimal rulebook náà lè ṣe bẹ́ẹ̀ náà, pẹ̀lú proof at most polynomially longer. Krajíček àti Pudlák conjectured ní 1989 pé **no optimal proof èyí exists**: whichever rulebook o fix, some other rulebook proves some family of true statements far sí i succinctly. Èyí jẹ́ ọ̀kan of central open conjectures of proof complexity, ó sì jẹ́ finite, complexity-theoretic cousin of Gödel’s incompleteness theorem: some true statements have no short proof nínú rulebook tí o fix — kì í ṣe torí wọ́n unprovable in principle, ṣùgbọ́n torí every fixed rulebook leave some short truths láìsí short proofs.

ìwé ìwádìí assume conjecture yìí (ní mildly stronger “infinitely often” form, standard when conjectures lò cryptographically). Payoff, nípa theorem of Krajíček àti Pudlák, concrete ni: fún every rulebook sequence of formulas kan wà tí genuinely unsatisfiable, tí rulebook kò lè refute pẹ̀lú short proofs — àti, crucially, efficient algorithm lè *generate* wọn. Last property yẹn, uniformity, ni ohun tó turn whole idea láti existence claim sí actual algorithm tí Alice lè run: decoys D rẹ̀ wá láti assembly line, kì í ṣe out of thin air.

Cryptographic move náà ni láti put shortage of proof power yẹn to iṣẹ́.

Ohun tí construction náà ń ṣe

Ẹ jẹ́ ká strip construction ìwé ìwádìí sí shape rẹ̀.

Fix rulebook kan — ZFC, say. Under proof-complexity assumption, sequence of formulas kan wà tí a lè efficiently generate, tí actually unsatisfiable, ṣùgbọ́n rulebook kò ní short proof pé wọ́n unsatisfiable.

Báyìí build one-message proof in form yìí:

> either gidi statement is satisfiable, tàbí this special hard formula is satisfiable.

Special hard formula náà kò satisfiable. Torí náà bí underlying proof machinery perfectly sound, accepting message náà ṣì mean gidi statement true. Èyí give perfect soundness.

Ṣùgbọ́n fún zero-knowledge-like security, imagine special hard formula *were* satisfiable. Witness rẹ̀ lè then be lò to simulate proofs láì mọ gidi witness. Formula náà kò satisfiable in reality — ṣùgbọ́n rulebook kò lè efficiently fi ẹ̀rí múlẹ̀ bẹ́ẹ̀. Torí náà kò lè efficiently fi ẹ̀rí múlẹ̀ pé simulator impossible.

Ìyẹn ni hinge. System kò hide secret nípa producing classical simulator. Ó hide secret, fún large class of observable security ìdánwò, behind rulebook’s inability to certify pé simulator absent.

Ohun tí ìwé ìwádìí claim

Main theorem wá ní layers. Core èsì ni èyí:

Under standard cryptographic assumption — existence of **non-interactive witness indistinguishable proofs**, well-studied objects tí follow láti several established assumption packages — àti lábẹ́ proof-complexity conjecture pé **no (infinitely often) optimal proof ètò exists**, ìwé ìwádìí construct, fún every choice of rulebook, one-message prover àti verifier fún NP/SAT pẹ̀lú **perfect soundness** àti no setup tí effectively zero-knowledge relative sí rulebook yẹn. (NP/SAT ni standard “hardest common denominator” ti puzzle-like ìṣòro; mega-Sudoku jẹ́ costume kan tí ó wọ̀.)

Fún broader claim nípa preserving falsifiable security properties, ìwé ìwádìí add ọ̀kan sí i standard assumption, derandomization belief **P = BPP** (roughly: randomness kò fún algorithms ní essential extra power).

Ní èdè tí kì í ṣe theorem:

- Proof náà jẹ́ ọ̀kan message.
- Kò sí trusted setup.
- False statements kò lè be fi ẹ̀rí múlẹ̀.
- Prover náà kì í ṣe classical zero-knowledge — kò ní simulator.
- Ṣùgbọ́n every falsifiable, game-based security consequence of classical zero-knowledge lè achieved nínú setting yìí.

"Falsifiable" ṣe pàtàkì. Ó túmọ̀ sí pé security ìkùnà lè ìdánwò nípa running adversary nínú game. Ọ̀pọ̀ cryptographic security definitions ní form yìí: lè adversary distinguish méjì encryptions, invert a function, recover a witness, tàbí win specified ìdánwò? Theorem náà give prover kan fún each falsifiable property, ọ̀kan at a àkókò. Single prover tó enjoy *every* falsifiable property at once likely impossible — old reusability ìkọ̀lù ("Bob lè fi hàn the proof to others") fúnra rẹ̀ jẹ́ falsifiable property, ó sì genuinely fail níbí. Proposal ìwé ìwádìí ni pé single prover kan plausibly lè cover gbogbo *natural* falsifiable properties — ones tí actually occur nínú cryptographic practice — ṣùgbọ́n apá yẹn jẹ́ conditional theorem tó rest lórí informal notion of "natural," plus explicit conjecture. Guarantee aimed at observable ìkùnà, kì í ṣe every philosophical tàbí simulation-based meaning of secrecy.

Concrete corollary kan tọ́ láti name: construction yields first non-interactive *witness hiding* proofs pẹ̀lú uniform prover — "proof of a puzzle does not help you rí its solution," pẹ̀lú no interaction àti no setup — object tó sound modest ṣùgbọ́n tí construction rẹ̀ resisted fún decades.

Ohun tí èyí kò sọ

Section yìí ni tó keep piece honest.

Kò sọ pé old impossibility theorems wrong. Construction avoids wọn nípa changing definition.

Kò give ordinary, classical zero-knowledge pẹ̀lú no interaction, no setup àti perfect soundness. ìwé ìwádìí explicitly sọ pé constructed prover kò ní simulator.

Kò mean proof kò lè be reused. One-message proof ṣì lè fi hàn fún ẹlòmíràn; ìwé ìwádìí kò preserve deniability-style properties. (Non-interactive zero-knowledge pẹ̀lú trusted setup ní limitation kan náà.)

Kò mean pé practical protocol ready fún deployment ni. Èyí jẹ́ complexity àbá-ẹ̀kọ́ àti cryptographic foundations. èsì depend lórí major assumptions láti proof complexity àti cryptography, construction náà sì nípa ohun tó ṣeé ṣe in principle.

Kò make "Gödel" magic security primitive. Gödel connection wá nípasẹ̀ proof ètò, optimal proof ètò àti finite analogues of incompleteness. Useful intuition kì í ṣe "incompleteness protects your password." Ó jẹ́: bí rulebook kò lè efficiently fi ẹ̀rí múlẹ̀ pé simulator impossible, ìkọ̀lù tí yóò require proof yẹn lè blocked at ìpele of security definitions.

Kí ló dé tí ó fi interesting síbẹ̀?

Cryptography sábà turn hardness sínú ààbò. Factoring hard, torí RSA-style assumptions useful. Lattice ìṣòro hard, torí lattice cryptography useful. Ní here hardness stranger: kì í ṣe “hard to compute a secret,” ṣùgbọ́n “hard to fi ẹ̀rí múlẹ̀ that a certain proof object kò lè exist.”

Ìdí nìyẹn tí ìwé ìwádìí fi feel unusual. Ó treat axioms àti rulebooks almost like cryptographic resources. Usual impossibility ṣọ pé tension wà láàárín soundness àti simulation. Move Ilango ni láti place tension náà behind proof-theoretic curtain: simulator absent, ṣùgbọ́n formal ètò kò lè efficiently expose absence náà.

Fún reader, surprising part kì í ṣe pé èyí máa replace today’s zero-knowledge ètò. Probably kò ní, at least not taara. Surprising part ni pé limitation láti mathematical logic lè be lò constructively: kì í ṣe wall nìkan, ṣùgbọ́n kind of cover.

Báwo ni ẹ̀rí ṣe lágbára tó?

Èyí jẹ́ theorem ìwé ìwádìí, torí náà “ẹ̀rí” túmọ̀ sí nǹkan yàtọ̀ sí biology tàbí astronomy ìwé ìwádìí. ìbéèrè kì í ṣe bóyá ìdánwò replicated. ìbéèrè ni bóyá definitions, assumptions àti proof chain ṣe àtìlẹ́yìn fún claim náà.

Proof náà formal, ìwé ìwádìí sì explicit nípa assumptions rẹ̀. Assumptions náà kì í casual. Non-interactive witness indistinguishable proofs jẹ́ standard objects nínú cryptography, wọ́n sì follow láti several established assumption packages. No-optimal-proof-system conjecture jẹ́ central conjecture nínú proof complexity. P = BPP jẹ́ standard derandomization belief tí a lo nìkan fún broader falsifiable-property theorem.

ìwé ìwádìí tún argue pé assumptions náà ni right price, kì í ṣe arbitrary scaffold: ó fi ẹ̀rí múlẹ̀ converse kan showing pé wọ́n essentially necessary — bí constructions bí èyí bá exist rárá, non-interactive witness indistinguishable proofs gbọ́dọ̀ exist, àti (granting standard one-way functions) no optimal proof ètò lè exist. Assumptions náà sì “win-win”: refuting any of them fúnra rẹ̀ yóò jẹ́ landmark discovery nínú proof complexity, cryptography tàbí complexity àbá-ẹ̀kọ́.

Ṣùgbọ́n torí èsì conditional, ìgbẹ́kẹ̀lé rẹ̀ conditional náà. Bí assumptions yẹn bá fail, interpretation theorem change. Bí wọ́n bá hold pàápàá, guarantee náà kì í ṣe full classical zero-knowledge; ó jẹ́ relaxed, proof-theoretic ẹ̀dà ìwé ìwádìí.

Torí náà right ìgbẹ́kẹ̀lé ni high pé ìwé ìwádìí establish coherent conditional possibility èsì; moderate pé assumptions rẹ̀ describe cryptographic world tí a ń gbé gidi; low fún any immediate practical consequence.

Kí nìdí tí ó fi ṣe pàtàkì?

ìwé ìwádìí náà open route kan tí a ro pé closed.

Classical àbá-ẹ̀kọ́ ṣọ pé: full zero-knowledge kò lè be ọ̀kan message láìsí setup, kò sì lè perfectly sound. ìwé ìwádìí Ilango ṣọ pé: bí a bá ask fún consequences of zero-knowledge tí a lè ìdánwò nínú security games, àti bí a bá allow

security definition to depend lórí ohun tí rulebook lè tàbí kò lè efficiently refute, nígbà náà púpọ̀ useful ìhùwàsí lè recovered — pẹ̀lú ọ̀kan message, no setup àti perfect soundness.

Èyí kì í ṣe small definitional tweak. Ó jẹ́ way yàtọ̀ láti think nípa cryptographic guarantees. Dípò asking nìkan what exists, ask what rulebook rẹ lè òfin out. Dípò treating unprovability bí philosophical nuisance, lò rẹ̀ as ìṣètò.

Practical world lè má change lọ́la. Ṣùgbọ́n conceptual map change. Formal sense kan wà báyìí níbi tí “no ọ̀kan lè efficiently fi ẹ̀rí múlẹ̀ that the secret leaked” lè lágbára tó láti recover many game-based protections tí a fẹ́ láti “the secret did not leak.”

Ìdí nìyẹn tí Gödel fi belong nínú title.

Àkótán kedere

Zero-knowledge proofs jẹ́ kí prover convince verifier pé statement true láì reveal witness. Classical impossibility èsì sọ pé zero-knowledge kò lè squeezed sí ọ̀kan message láìsí setup, kò sì lè ní perfect soundness. ìwé ìwádìí Rahul Ilango kò refute impossibilities yẹn. Ó define weaker notion, effectively zero-knowledge: dípò requiring pé simulator really exists, ó require pé chosen proof ètò — formal rulebook bí ZFC — kò lè efficiently fi ẹ̀rí múlẹ̀ pé no simulator exists. Under major assumptions láti cryptography (non-interactive witness indistinguishable proofs) àti proof complexity (no optimal proof ètò exists), ìwé ìwádìí construct one-message provers fún NP/SAT pẹ̀lú no setup àti perfect soundness tí achieve falsifiable, game-based consequences of zero-knowledge property by property. Single prover tó cover gbogbo “natural” properties bẹ́ẹ̀ jẹ́ further, partly conjectural extension — covering literally every falsifiable property likely impossible, torí proofs remain reusable. èsì náà theoretical àti conditional, kì í ṣe deployed primitive, ṣùgbọ́n ó fi hàn tuntun way láti lò proof-theoretic unprovability as cryptographic resource.

Àyẹ̀wò láìsí àṣejù

Ohun tí ìwé ìwádìí fi hàn: Under stated assumptions, a lè build one-message, no-setup, perfectly sound provers fún NP/SAT tí effectively zero-knowledge relative sí any chosen proof ètò, tí wọ́n sì achieve each falsifiable game-based consequence of classical zero-knowledge.

Ohun tó ṣeé gbà ṣùgbọ́n tí a kò fi ẹ̀rí múlẹ̀ unconditionally: Pé needed proof-complexity àti cryptographic assumptions hold. Wọ́n serious, well-studied assumptions — ìwé ìwádìí sì fi hàn pé wọ́n essentially necessary as well as sufficient — ṣùgbọ́n assumptions ni wọ́n ṣì jẹ́.

Ohun tí kò fi hàn: Classical zero-knowledge pẹ̀lú no interaction, no setup àti perfect soundness; practical ètò ready fún deployment; deniability tàbí non-reusability of proofs; tàbí pé Gödel’s incompleteness theorem fúnra rẹ̀ secure cryptography.

Main limitations: Guarantee náà jẹ́ relaxation of zero-knowledge; broadest ẹ̀dà depend lórí multiple assumptions; single-universal-prover claims remain partly conjectural; èsì náà sì primarily foundational.

Confidence wo ni gbogbogbò reader yẹ kí ó ní? High pé èyí jẹ́ important conditional àbá-ẹ̀kọ́ èsì bí a bá accept definitions. Moderate pé assumptions capture reality. Low fún immediate practical deployment. Safe takeaway ni: ìwé ìwádìí náà kò break zero-knowledge impossibilities; ó rí tuntun proof-theoretic way around parts of them tí matter fún many security games.

Àwọn orísun

Ilango, IACR ePrint 2025/1296
`https://eprint.iacr.org/2025/1296`

FOCS 2025, DOI 10.1109/FOCS63196.2025.00058
`https://doi.org/10.1109/FOCS63196.2025.00058`