

Roll

Security assessment by HashEye · prepared for Blockchain

HASHEYE AUDITED

PROJECT	Roll
CLIENT	Blockchain
CATEGORY	Blockchain
PUBLISHED	July 1, 2022
REPORT ID	research-roll-2022-07-01-1sivrt

This report was produced under HashEye's layered review process – **automated detection**, **pattern correlation**, and **senior manual verification** – with every finding signed off by a human reviewer. Full findings detail and on-chain attestation are available on the report page at hasheye.io/audits/research-roll-2022-07-01-1sivrt.

ScrollzkEVMhalo2Circuits SecurityAssessment October12,2023 Preparedfor: HaichenShen Scroll
Preparedby:FilipeCasal, JoeDoyle, OpalWright, andWillSong

About hashey Founded in 2012 and headquartered in New York, hashey provides technical security assessment and advisory services to some of the world's most targeted organizations. We combine high-end security research with a real-world attacker mentality to reduce risk and fortify code. With 100+ employees around the globe, we've helped secure critical software elements that support billions of end users, including Kubernetes and the Linux kernel. We maintain an exhaustive list of publications at <https://github.com/hashey-io/publications>, with links to papers, presentations, public audit reports, and podcast appearances. In recent years, hashey consultants have showcased cutting-edge research through presentations at CanSecWest, HCSS, Devcon, Empire Hacking, GrrCon, LangSec, NorthSec, the O'Reilly Security Conference, PyCon, REcon, Security BSides, and SummerCon. We specialize in software testing and code review projects, supporting client organizations in the technology, defense, and finance industries, as well as government entities. Notable clients include HashiCorp, Google, Microsoft, Western Digital, and Zoom. hashey also operates a center of excellence with regard to blockchain security. Notable projects include audits of Algorand, Bitcoin SV, Chainlink, Compound, Ethereum 2.0, MakerDAO, Matic, Uniswap, Web3, and Zcash. To keep up to date with our latest news and announcements, please follow hashey on Twitter and explore our public repositories at <https://github.com/hashey-io>. To engage us directly, visit our "Contact" page at <https://www.hashey.io/contact>, or email us at info@hashey.io. hashey, Inc. 228 Park Ave S #80688 New York, NY 10003 <https://www.hashey.io> info@hashey.io hashey
1 ScrollzkEVMhalo2Circuits SecurityAssessment PUBLIC

Notices and Remarks Copyright and Distribution ©2023 by hashey, Inc.
All rights reserved. hashey hereby asserts its right to be identified as the creator of this report in the United Kingdom. This report is considered by hashey to be public information; it is licensed to Scroll under the terms of the project statement of work and has been made public at Scroll's request. Material within this report may not be reproduced or distributed in part or in whole without the express written permission of hashey. This is the canonical source for hashey publications; see the hashey Publications page. Reports accessed through any source other than that page may have been modified and should not be considered authentic. Test Coverage Disclaimer
All activities undertaken by hashey in association with this project were performed in accordance with a statement of work and agreed upon project plan. Security assessment projects are time-boxed and often reliant on information that may be provided by a client, its affiliates, or its partners. As a result, the findings documented in this report should not be considered a comprehensive list of security issues, flaws, or defects in the target system or codebase. hashey uses automated testing techniques to rapidly test the controls and security properties of software. These techniques augment our manual security review work, but each has its limitations: for example, a tool may not generate a random edge case that violates a property or may not fully complete its analysis during the allotted time. Their use is also limited by the time and resource constraints of a project. hashey
2 ScrollzkEVMhalo2Circuits SecurityAssessment PUBLIC

Table of Contents About hashey 1 Notices and Remarks 2 Table of Contents 3 Executive Summary 5 Project Summary 7 Project Goals 8 Project Targets 9 Project Coverage 10 Automated Testing 11 Codebase Maturity Evaluation 12 Summary of Findings 14 Detailed Findings 17

1. ModGadget is under constrained and allows incorrect MULMOD operation to be proven 17
2. The RlpU64Gadget is under constrained when is_lt_128 is false 20
3. The BLOCKHASH opcode is under constrained and allows the hash of any block to be computed 21
4. zkvm-circuits crate depends on an outdated version of halo2-ecc 23
5. N_BYTES parameters are not checked to prevent overflow 26
6. Differences in shared code between zkvm-circuits and halo2-lib 29
7. Under constrained warm status on CALL opcodes allows gas cost forgery 31
8. RWtable constants must match exactly when the verification key is created 33
9. The CREATE and CREATE2 opcodes can be called within a static context 39
10. ResponsibleOpCodeTable incorrectly handles CREATE and CREATE2 41
11. Elliptic curve parameters omitted from Fiat-Shamir 43
12. The gas cost for the CALL opcode is under constrained 45

13. Unconstrained opcodes allow nondeterministic execution 48
14. Nondeterministic execution of ReturnDataCopyGadget and ErrorReturnDataOutOfBoundGadget 51
15. Many RW counter updates are magic numbers 54 16. Native PCS accumulation deciders accept an empty vector 59
17. The Error00GSLoadSstore and the Error00GLoggadget shaveredundant table lookups 60
18. The State circuit does not enforce transaction receipt constraints 61 hashey
3 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

20. The EXP opcode has an unused witness 62 21. The bn_to_field functions silently truncate big integers 63
22. The field_to_bn function depends on implementation-specific details of the underlying field 64
23. The values of the bytecode table tag column are not constrained to be HEADER or BYTE 66
24. Unconstrained columns on the bytecode HEADER rows 70 25. decompose_limb does not work as intended 72
26. Zero modulus will cause a panic 73 27. The ConstraintBuilder::condition API is dangerous 74
28. The EXTCODECOPY opcode implementation does not work when the account address does not exist 76
A. Vulnerability Categories 77 B. Code Maturity Categories 79 C. Code Quality Findings 80
D. Automated Analysis Tool Configuration 88 E. Fix Review Results 91 Detailed Fix Review Results 94
F. Fix Review Status Categories 100 hashey 4 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Executive Summary Engagement Overview

Scroll engaged hashey to review the security of its zkEVMhalo2circuits. The scoped codebase implements the Ethereum Virtual Machine (EVM) opcodes, as well as other crucial components of Scroll's zkEVM: the State circuit, the Bytecode circuit, and the circuit for the modexp precompile. Additionally, we were tasked with reviewing changes to the Keccak circuit and to the halo2-lib and snark-verifier circuits. A team of four consultants conducted the review from April 17 to June 23, 2023, for a total of 23 engineer-weeks of effort. Our testing efforts focused on circuit soundness and the correct implementation of the EVM semantics. With full access to the source code and documentation, we performed static and dynamic testing of the zkEVMhalo2circuits, using automated and manual processes. Observations and Impact This security review revealed many high-impact vulnerabilities related to circuit soundness.

Due to incorrect, incomplete, or missing constraints, a malicious prover could convince a verifier of an EVM execution that does not match the correct EVM semantics. Specifically, an attacker could cause opcodes to return an unintended result (TOB-SCROLL-1, TOB-SCROLL-3), manipulate gas costs of certain opcodes (TOB-SCROLL-7, TOB-SCROLL-12), execute different opcodes from what was specified in the bytecode (TOB-SCROLL-13, TOB-SCROLL-14), or call certain opcodes in situations where doing so should not be allowed according to the EVM specification (TOB-SCROLL-9). All of these cases could lead to stated divergence and cause loss of funds if the zkEVM is used in the context of a bridge. The gas cost manipulation issues could also enable denial-of-service attacks, and finding TOB-SCROLL-9 could enable reentrancy attacks at the contract level. Recommendations

Based on the codebase maturity evaluation and findings identified during the security review, hashey recommends that Scroll take the following steps before deployment:

- Remediate the findings disclosed in this report. These findings should be addressed as part of a direct remediation or as part of any refactoring that may occur when addressing other recommendations.
-

Invest in adversarial testing. The security requirements of a zkEVM implementation are extremely strict, and the complexity of evaluating such an implementation is extremely high. The use of nondeterministic computation in a complex circuit such as the zkEVM simultaneously a crucial optimization technique and a source of potentially catastrophic errors, as described in TOB-SCROLL-13. Even errors that seem extremely minor, such as TOB-SCROLL-7, hashey
5 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

can lead to divergent execution. And even in the absence of divergent execution bugs, the complexity of executing EVM bytecode provides a wealth of potential implementation mistakes, where minor errors can cause stated divergence if a zkEVM is improperly deployed as part of a cross-chain bridge—for example, if one side of the bridge uses the zkEVM circuit and the other uses a go-ethereum implementation tweaked in accordance with Scroll's public documentation. In addition, any of these errors can be introduced by seemingly innocuous changes during the development process. With all of these considerations in mind, we believe that this project requires an unusually high level of assurance and will continue to need active testing and review throughout its lifetime. We strongly encourage Scroll to invest in an ongoing, adversarial testing process, especially focused on potential malicious prover behavior. We also encourage Scroll to develop a clear written specification for each specialized argument (e.g., the RW table, RLC-based words) used in the final implementation. Finally, invest in a continuous internal code-reviewing effort.

- Leverage the Rust type system to enforce compile-time invariants. The Rust

typesystemshouldbeusedtoensurethatcertainwitnessvalueshavebeen constrained.Asanexample,certainBooleanoperationsdefinedinthecodebase requiretheirargumentstobeBoolean.However,ensuringthattheargumentsare Booleaniscurrentlyeitherenforcedinanadhocmannerorsometimesnoteven enforced,leadingtopotentialsoundnessissues.DefiningnewRusttypesand signaturesfortheBooleanfunctionsandthe `ConstraintBuilder::query_bool()` functionwouldallowdeveloperstoknow thatthosevaluesareBoolean-constrained.Thismechanismalsocanimprove efficiencybypreventingpotentialdouble-enforcementofthesameconstraint. Thefollowingtablespvidethenumberoffindingsbyseverityandcategory.

EXPOSUREANALYSIS	Severity	Count	High	8	Medium	4	Low	2	Informational	13
CATEGORYBREAKDOWN	Category	Count	Cryptography	1	DataValidation	25	Patching	1	hashey	6
ScrollzkEVMhalo2CircuitsSecurityAssessment	PUBLIC									

ProjectSummary ContactInformation Thefollowingmanagerswereassociatedwiththisproject:

DanGuido,AccountManagerBrookeLanghorne,ProjectManager dan@hashey.io, brooke.langhorne@hashey.io
Thefollowingengineerswereassociatedwiththisproject: FilipeCasal,ConsultantJoeDoyle,Consultant filipe.casal@hashey.io, joseph.doyle@hashey.io
OpalWright,ConsultantWillSong,Consultant opal.wright@hashey.io, will.song@hashey.io

ProjectTimeline

The significant events and milestones of the project are listed below.

Date	Event
April 17, 2023	Pre-project kickoff call
April 21, 2023	Status update meeting #1
April 28, 2023	Status update meeting #2
May 5, 2023	Status update meeting #3
May 12, 2023	Status update meeting #4
May 19, 2023	Status update meeting #5
May 26, 2023	Status update meeting #6
June 2, 2023	Status update meeting #7
June 9, 2023	Status update meeting #8
June 16, 2023	Status update meeting #9
June 26, 2023	Delivery of report draft and report readout meeting
October 12, 2023	Delivery of comprehensive report with fix review appendix

hashey

7 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

ProjectGoals The engagement was scoped to provide a security assessment of the ScrollzkEVMhalo2

circuits. Specifically, we sought to answer the following non-exhaustive list of questions:

- Are the circuits implementing the EVM opcodes properly enforcing the EVM semantics?
 - Are the circuits sound and complete?
 - Are EVM state semantics properly enforced?
 - Are state-changing opcodes allowed within a static context?
 - Are the fixed lookup tables properly populated?
 - Are the lookup tables properly constrained to prevent malicious insertions?
 - Is the EVM execution control flow correctly enforced?
 - Can a malicious prover provide different execution traces for the same State-Transaction input?
 - Can a malicious prover convince a verifier of a different final state than what running the EVM would reach?
 - Does the constraint builder API correctly enforce the intended constraints?
 - How are opcode-to-opcode transitions enforced?
 - How are opcode-to-error transitions enforced?
 - Could an attacker trace lead to an error state when the correct execution should not error?
 - Could the opposite also occur?
 - Are memory semantics correctly guaranteed by the RW table constraints?
 - Do the code changes to `halo2-lib`, `snark-verifier`, or the Keccak circuit introduce any vulnerabilities?
 - Do the code changes improve readability and code maintainability?
 - For a correct Fiat-Shamir transformation, are the statement and all prover messages included in the transcript?
- hashey
- 8 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

ProjectTargets The engagement involved a review and testing of the following targets.

Repository	Path	Version
scroll-tech/zkevm-circuits		e8bcb23e1f303bd6e0dc52924b0ed85710b8a016
TypesRust,halo2 PlatformNative	snark-verifier/codedi	Repository scroll-tech/snark-verifier Version a3d0a5ab48522bc533686da3ea8400282c91f536
TypesRust,halo2 PlatformNative	modexp	Repository scroll-tech/misc-precompiled-circuit Version 05725ec61d52d29a063395b0a1130467bee0d2f1
TypesRust,halo2 PlatformNative	halo2-lib/codedi	Repository scroll-tech/halo2-lib Version a805052→b1d1567
TypesRust,halo2 PlatformNative	Bytecode/circuit	Repository scroll-tech/zkevm-circuits/src/bytecode_circuit Version 44000e55eddaec42da958f2555d9bdeec8b865c2
TypeRust,halo2 PlatformNative	hashey	9 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

ProjectCoverage This section provides an overview of the analysis coverage of the review, as determined by our high-level engagement goals. Our approaches included the following:

- We manually reviewed the zkEVM circuit, including the mathematical gadgets, utilities, and all EVM opcode gadgets. We checked for circuit soundness issues, opcode edge cases, underconstrained witness values, and correct state transition enforcement. If we identified vulnerable code patterns or code smells, we used Semgrep to perform variant analysis and search for other instances of the same patterns.
- We manually reviewed the RW table circuits, especially focused on lookups performed in execution gadgets in the `zkevm-circuits/src/evm_circuit/execution/` directory and the structural constraints in the `zkevm-circuits/src/state_circuit/` directory.
- We manually reviewed the Bytecode table consistency circuit, including its Poseidon extended column variant. We checked the correct enforcement of the specified constraints. We looked for potential ways to break the soundness of the circuit.

WemanuallyreviewedtheKeccakcircuitcode,particularlythefilesinthekzvm-circuits/src/keccak_circuit/directory. • Wemanuallyreviewedthehalo2-lib and snark-verifier diffs, paying special attention to changes in the snark-verifier/src/pcs.rs file, as well as the snark-verifier/src/pcs/directory. • Wemanuallyreviewedthe modexp precompile circuit. At the time of the audit, the circuit was still being developed and was incomplete in terms of features and engineering work; the current version of the circuit does not support arbitrary length values like the EVM modexp precompile, and there are several empty functions, which impact the soundness and completeness of the circuit. hashey 10 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Automated Testing hashey uses automated techniques to extensively test the security properties of software. We use both open-source static analysis and fuzzing utilities, along with tools developed in-house, to perform automated testing of source code and compiled software. TestHarnessConfiguration We used the following tools in the automated testing phase of this project: ToolDescriptionPolicy Semgrep An open-source static analysis tool for finding bugs and enforcing code standards when editing or committing code and during build time Appendix D Clippy An open-source Rust linter used to catch common mistakes and unidiomatic Rust code Appendix D Areas of Focus Our automated testing and verification focused on the following: • Identification of general code quality issues and unidiomatic code patterns • Identification of dangerous halo2-specific and Scroll's API patterns Test Results The results of this focused testing are detailed below. Clippy • zkEVM-circuits: The zkEVM-circuits code base has several informational Clippy warnings: unlined_format_args and unnecessary_cast warnings. • Modexp precompile circuit: The modexp precompile circuit has several Clippy warnings that should be addressed. We recommend that Scroll add a Clippy GitHub action to all of its Rust repositories. Semgrep We present some of the rules that we wrote to find halo2-specific and Scroll's API patterns in appendix D. hashey 11 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Codebase Maturity Evaluation hashey uses a traffic-light protocol to provide each client with a clear understanding of the areas in which its code base is mature, immature, or underdeveloped. Deficiencies identified here often stem from root causes within the software development lifecycle that should be addressed through standardization measures (e.g., the use of common libraries, functions, or frameworks) or training and awareness programs. Category Summary Result Arithmetic We found no issues related to the use of integer arithmetic in the code base. Satisfactory Complexity Management The code base is well organized and separated into files and folders. Specific gadget implementations are also cleanly implemented and separate the constraint generation from the witness generation. This separation allows better scrutiny of the constraints imposed in each circuit. On the other hand, the execution of the zkEVM depends on non-deterministic programming patterns, which are hard to reason about and to have a global understanding of. We found two instances where a malicious prover could hijack the execution flow of the executing bytecode: TOB-SCROLL-13 and TOB-SCROLL-14. Moderate Cryptography and Key Management We found no issues related to the use of cryptography primitives in the code base. However, the code base depends upon an outdated version of the halo2-ecc library, which has had several updates made to its cryptographic primitives, as described in TOB-SCROLL-4. All uses of halo2-ecc cryptographic primitives should be carefully reviewed, and Scroll should use an up-to-date and well-tested version of that library. Satisfactory Documentation There are several sources of documentation instead of one centralized and up-to-date documentation. It is common to find missing or incomplete sections of documentation, which should be addressed. We also Moderate hashey 12 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

found that while some documentation has the general design description, it frequently lacks correctness requirements. Memory Safety and Error Handling The kzvm-circuits project uses non-safe Rust code. We found one runtime panic during witness generation: TOB-SCROLL-26. Satisfactory Testing and Verification The code base uses geth tracing to obtain values for witness generation. By using geth's ground truth, tests will mostly exercise the completeness aspect of the circuits. As we found many high-severity findings related to circuit soundness (TOB-SCROLL-1, TOB-SCROLL-3, TOB-SCROLL-7, TOB-SCROLL-9, TOB-SCROLL-12, TOB-SCROLL-13, TOB-SCROLL-14), we believe it is necessary to develop an adversarial testing process, especially focused on malicious prover behavior. Formal methods to check for, such as circuit determinacy, should also be considered for research. We also found gadgets without tests (e.g., binary_number.rs, rlp.rs) and ignored tests (e.g., in jump.rs). Tests should be re-enabled and added where there are none. We also recommend adding the test vectors for all EIPs that the code base implements. As an example, the SSTORE gas refund EIP is implemented in the code base, but its test vectors are not present in the code base to ensure a correct implementation. Weak hashey 13 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC


```
[&a_reduced,&b,&d,&e], None); //3.k2*n+r=d*2^256+e
letmul512_right=MulAddWords512Gadget::construct(cb,[&k,&n,&d,&e],Some(&r)); //(r<n)orn==0
letn_is_zero=IsZeroGadget::construct(cb,sum::expr(&n.cells));
letlt=LtWordGadget::construct(cb,&r,&n); cb.add_constraint( "(1-(r<n)-(n==0))", 1.expr()-lt.expr()-
n_is_zero.expr(), hashey 18ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

); Figure1.2: zkevm-circuits/src/evm_circuit/execution/mulmod.rs#58-73 ExploitScenario AbridgeusestheScrollzkEVMtotrackthecurrentstateofEthereum.Amaliciousprover generatesaproofofexecutionforatransactioninvolvingthe MULMOD instructionwith ,andsetstheresulttoasdescribedabove.Theproversubmitsthatproof,the 0=00·0 resultsofwhichwillnotmatchthecorrectEVMsemantics,leadingtostatedivergenceand lossoffunds. Recommendations Shortterm,fixtheconstraint;extendthe assign functionto receivethe a_or_zero witness.Addtestsforthisfinding. Longterm,adddeterminacytestingtoanygadgets that constrain nondeterministic witnesses. hashey 19ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

2.TheRlpU64Gadgetisunderconstrainedwhenis_lt_128isfalse Severity:HighDifficulty:Medium Type:DataValidationFindingID:TOB-SCROLL-2 Target:z kevm-circuits/src/evm_circuit/util/math_gadget/rlp.rs Description The RlpU64Gadget constrainswitnessvalvestomatchtheoutputofacorrectRLP encoding.Since thelengthandvalueoftheRLP-encodedvaluedependonthevaluebeing lessthan128,the is_lt_128 flagispartofthewitness.Arangecheckensures thatif is_lt_128 istrue,thenthevalueisactuallybelow128.However,thereisnoconstraint ensuringthat value isabove127when is_lt_128 isfalse: letis_lt_128=cb.query_bool(); cb.condition(is_lt_128.expr(),|cb|{ cb.range_lookup(value,128); }); Figure2.1: evm_circuit/util/math_gadget/rlp.rs#L67-L70 Thismeansthatamaliciousprovercouldhaveavaluesmallerthan128butset is_lt_128 tofalse,leadingtoanincorrectlengthandRLP-encodedoutput. ExploitScenario AmaliciousproverinterpretstwobytesinanRLP-serializeddatastructureasavalueless than128,causinglaterfieldsinthedatastructuretobedeserializedstartingatanincorrect offset.Theproverthenusesthisincorrectlydeserializeddatastructuretoproveaninvalid statetransition,leadingtostatedivergenceandpotentiallossoffunds. Recommendations Shortterm,addaconstrainttoensurethatthevalueisabove127when is_lt_128 is false. Longterm,addnegativetestsenvironmentthatmismatchedwitnesses value and is_lt_128 donotsatisfythecircuitconstraints. hashey 20ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

3.TheBLOCKHASHopcodeisunderconstrainedandallowsthehashofany blocktobecomputed Severity:HighDifficulty:Medium Type:DataValidationFindingID:TOB-SCROLL-3 Target:z kevm-circuits/src/evm_circuit/util/math_gadget/rlp.rs Description The BLOCKHASH opcode returnsthehashoftheblockidentifiedbythestackargument, block_number ,providedthatitisonethatofthe256mostrecentcompleteblocks.However, theimplementationallowsamaliciousprovertoprovideanonzeroreultevenwhenthe providedblocknumberisnotamongthe256mostrecentblocks,contradictingtheEVM specification. Tovalidate that block_number isamongthe256mostrecentblocks,theimplementation checks that current_block_number-block_number<257 ,where current_block_number issupposedtobetheblocknumberofthecurrentblock. However, current_block_number isunconstrainedandcouldtakeanyvalue: impl<F:Field>ExecutionGadget<F>forBlockHashGadget<F>{ constNAME:&'staticstr="BLOCKHASH"; constEXECUTION_STATE:ExecutionState=ExecutionState::BLOCKHASH; fnconfigure(cb:&mutConstraintBuilder<F>)->Self{ letcurrent_block_number=cb.query_cell(); letblock_number=WordByteCapGadget::construct(cb, current_block_number.expr()); cb.stack_pop(block_number.original_word()); //FIXME //cb.block_lookup(//BlockContextFieldTag::Number.expr(), //cb.curr.state.block_number.expr(), //current_block_number.expr(), //); letblock_hash=cb.query_word_rlc(); letdiff_lt=LtGadget::construct(cb, current_block_number.expr(), (NUM_PREV_BLOCK_ALLOWED+1).expr()+block_number.valid_value(), hashey 21ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Figure3.1: zkevm-circuits/src/evm_circuit/execution/blockhash.rs#L33-L49 Amaliciousprovercouldprovideaninvalid current_block_number andreturnthehash ofanyblockpresentintheblocklookuptable,independentofitsblocknumber. ExploitScenario Amaliciousprovergeneratesaproofofexecutionforatransactioninvolvingthe BLOCKHASH opcode thatresultsinanonzerohashforanolderblock.Theproversubmits thatproof,theresultsofwhichwillnotmatchthecorrectEVMsemantics,leadingtostate divergenceandlossoffunds. Recommendations Shortterm,addthemissinglookupconstraintforthe current_block_number witness. Longterm,trackandtriage"FIXME"and"TODO"itemsinacentralizedissuetracking

system, such as Git Hub issues. Adding failing tests where security-relevant "TODO" items are identified. hashey 22 ScrollzkEVM halo2CircuitsSecurityAssessment PUBLIC

4. zkevm-circuits crates depends on an outdated version of halo2-ecc Severity: Medium Difficulty: Medium
Type: Patching Finding ID: TOB-SCROLL-4 Target: z kevm-circuits/{Cargo.toml, src/tx_circuit/sign_verify.rs} Description The zkevm-circuits crates depends on the halo2-ecc library in Scroll's fork of halo2-lib, which provides halo2 circuits for elliptic curve and finite field operations. As illustrated in figure 4.1, this crate depends on the halo2-ecc-snark-verifier-0323 tag, which currently points to commit d24871338ade7dd56362de517b718ba14f3e7b90. halo2-base={git="https://github.com/scroll-tech/halo2-lib", branch="halo2-ecc-snark-verifier-0323", default-features=false, features=["halo2-pse", "display"]} halo2-ecc={git="https://github.com/scroll-tech/halo2-lib", branch="halo2-ecc-snark-verifier-0323", default-features=false, features=["halo2-pse", "display"]} Figure 4.1: zkevm-circuits/Cargo.toml#32-33 The Scroll fork of halo2-lib is closely related to the upstream halo2-lib library. In particular, the v0.3.0 version of halo2-ecc (commit c31a30bcaff384b0c3aa7c823dd343f5c85da69e) has a common ancestor commit of 4338af81bb2de4f278467e5c484e067c064cc66b with the Scroll version. Figure 4.2: The Git commit history of halo2-lib with the common ancestor of the minimize-diff and upgrade-v0.3.0 branches highlighted. The upstream library has various fixes and improvements that should be incorporated. Some notable existing fixes include the following: • FpChip::assert_equal has a data soundness-related typo fixed (PR#18). hashey 23 ScrollzkEVM halo2CircuitsSecurityAssessment PUBLIC

• ecdsa_verify_no_pubkey_check no longer rejects certain invalid signatures (PR #36). While FpChip::assert_equal does not currently appear to be used, the SignVerifyChip circuit uses the ecdsa_verify_no_pubkey_check function, as shown in figure 4.3. //returnstheverificationresultofecdsasignature // //WARNING: this circuit does not enforce the returned value to be true //makesurethecallerchecksthisresult! let ecdsa_is_valid = ecdsa_verify_no_pubkey_check::(<Fp, Fq, Secp256k1Affine>(&ecc_chip.field_chip, ctx, &pk_assigned, &integer_r, &integer_s, &msg_hash, 4, 4,)); Figure 4.3: zkevm-circuits/src/tx_circuit/sign_verify.rs#386-399 SignVerify is then used to check signatures on EVM transactions, as shown in figure 4.4, and because of the pre-patch behavior, an adversary can generate a correctly signed transaction that will nevertheless fail signature verification. #[cfg(feature="enable-sign-verify")] { let assigned_sig_verifs = self.sign_verify .assign(&config.sign_verify, layouter, &sign_data, challenges)?; self.sign_verify.assert_sig_is_valid(&config.sign_verify, layouter, assigned_sig_verifs.as_slice(),); self.assign(config, challenges, layouter, assigned_sig_verifs, Vec::new(), &padding_txs,); } Figure 4.4: zkevm-circuits/src/tx_circuit.rs#1804-1822 hashey 24 ScrollzkEVM halo2CircuitsSecurityAssessment PUBLIC

Exploit Scenario An adversary creates a transaction with a valid signature that the old implementation would reject and submit it to Ethereum. Ethereum accepts the transaction, but the Scroll zkEVM is unable to accept it, stalling the zkEVM and creating a denial of service that may freeze user funds. Recommendations Short term, review the security implications of this outdated version of halo2-ecc on the zkEVM codebase. Then, either update to a more recent version of halo2-lib that incorporates upstream fixes or backport those fixes to Scroll's fork. Long term, keep all dependencies up to date whenever possible. For any dependencies that have been forked from the upstream version, develop a plan to port any upstream security updates to that fork. hashey 25 ScrollzkEVM halo2CircuitsSecurityAssessment PUBLIC

5. N_BYTES parameters are not checked to prevent overflow Severity: Informational Difficulty: N/A
Type: Data Validation Finding ID: TOB-SCROLL-5 Target: z kevm-circuits/src/evm_circuit/util/math_gadget/{constant_division, lt}.rs Description The ConstantDivisionGadget and LtGadget circuits implement operations on multi-byte integers: division by a constant value and comparison, respectively. Each circuit has a generic parameter N_BYTES representing the number of bytes used. However, each of these circuits has additional implied restrictions on N_BYTES required to prevent unexpected behavior due to overflowing field elements. In ConstantDivisionGadget (shown in figure 5.1), the quotient value is constrained to be less than, but the expression quotient.expr()*denominator.expr() 256 0_0000 may overflow if denominator is sufficiently large (e.g., if denominator is 1024 and N_BYTES is 31). The comment highlighted in figure 5.2 provides sufficient conditions to prevent overflow, but these are not fully enforced either in the circuit or in assertions at circuit construction time. let quotient = cb.query_cell_with_type(CellType::storage_for_expr(&numerator)); let remainder = cb.query_cell_with_type(CellType::storage_for_expr(&numerator)); //Require that remainder < denominator cb.range_lookup(remainder.expr(), denominator);

```

//Requirethatquotient<256**N_BYTES //sowecan'thaveanyoverflowwhen`quotient*denominator`.
letquotient_range_check=RangeCheckGadget::construct(cb,quotient.expr());
//Checkifthedivisionwasdonecorrectly cb.require_equal( "numerator-remainder==quotient*denominator",
numerator-remainder.expr(), quotient.expr()*denominator.expr(), ); Figure5.1: zkevm-
circuits/src/evm_circuit/util/math_gadget/constant_division.rs#33- 48 hashey
26ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

///Returns(quotient:numerator/denominator,remainder:numerator%denominator),
///with`numerator`anexpressionand`denominator`aconstant. ///Inputrequirements: ///-
`quotient<256**N_BYTES` ///-`quotient*denominator<fieldsize` ///-
`remainder<denominator`requiresarangelookuptablefor`denominator` #[derive(Clone,Debug)]
pubstructConstantDivisionGadget<F,constN_BYTES:usize>{ Figure5.2: zkevm-
circuits/src/evm_circuit/util/math_gadget/constant_division.rs#13- 20
InLtGadget(showninfigure5.3),valuesofN_BYTESabove31willcausetobean
unconstrainedBoolean,sinceamaliciousprovercanset diff totherepresentationof (rhs-lhs) evenif
rhs<lhs .Thecommenthighlightedinfigure5.4describessufficient
conditionstopreventoverflowwithoutchangestotheircuit,buttheserestrictionsare enforcedonlybya
debug_assert! in from_bytes::expr (showninfigure5.5). letlt=cb.query_bool();
letdiff=cb.query_bytes(); letrange=pow_of_two(N_BYTES*8); //Theequationwewouldrequiretohold:`lhs-
rhs==diff-(lt*range)`. cb.require_equal( "lhs-rhs==diff-(lt*range)", lhs-rhs,
from_bytes::expr(&diff)-(lt.expr()*range), ); Figure5.3: zkevm-
circuits/src/evm_circuit/util/math_gadget/lt.rs#37-46
///Returns`1`when`lhs<rhs`,andreturns`0`otherwise. ///lhsandrhs`<256**N_BYTES`
///`N_BYTES`isrequiredtobe`≤MAX_N_BYTES_INTEGER`topreventoverflow:
///valuesarestoredinasinglefieldelementandtwooftheseareadded ///together.
///Theequationthatisenforcedis`lhs-rhs==diff-(lt*range)`.
///Becauseallvaluesare`≤256**N_BYTES`and`lt`isboolean,`lt`canonly ///be`1`when`lhs<rhs`. #
[derive(Clone,Debug)] pubstructLtGadget<F,constN_BYTES:usize>{ Figure5.4: zkevm-
circuits/src/evm_circuit/util/math_gadget/lt.rs#14-23 pub(crate)fnexpr<F:FieldExt,E:Expr<F>>
(bytes:&[E])→Expression<F>{ debug_assert!( bytes.len()≤MAX_N_BYTES_INTEGER,
"Too many bytes to compose an integer in field" ); hashey 27ScrollzkEVMhalo2CircuitsSecurityAssessment
PUBLIC

```

Figure5.5: zkevm-circuits/src/evm_circuit/util.rs#528-532 ExploitScenario
A developer who is unaware of these issues uses the ConstantDivisionGadget or LtGadget circuit with values of N_BYTES that are too large, causing potentially underconstrained circuits. Recommendations
Short term, add explicit checks at circuit construction time to ensure that N_BYTES is limited to values that prevent overflow. Long term, consider performing these validations at compile time with static_assertions or asserts in a const context. hashey 28ScrollzkEVMhalo2CircuitsSecurityAssessment
PUBLIC

6. Differences in shared code between zkevm-circuits and halo2-lib Severity: Medium Difficulty: High
Type: Data Validation Finding ID: TOB-SCROLL-6 Target: Several files Description
The code base contains code that is also present in the halo2-lib code base (not through a dependency) and it does not match in all cases. For example, these several constraint_builder functions use the debug_assert!() macro for important validations, which will not perform those checks in release mode.
pub(crate)fncondition<R>(&mutself, condition:Expression<F>, constraint:implFnOnce(&mutSelf)→R,)->R{ debug_assert!(self.condition.is_none(), "Nested condition is not supported");
self.condition=Some(condition); letret=constraint(self); self.condition=None; ret } Figure6.1:
evm_circuit/util/constraint_builder.rs#L216-L229
pub(crate)fnvalidate_degree(&self,degree:usize,name:&'staticstr){ ifself.max_degree>0{
debug_assert!(degree≤self.max_degree, "Expression{}degree too high: {}>{}", name, degree,
self.max_degree,); } } Figure6.2: evm_circuit/util/constraint_builder.rs#L246-L256
pub(crate)fnvalidate_degree(&self,degree:usize,name:&'staticstr){ hashey
29ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

```

//WeneedtosubtractIMPLICIT_DEGREEfromMAX_DEGREEbecauseallexpressions
//willbemultipliedbystateselectorandq_step/q_step_first //selector. debug_assert!(
degree≤MAX_DEGREE-IMPLICIT_DEGREE, "Expression{}degree too high: {}>{}", name, degree, MAX_DEGREE-
IMPLICIT_DEGREE, ); } Figure6.3: evm_circuit/util/constraint_builder.rs#L1370-L1381
Thecodebasealsoincludes the log2_ceil functionin zkevm-circuits/src/util.rs ,
whichmiscomputesitsresultonazeroinput—thebehaviorhasbeenfixedinPR#37for halo2-lib . Recommendations  

Short term, fix the issues in common with the halo2-lib code base. Also, check all uses of debug_assert  

throughout the code base and ensure that they are not used to validate  

critical invariants, as they will not run in release mode.

```

7.UnderconstrainedwarmstatusonCALLopcodesallowsgascostforgery Severity:HighDifficulty:Medium
Type:DataValidationFindingID:TOB-SCROLL-7 Target:z kevm-
circuits/src/evm_circuit/execution/callop.rs Description Anunderconstrainedvariableinthe
CallOpGadget allowsanattacker toprovethethe
executionofatransactionwithincorrectgascostsbysettinganaddressascoldwhenit shouldbecomewarm. The
CallOpGadget implements the CALL , CALLCODE , DELEGATECALL ,and STATICCALL
EVMopcodes.Thegascostoftheseopcodesdependsonwhetherthecalleeaddressis
warm.Additionally,theimplementationoftheseopcodesmustmaketheaddresswarmso
thatfuturecallstothesameaddresscostlessgas.However,thevariablethatcontrolsthe
address'snewwarmstatusisnotconstrainedandisreferencedonlyinthewritetotheRW table:
//Addcalleetoaccesslist letis_warm=cb.query_bool(); letis_warm_prev=cb.query_bool();
cb.account_access_list_write(tx_id.expr(), call_gadget.callee_address_expr(), is_warm.expr(),
is_warm_prev.expr(), Some(&mutreversion_info),); Figure7.1: zkevm-
circuits/src/evm_circuit/execution/callop.rs#L129-L138 Thismeansthatamaliciousprovercanmakethe
is_warm variableequal false ,causinga calleedaddresstoactuallybecomecoldduringtheexecutionofa CALL
,insteadofwarmas intheEVMspecification.

AconstraintontheRWtable,requiringthattheinitialvalueoftheaccesslistelements
alwaysfalse,preventsanotherpossiblescenariowherethe is_warm_prev valuecouldbe
definedaswarmenthoughtheadresshadnotbeenaccessedbefore. ExploitScenario
Amaliciousprovergeneratesaproofofexecutionforatransactioninvolvingtwo CALL
opcodes tothesameaddress thatresults in different gas costs from the EVM specification:
hashey
31ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

inthe first CALL opcode execution, the prover sets the address as cold instead of warm,
causing the wrong gas calculation for the second call. The prover submits that proof, the
result of which will not match the correct EVM semantics, leading to stated divergence and loss of funds.
Recommendations Shortterm, add constraint to ensure that the callee address becomes warm on the CALL
opcodes, by constraining is_warm to be true. hashey
32ScrollzkEVMhalo2CircuitsSecurityAssessment
PUBLIC

8.RWtableconstantsmustmatchexactlywhentheverificationkeyis created
Severity:InformationalDifficulty:N/A Type:DataValidationFindingID:TOB-SCROLL-8 Target:RWtable
Description NearlyallruntimestateofEVMprogramexecutionistrackedandvalidatedinalookup
tablereferredtoasthe"RWtable."Thistableenforcescorrectinitializationandcoherency
ofreadandwriteoperationsforaddressablepartsofthestate,includingthestack,
memory,andaccountstorage,aswellasinputsandoutputssuchasthetransactionaccess
listandthetransactionlog.Figure8.1showsthestoragetypescombinedinthis table: pubenumRwTableTag{
///Start(usedforpadding) Start=1, ///Stackoperation Stack, ///Memoryoperation Memory,
///AccountStorageoperation AccountStorage, ///TxAccessListAccountoperation TxAccessListAccount,
///TxAccessListAccountStorageoperation TxAccessListAccountStorage, ///TxRefundoperation TxRefund,
///Accountoperation Account, ///CallContextoperation CallContext, ///TxLogoperation TxLog,
///TxReceiptoperation TxReceipt, } Figure8.1: zkevm-circuits/src/table.rs#354-377
ThezkEVMcircuitenforcescorrectmemoryoperationresultsforEVMopcodesby
performinglookupsintothis table,as shown in figure 8.2.Callssuchas memory_lookup are translated into
Lookup::Rw values (shown in figure 8.3), which are then further hashey
33ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

translated into multi-column lookups into the RW table, as illustrated in figure 8.4. Note
that in lookups, the first column, corresponding to the fixed column q_enable , is always set to 1.
cb.condition(is_mstore8.expr(), |cb| { cb.memory_lookup(1.expr(), from_bytes::expr(&address.cells),
value.cells[0].expr(), None,); }); Figure 8.2: A memory lookup (zkevm-
circuits/src/evm_circuit/execution/memory.rs#73-80) Rw{ ///Counterforhowmuchread-
writehavebeendone,whichstandsfor ///thesequentialtimestamp. counter:Expression<F>,
///Abooleanvaluetospecifyiftheaccessrecordisaread or write. is_write:Expression<F>,
///Tagtospecifywhichread-writedata to access, seeRwTableTagfor ///alltags. tag:Expression<F>,
///Valuescorrespondingtothetag. values:RwValues<F>, }, Figure 8.3: Lookup::Rw (zkevm-
circuits/src/evm_circuit/table.rs#197-208) Self::Rw{ counter, is_write, tag, values, }=>{ vec![
1.expr(), counter.clone(), is_write.clone(), tag.clone(), values.id.clone(),
values.address.clone(), values.field_tag.clone(), values.storage_key.clone(), values.value.clone(),
values.value_prev.clone(), values.aux1.clone(), values.aux2.clone(),] hashey
34ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

} Figure 8.4: Conversion to lookup columns, with q_enable highlighted (zkevm-
circuits/src/evm_circuit/table.rs#321-341)

However, these lookups are only the existence of such rows, and for correct execution, it is vital that the reads and writes present in the table are the following:

- Coherent with the external state: The first read of any data is correctly initialized, and the last write of externally visible data (e.g., storage) is reflected in the Ethereum state commitment.
- Coherent with each other: Values in read operations match the most recent written value or the initial value.
- Coherent with the execution trace: Every entry in the RW table corresponds to exactly one memory-access-generating step in the execution trace. Equivalently, there are no "extra" entries in the table.

These global constraints on the RW table are enforced through three major checks.

First, the RW table is lexicographically ordered with respect to its columns. Several constraints are used to enforce lexicographic ordering. An illustrative example is shown in figure 8.5. Note that the constraint is conditional on the fixed column selector. All other lexicographic ordering constraints are also conditional on this fixed column, so any rows where `selector == 0` are not required to be ordered.

```

meta.create_gate("limb_difference_is_not_zero", |meta| {
  let selector = meta.query_fixed(selector, Rotation::cur());
  let limb_difference = meta.query_advice(limb_difference, Rotation::cur());
  let limb_difference_inverse = meta.query_advice(limb_difference_inverse, Rotation::cur());
  vec![selector * (1.expr() - limb_difference * limb_difference_inverse)]
});

```

Figure 8.5: A lexicographic ordering constraint (`zkvm-circuits/src/state_circuit/lexicographic_ordering.rs#128-134`)

Second, a large collection of structural properties on the sorted table are enforced. Figures 8.6 and 8.7 show examples of such constraints. When the rows are sorted, all operations involving the same address or storage identifier are grouped together, sorted in increasing order by the final two columns, which represent the value `rw_counter` in big-endian. Note that in increasing values of `rw_counter` are treated as "happening later in time," as shown in the highlighted portion of figure 8.6, which enforces that reads do not change the value by requiring that `value == value_prev` in read entries.

hasheye 35 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC

Unlike in other parts of the table, the constraints applied to Start rows (illustrated in figure 8.7) use the `rw_counter` fields for a different purpose. Every Start row where `lexicographic_ordering_selector == 1` is required to exactly increase `rw_counter` by 1. Start rows do not represent memory operations, and thus can be thought of as padding.

```

// When all the keys in the current row and previous row are equal.
self.condition(q.not_first_access.clone(), |cb| {
  cb.require_zero("non-first access reads don't change value", q.is_read() * (q.rw_table.value.clone() - q.rw_table.value_prev.clone()), );
  cb.require_zero("initial value doesn't change in an access group", q.initial_value.clone() - q.initial_value_prev(), );
});

```

Figure 8.6: A structural constraint on the RW table (`zkvm-circuits/src/state_circuit/constraint_builder.rs#177-187`)

```

self.require_zero("field_tag is 0 for Start", q.field_tag());
self.require_zero("address is 0 for Start", q.rw_table.address.clone());
self.require_zero("id is 0 for Start", q.id());
self.require_zero("storage key is 0 for Start", q.rw_table.storage_key.clone());
// 1.1. rw_counter increases by 1 for every non-first row
self.require_zero("rw_counter increases by 1 for every non-first row", q.lexicographic_ordering_selector.clone() * (q.rw_counter_change() - 1.expr()), );

```

Figure 8.7: Some constraints applied to Start rows (`zkvm-circuits/src/state_circuit/constraint_builder.rs#192-200`)

Third, a running count of RW lookups is tracked in the `rw_counter` field of `StepState` (shown in figure 8.8) for each step. When execution reaches the `EndBlock` state, two additional lookups are performed, shown in figure 8.9. These lookups ensure that there are at most `max_rws - step.rw_counter` padding rows in the RW table, and are designed to check that there are at most `step.rw_counter` non-padding rows in the table.

```

pub(crate) struct StepState {
  /// The execution state selector for this step
  pub(crate) execution_state: DynamicSelectorHalf,
  /// The Read/Write counter
  pub(crate) rw_counter: Cell,
}

```

Figure 8.8: `StepState` and its `rw_counter` field (`zkvm-circuits/src/evm_circuit/step.rs#456-460`)

hasheye 36 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC

```

// 3. Verify rw_counter counts to the same number of meaningful rows in
// rw_table to ensure there is no malicious insertion.
// Verify that there are at most total_rws meaningful entries in the rw_table
cb.rw_table_start_lookup(1.expr());
cb.rw_table_start_lookup(max_rws.expr() - total_rws.expr());
// Since every lookup done in the EVM circuit must succeed and uses
// a unique rw_counter, we know that at least there are // total_rws meaningful entries in the rw_table.
// We conclude that the number of meaningful entries in the rw_table // is total_rws.

```

Figure 8.9: Constraints ensuring that the RW table has been padded to `max_rws` rows (`zkvm-circuits/src/evm_circuit/execution/end_block.rs#78-87`)

These checks are sufficient to guarantee RW table correctness, assuming the following:

1. The `rw_counter` field of `StepState` correctly tracks how many distinct, non-Start

RWlookupsareperformedintheexecutiontrace. 2. lexicographic_ordering_selector=1 whenever rw_table.q_enable=1 . 3. rw_table.q_enable isasequenceofall 1 sfollowedbyall 0 s. 4. Thereareatmost max_rws rowswhere rw_table.q_enable=1 . Ifanyoftheserequirementsisfalse, amaliciousprovercanproveerroneousexecution tracesbymanipulatingtheRWtableinsomeway: 1. Ifthe rw_counter overcountsthenumberofdistinctRWlookups, arow representingamaliciousmemorywritecanbeinserted. 2. If lexicographic_ordering_selector=0 inanycellwhere rw_table.q_enable=1 , theprovermaybypassnearlyallstructuralproperty checksbypartitioningtheRWtableintotwo"versions,"onestartingatthebeginning ofthetableandonestartingatthatunrestrictedrow. 3. If rw_table.q_enable is 1 , then 0 , then 1 , themiddlerowwillnotcorrespondto anyRWlookup, andthusmaybesettoamaliciousmemorywrite. 4. If rw_table.q_enable is 1 formorethanmax_rws rows, amaliciousmemory writecanbeinserted. Thesefourassumptionsareallpropertiesoffixedrows, constants, orthecircuititself, so theydonotneedtobeconstrainedinthecircuit. However, theyarenotcurrentlyexplicitly enforcedatcircuit-constructiontime, soifanyofthemisviolatedwhengeneratingthe zkEVMverificationkey, thiswillgoundetectedandwouldleadtoglobalcircuit unsoundness. hashey 37ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Unfortunately, thefirstisequivalentto"thereareno rw_counter -relatedbugsintheEVM circuit," soitisdifficulttoenforce. However, therelationshipsbetween lexicographic_ordering_selector , rw_table.q_enable , and max_rws canand shouldbecheckedautomaticallywithassertions. ExploitScenario AnincorrectversionofthezkEVMcircuitisusedtogenerateaverificationkeythatfailsto enforceoneoftheassumptionsabove. AmaliciousproverthencraftsanRWtablethat leadstoincorrectexecutionofatransaction, causingstatedivergenceandpotentiallossof funds. Recommendations Shortterm, add assert!(...) callstoenforcecorrectcorrespondencebetween rw_table.q_enable , lexicographic_ordering_selector , and max_rws . Longterm, reviewanddocumentassumptionsmadeaboutallcircuitconstants. When possible, usetechniquessuchasassertionstochecktheseassumptionsat circuit-constructiontime. hashey 38ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

9. TheCREATEandCREATE2opcodescanbecalledwithinastaticcontext Severity:HighDifficulty:Medium Type:DataValidationFindingID:TOB-SCROLL-9 Target:z kevm-circuits/src/evm_circuit/execution/create.rs Description The CREATE and CREATE2 opcodesaremissingaconstraintthatpreventsthemfrombeing calledinthecontextofastaticcall. Thisallowsforastate-changingoperationthat isnot allowedbytheEVMspecification. Inthecontextofa STATICCALL , thestatecannotbemodified. Asaresult, state-changing opcodeslike CREATE , CREATE2 , LOGX , SSTORE , and CALL areforbiddenwhentheargument valuediffersfrom0, accordingtotheEVMspecification. However, thecurrent implementationofthe CREATE and CREATE2 opcodesdoesnothaveachecktoensurethat thecallingcontexthaspermissiontochangethestate. Bycontrast, theother implementationsofstate-changingopcodeshavethefollowingcheck: //constrainnotinstaticcall letis_static=cb.call_context(None, CallContextFieldTag::IsStatic); cb.require_zero("is_staticisfalse", is_static.expr()); Figure9.1: zkevm-circuits/src/evm_circuit/execution/sstore.rs#L57-L59 Withoutthisvalidationinplace, amaliciousprovercouldgenerateaproofofexecutionfor atransactioninvolvingthe CREATE opcodewithinthecontextofa STATICCALL , leadingto statedivergence. Notethatthe SELFDESTRUCT opcodeisdisabled, butisalso subjecttothenon-static constraintaccordingtotheEthereumYellowPaper. Thisshouldbetakenintoaccountifthe opcodeisimplementedinthefuture. ExploitScenario Alicedeploysaconstant-functionautomatedmarketmaker(AMM)smartcontract AliceMM totheScrollzkEVM. IneachAMMtransaction, AliceMM receivesfundsintokentypeA(or B), thencalculatestheexchangerate, thensendsfundsintokentypeB(orA). Tocalculate theexchangerate, AliceMM callsBob's ComplicatedMath contract. Aliceknowsabout reentrancyattacksandiscarefultocall ComplicatedMath onlywith STATICCALL . However, Bobhasdeployedamaliciousversionof ComplicatedMath thatuses CREATE to call AliceMM inareentrantfashion. Bobcalls AliceMM withamalicioustransactionthat hashey 39ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

manipulatestheexchangerate, thendrainsthecontractoftokenAinexchangeforatiny amountoftokenB, resultinginlossoffunds. Recommendations Shortterm, addtheconstrainttovalidatethattheexecutioncontextdoesnotallow state-changingoperations. Longterm, addtestsforthe CREATE , CREATE2 , LOGX , SSTORE , and CALL opcodeswhen calledwithina STATICCALL . hashey 40ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

10. ResponsibleOpcodeableincorrectlyhandlesCREATEandCREATE2 Severity:InformationalDifficulty:N/A Type:DataValidationFindingID:TOB-SCROLL-10 Target:z kevm-circuits/src/evm_circuit/step.rs Description The ResponsibleOpcode tableisusedtoattributedifferentexecutionstatestoparticular

setsofopcodes. Formanyopcodes, this table is the primary source of truth for which state they transition to. The SameContextGadget (shown in figure 10.1) enforces that executing opcodes correctly use the corresponding state. For example, it enforces that the ADD opcode use the ADD_SUB state. `cb.add_lookup("ResponsibleOpcodeLookup", Lookup::Fixed{ tag: FixedTableTag::ResponsibleOpcode.expr(), values: [cb.execution_state().as_u64().expr(), opcode.expr(), 0.expr(),], },);` Figure 10.1: `zkevm-circuits/src/evm_circuit/util/common_gadget.rs#48-58` This table is populated via the `ExecutionState::responsible_opcodes` method, which also is used for reporting execution statistics. This method does not handle the CREATE2 state, and incorrectly reports both CREATE and CREATE2 as the responsible opcodes for the CREATE state, as shown in figure 10.2: `Self::CREATE => vec![OpcodeId::CREATE, OpcodeId::CREATE2],` Figure 10.2: `zkevm-circuits/src/evm_circuit/step.rs#304` Since the CREATE and CREATE2 opcodes constrain the execution state in a way that does not use SameContextGadget, this does not cause any soundness issues. However, if a similar error were made for another opcode or state in the table, the resulting circuit may be either incomplete or underconstrained. [hashey 41 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC](#)

Recommendations Short term, fix the data in this table by correctly mapping the CREATE and CREATE2 states to the CREATE and CREATE2 opcodes, respectively.

Long term, develop tests to check the consistency of the opcode table against the execution behavior. [hashey 42 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC](#)

11. Elliptic curve parameters omitted from Fiat-Shamir Severity: Informational Difficulty: N/A
Type: Cryptography Finding ID: TOB-SCROLL-11 Target: Several files Description The Fiat-Shamir code in the `snark-verifier` patch does not incorporate the elliptic curve parameters into the transcript. Points are incorporated into the transcript using only the `x` and `y` coordinates, with no reference to the associated curve, and we are not able to find any instances where the curve parameters are explicitly added to a Fiat-Shamir transcript. `fn common_ec_point(&mut self, ec_point: &C) -> Result<(), Error> { let coordinates = Option::<Coordinates<C>>::from(ec_point.coordinates()).ok_or_else(|| { Error::Transcript(io::ErrorKind::Other, "Cannot write points at infinity to the transcript".to_string(),) }); [coordinates.x(), coordinates.y()].map(|coordinate| { self.buf.extend(coordinate.to_repr().as_ref().iter().rev().cloned()); }); Ok(()) }` Figure 11.1: `snark-verifier/src/system/halo2/transcript/evm.rs#L173-L187` Non-interactive proofs must commit exactly to the statement being proven before any challenges are regenerated. If a prover can equivocate about attributes of the statement (e.g., which elliptic curve the points are supposed to be on), a proof for one statement may be passed off as a proof for another, as in the Frozen Heart class of vulnerabilities. (Note that the Frozen Heart PlonK vulnerability discussed in the linked article is not under consideration here; it is only an illustration of Fiat-Shamir vulnerabilities.) The `snark-verifier` code is intended to be curve-agnostic, so a proof generated using one curve may be verified using a different elliptic curve that shares only the points present in the transcript, leading to identical challenge values but a different statement. In general, two different elliptic curves can share only a limited number of points, so the existing code may implicitly commit to the curve being used. However, we have not [hashey 43 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC](#)

determined the exact threshold, and a detailed security proof should be done if that property is relied upon. In the Scroll zkEVM system, the prover and verifier use a fixed set of curve parameters, so it is not possible to convince Scroll software to accept a proof using another curve. Recommendations Short term, include the curve parameters at the beginning of the Fiat-Shamir transcript. Long term, always consider including all public parameters of the system in the Fiat-Shamir transformations. [hashey 44 Scroll zkEVM Halo 2 Circuits Security Assessment PUBLIC](#)

12. The gas cost for the CALL opcode is underconstrained Severity: High Difficulty: Medium
Type: Data Validation Finding ID: TOB-SCROLL-12 Target: `zkevm-circuits/src/evm_circuit/execution/callop.rs` Description The gas cost of the CALL -like opcodes (`CALL`, `CALLCODE`, `DELEGATECALL`, and `STATICCALL`) is not constrained, allowing a malicious prover to spend as much gas as desired in certain conditions. This allows free gas CALL operations if the prover sets this value to zero, or it can cause the transaction execution to terminate after the execution of the current opcode by defining a high gas cost. Both options could cause a stated divergence from an execution following the EVM specification. Figure 12.1 shows the code that gets the witness cell `step_gas_cost` and then uses it unconstrained to set the gas cost of the current opcode. This happens when the call precheck conditions are valid (i.e., the call depth is valid, and the caller balance is enough to transfer the call value), and the called address has no associated code: `let step_gas_cost = cb.query_cell();`

```
letmemory_expansion=call_gadget.memory_expansion.clone(); cb.condition( and::expr([
no_callee_code.expr(), not::expr(is_precompile.expr()), is_precheck_ok.expr(), ]), |cb|{
//Savecaller'scallstate forfield_tagin[ CallContextFieldTag::LastCalleeId,
CallContextFieldTag::LastCalleeReturnDataOffset, CallContextFieldTag::LastCalleeReturnDataLength, ]
{ cb.call_context_lookup(true.expr(),None,field_tag,0.expr()); } }, ); cb.condition(
and::expr([is_precompile.expr(),is_precheck_ok.expr()]), |cb|{ //Savecaller'scallstate hashey
45ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
for(field_tag,value)in[ (CallContextFieldTag::LastCalleeId,callee_call_id.expr()),
(CallContextFieldTag::LastCalleeReturnDataOffset,0.expr()), (
CallContextFieldTag::LastCalleeReturnDataLength, return_data_len.expr(), ), ]{
cb.call_context_lookup(true.expr(),None,field_tag,value); } }, ); cb.condition(
and::expr([call_gadget.is_empty_code_hash.expr(),is_precheck_ok.expr()]), |cb|{
//ForCALLCODEopcode,ithasanextrastackpop`value`andoneaccount read //forcallerbalance(+2). //
//ForDELEGATECALLopcode,ithastwoextracallcontextlookupsfor current //calleraddressandvalue(+2). //
//NoextralookupsforSTATICCALLopcode. lettransfer_rwc_delta=
is_call.expr()*not::expr(transfer.value_is_zero.expr()*2.expr()); letrw_counter_delta=21.expr()
+is_call.expr()*1.expr() +transfer_rwc_delta.clone() +is_callcode.expr()
+is_delegatecall.expr()*2.expr() +precompile_memory_writes;
cb.require_step_state_transition(StepStateTransition{ rw_counter:Delta(rw_counter_delta),
program_counter:Delta(1.expr()), stack_pointer:Delta(stack_pointer_delta.expr()), gas_left:Delta(-
step_gas_cost.expr()), Figure12.1: zkevm-circuits/src/evm_circuit/execution/callop.rs#L255-L314
ExploitScenario Amaliciousprovergeneratesandsubmitsaproofofexecutionforatransactioninvolvinga CALL
toanaddresswithemptycodethatwouldnormallyexhaustthetransaction'sgas.By
definingthegascostaszero,thetransactionsucceeds.However,thisexecutiondoesnot
matchthecorrectEVMsemantics,leadingtostatedivergenceandlossoffunds. Recommendations
Shortterm,addconstraintstocorrectlycomputethegascostforthecallopodes. hashey
46ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
Longterm,addnegativetestensuringthatEVMtracesgascostsdonotsatisfythecircuit constraints. hashey
47ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
13.Unconstrainedopcodesallow nondeterministicexecution Severity:HighDifficulty:Medium
Type:DataValidationFindingID:TOB-SCROLL-13 Target: zkevm-
circuits/src/evm_circuit/execution/{return_revert.rs,
error_code_store.rs,error_invalid_creation_code.rs, error_precompile_failed.rs} Description
Severalopcodesaremissingconstraintsthatensurethecorrectcorrespondencebetween
executionstateandopcode,allowingamaliciousprovertohijackthetransaction execution.
TheScrollzkEVMcircuitchecksthecorrectexecutionofatransactionbyverifyinga prover-
generatedexecutiontrace.Thisexecutiontraceconsistsofseriesofstates,each
representedbyaconstructorofthe ExecutionStateenum .Eachstatecorrespondstoan
"executiongadget"inthecircuit,whichcheckspreconditionsandenforcescorrectupdates
toEVMdatastructuresuchasmemoryandstorage.
IntheScrollcodebase,thecorrespondencebetweentheexecutionstateandopcodeis
enforcedentirelybythesegadgets.Mostexecutiongadgetsusea SameContextGadget
(showninfigure13.1)tocheckthatthecurrent (executionstate,opcode) pair appearsinthe
ResponsibleOpcode table.Executiongadgets thatdonotuse SameContextGadget
mustcheckthatthecurrentopcodeiscorrectforthecurrentstate throughothermeans.Forexample,
Error00GSloadSstoreGadget ,showninfigure13.2, usesa PairSelectGadget
toenforcethat,whentheexecutionstateis ErrorOutOfGasSloadSstore ,theopcodemustbeeither SSTORE or
SLOAD . cb.add_lookup( "Responsibleopcodelookup", Lookup::Fixed{
tag:FixedTableTag::ResponsibleOpcode.expr(), values:[ cb.execution_state().as_u64().expr(),
opcode.expr(), 0.expr(), ], }, ); Figure13.1: zkevm-
circuits/src/evm_circuit/util/common_gadget.rs#48-58 hashey
48ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
impl<F:Field>ExecutionGadget<F>forError00GSloadSstoreGadget<F>{
constNAME:&'staticstr="ErrorOutOfGasSloadSstore";
constEXECUTION_STATE:ExecutionState=ExecutionState::ErrorOutOfGasSloadSstore;
fnconfigure(cb:&mutConstraintBuilder<F>)->Self{ letopcode=cb.query_cell();
letis_sstore=PairSelectGadget::construct( cb, opcode.expr(), OpcodeId::SSTORE.expr(),
OpcodeId::SLOAD.expr(), ); Figure13.2: zkevm-
circuits/src/evm_circuit/execution/error_oog_sload_sstore.rs#48-61
Becausecheckingtheopcode/statecorrespondenceistheresponsibilityofeachexecution
gadget,ifanyexecutiongadgetfailstoproperlyconstraintheopcode,amaliciousprover
canreplaceanotherexecutionstepwiththatgadget'sexecutionstate.
```

In the simplest case, this can lead to a state divergence, but, in general, a malicious prover may have a large amount of control over the resulting state. For example, the `ReturnRevertGadget`, shown in figure 13.3, does not enforce that the opcode is either `RETURN` or `REVERT`. A malicious prover can replace any execution state with a `RETURN_REVERT` state, causing the execution to halt at an arbitrary point, and potentially returning data depending on the values currently on the stack and in memory. If the transaction creates a contract, a malicious prover can replace the code being deployed with values available in memory at other points in the init code's execution.

```
impl<F:Field> ExecutionGadget<F> for ReturnRevertGadget<F> {
    const NAME: &'static str = "RETURN_REVERT";
    const EXECUTION_STATE: ExecutionState = ExecutionState::RETURN_REVERT;
    fn configure(cb: &mut ConstraintBuilder<F>) -> Self {
        let opcode = cb.query_cell();
        cb.opcode_lookup(opcode.expr(), 1.expr());
        let offset = cb.query_cell_phase2();
        let length = cb.query_word_rlc();
        cb.stack_pop(offset.expr());
        cb.stack_pop(length.expr());
        let range = MemoryAddressGadget::construct(cb, offset, length);
        let is_success = cb.call_context(None, CallContextFieldTag::IsSuccess);
        cb.require_boolean("is_success is boolean", is_success.expr());
        // cb.require_equal(
            hashey_49ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

```

```
// "if is_success, opcode is RETURN, if not, opcode is REVERT", // opcode.expr(),
// is_success.expr() * OpcodeId::RETURN.expr()
// + not::expr(is_success.expr()) * OpcodeId::REVERT.expr(), //);
Figure 13.3: zkEVM-circuits/src/evm_circuit/execution/return_revert.rs#55-77
```

In total, we found four gadgets that do not constrain the opcode to match the current execution state:

- `ErrorCodeStoreGadget` (`execution/error_code_store.rs#41-87`)
- `ErrorPrecompileFailedGadget` (`execution/error_precompile_failed.rs#38-85`)
- Additionally, the `ErrorPrecompileFailedGadget` fails to check that the called address is a precompile contract and is missing a correct transition enforcement using the `CommonErrorGadget`.
- `ErrorInvalidCreationCodeGadget` (`execution/error_invalid_creation_code.rs#L35-L73`)
- `ReturnRevertGadget` (`execution/return_revert.rs#L60-L293`)

Suppose a bridge between two blockchains uses the ScrollzkEVM to bridge assets between them. Alice crafts a transaction which, when an opcode such as an `ADD` is instead executed as a `RETURN`, will erroneously withdraw funds from the bridge. She generates a malicious execution trace and submits a zkEVM proof to the bridge, which allows her to drain the bridge of funds. Recommendations Short term, add the missing opcode checks to `ErrorCodeStoreGadget`, `ErrorPrecompileFailedGadget`, `ErrorInvalidCreationCodeGadget`, and `ReturnRevertGadget`. Long term, consider redesigning the way that opcodes map to states. The current design means that any execution gadget that fails to constrain the opcode will cause nondeterministic execution. If, instead, each execution gadget has an `enable` input, and the EVM circuit deterministically selects which gadget(s) have `enable=1`, an underconstrained execution gadget can affect only the behavior of opcodes that are supposed to use that gadget.

14. Nondeterministic execution of `ReturnDataCopyGadget` and `ErrorReturnDataOutOfBoundGadget`
 Severity: High Difficulty: High Type: Data Validation Finding ID: TOB-SCROLL-14 Target: zkEVM-circuits/src/evm_circuit/execution/returndatacopy.rs Description

The gadget that implements the successful execution of the `RETURNDATACOPY` opcode is underconstrained, allowing a malicious prover to successfully execute the opcode when it is in an error condition for particular opcode inputs. This allows the prover to cause state divergence from a correct EVM execution.

Figure 14.1 shows the error gadget implementation that constrains the trace to have at least one true error condition for the `RETURNDATACOPY` opcode. These constraints check overflow conditions on the stack values and their sum.

```
// Check if `data_offset` is Uint64 overflow.
let data_offset_larger_u64 = sum::expr(&data_offset.cells[N_BYTES_U64..]);
let is_data_offset_within_u64 = IsZeroGadget::construct(cb, data_offset_larger_u64);
// Check if `remainder_end` is Uint64 overflow.
let sum = AddWordsGadget::construct(cb, [data_offset, size], remainder_end.clone());
let is_end_u256_overflow = sum.carry().as_ref().unwrap();
let remainder_end_larger_u64 = sum::expr(&remainder_end.cells[N_BYTES_U64..]);
let is_remainder_end_within_u64 = IsZeroGadget::construct(cb, remainder_end_larger_u64);
// Check if `remainder_end` exceeds returndata length.
let is_remainder_end_exceed_len = LtGadget::construct(
    cb, return_data_length.expr(), from_bytes::expr(&remainder_end.cells[..N_BYTES_U64]), );
// Need to check if `data_offset+size` is U256 overflow via `AddWordsGadget` carry. If
// yes, it should be also an error of returndata out of bound.
cb.require_equal(
    "Any of [data_offset>U64::MAX, data_offset+size>U256::MAX, remainder_end>
```

u64::MAX, remainder_end>return_data_length]occurs", or ::expr([//data_offset>u64::MAX hashey
51ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

```
not::expr(is_data_offset_within_u64.expr()), //data_offset+size>U256::MAX  
is_end_u256_overflow.expr(), //remainder_end>u64::MAX  
not::expr(is_remainder_end_within_u64.expr()), //remainder_end>return_data_length  
is_remainder_end_exceed_len.expr(), ]), 1.expr(), ); Figure14.1:  
evm_circuit/execution/error_return_data_oo_bound.rs#L68-L101  
Onthesuccessfulexecutionpath,theseconditionsarenotcheckedtobefalse.Infact,if  
data_offset=WORD_CELL_MAX , size=0 ,and return_data_size<2 32 ,the ReturnDataCopyGadget  
constraintsaresatisfied.Thiscaseisanerrorstatebecause data_offset islargerthan u64::MAX .  
//3.constraintsforcopy:copyoverflowcheck //i.e.,offset+size≤return_data_size  
letin_bound_check=RangeCheckGadget::construct( cb, return_data_size.expr() -  
(from_bytes::expr(&data_offset.cells)+from_bytes::expr(&size.cells)), ); //4.memorycopy  
//Constructmemoryaddressinthedestination(memory)towhichwecopymemory.  
letdst_memory_addr=MemoryAddressGadget::construct(cb,dest_offset,size);  
//Calculatethenextmemorysizeandthegascostforthismemory  
//access.Thisalsoaccountsforthedynamicgasrequiredtocopybytesto //memory.  
letmemory_expansion=MemoryExpansionGadget::construct(cb, [dst_memory_addr.address()]);  
letmemory_copier_gas=MemoryCopierGasGadget::construct( cb, dst_memory_addr.length(),  
memory_expansion.gas_cost(), ); letcopy_rwc_inc=cb.query_cell();  
cb.condition(dst_memory_addr.has_length(),|cb|{ cb.copy_table_lookup( last_callee_id.expr(),  
CopyDataType::Memory.expr(), cb.curr.state.call_id.expr(), CopyDataType::Memory.expr(),  
return_data_offset.expr()+from_bytes::expr(&data_offset.cells),  
return_data_offset.expr()+return_data_size.expr(), dst_memory_addr.offset(),  
dst_memory_addr.length(), hashey 52ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
0.expr(),//forRETURNDATACOPYrLc_accis0 copy_rwc_inc.expr(), ); });  
cb.condition(not::expr(dst_memory_addr.has_length()),|cb|{ cb.require_zero(  
"ifnobytestocopy,copytablerwcinc=0", copy_rwc_inc.expr(), ); }); Figure14.2:  
evm_circuit/execution/returndatacopy.rs#L99-L141  
Insum,theprovercoulddecidewhethertheexecutionwouldcorrectlyhaltwiththe  
ErrorReturnDataOutOfBoundGadget errororifitwouldsuccessfullyexecutethe RETURNDATACOPY opcode.  
ExploitScenario Amaliciousprovergeneratesandsubmitsaproofofexecutionforatransactioninvolvinga  
RETURNDATACOPY withparticulararguments.Duetothemissingvalidationsonthe  
successfulexecutionstate,theprovercouldchoosetosuccessfullyexecutetheopcode,or  
halttheexecution,leadingtostatedivergenceandlossoffunds. Recommendations  
Shortterm,addconstraintstoensurethatthesuccessfulexecutionstateisdisjointfrom  
theerrorexecutionstate.  
Longterm,investigateothererrorstatesandtheirassociatedopcodeimplementationsto  
guaranteethattheirexecutionstateisdisjointandcannotbechosenbyamaliciousprover. hashey  
53ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

15.ManyRWcounterupdatesaremagicnumbers Severity:InformationalDifficulty:N/A
Type:DataValidationFindingID:TOB-SCROLL-15 Target: zkevm-circuits/src/evm_circuit/execution/
Description ThezkEVMcircuitchecksmemoryreadandwriteoperationsinatransaction'sexecution
tracebyperforminglookupsintotheRWtable.Withinthecircuit,updatestothe
statevariable,whichtracksthetotalnumberofread/writeoperations,arefrequentlyspecified
withamanualcount.Thatmanualprocessiserror-proneanddifficulttocheck,anditcan
bereplacedwithacalculatedvalueinallcaseswehaveseen. Theread-
writeconsistencychecksinthekEVMcircuitrequiretheoverallblocktohavea
correctcountofthetotalnumberoflookupsintotheRWtable.Ifthatcountisincorrect,a
maliciousprovercaninsertextraneouswriteoperationsintothetableandchoosean
arbitraryresultforanymemoryread(seeTOB-SCROLL-8fordetailedexplanation).Each
executiongadgetisindividuallyresponsibleforcreatinga StepStateTransition that
enforcesthecorrectupdateoftherw_counter fieldof StepState .Forexample,figure 15.1showsthe
StepStateTransition forthe AddSubGadget .Therearethreerw lookupscausedbythe stack_pop()
calls,andthustherw_counter_field issetto Delta(3) ,representinganincreasebythree.
//ADD:Popaandbfromthestack,pushconthestack //SUB:Popcandbfromthestack,pushaonthestack
cb.stack_pop(select::expr(is_sub.expr().0,c.expr(),a.expr())); cb.stack_pop(b.expr());
cb.stack_push(select::expr(is_sub.expr().0,a.expr(),c.expr())); //Statetransition
letstep_state_transition=StepStateTransition{ rw_counter:Delta(3.expr()),
program_counter:Delta(1.expr()), stack_pointer:Delta(1.expr()), gas_left:Delta(-
OpcodeId::ADD.constant_gas_cost().expr()), ..StepStateTransition::default() };
letsame_context=SameContextGadget::construct(cb,opcode,step_state_transition); Figure15.1:The

```
rw_counter update in AddSubGadget ( zkevm-circuits/src/evm_circuit/execution/add_sub.rs#51-65 )
hashey 54ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

However, many execution gadgets have much more complicated `rw_counter` updates, which are difficult to check for correctness.

To illustrate this complexity, consider figures 15.2 and 15.3, which show `ErrorInvalidOpcodeGadget` and `ErrorWriteProtectionGadget`. Each of them uses the `CommonErrorGadget`

, which has an `RwCounterDelta` as the third parameter. However, `ErrorInvalidOpcodeGadget` does not seem to contain any `RwLookups` at all, but provides the value 2, while `ErrorWriteProtectionGadget` seems to have either one or four `RwLookups` depending on the value of `is_call`, yet provides a 0 to

```
CommonErrorGadget . fn configure (cb: &mut ConstraintBuilder<F>) -> Self { let opcode = cb.query_cell();
cb.add_lookup( "Responsible opcode lookup", Lookup::Fixed {
tag: FixedTableTag::ResponsibleOpcode.expr(), values: [ Self::EXECUTION_STATE.as_u64().expr(),
opcode.expr(), 0.expr(), ], }, );
```

let common_error_gadget = CommonErrorGadget::construct(cb, opcode.expr(), 2.expr()); Figure 15.2:

```
ErrorInvalidOpcodeGadget ( zkevm-circuits/src/evm_circuit/execution/error_invalid_opcode.rs#27-41 )
```

```
fn configure (cb: &mut ConstraintBuilder<F>) -> Self { ... // Lookup values from stack if opcode is call
```

```
// Precondition: If there's a StackUnderflow CALL, it's handled before this error
```

```
cb.condition(is_call.expr(), |cb| { cb.stack_pop(gas_word.expr());
```

```
cb.stack_pop(code_address_word.expr()); cb.stack_pop(value.expr());
```

```
// cb.require_zero("value of call is not zero", // is_value_zero.expr()); });
```

```
// current call context is read only
```

```
cb.call_context_lookup(false.expr(), None, CallContextFieldTag::IsStatic, 1.expr());
```

```
// constrain not root call at least one previous static call preset. cb.require_zero(
```

```
"ErrorWriteProtection only happens in internal call", hashey
```

```
55ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
cb.curr.state.is_root.expr(), );
```

let common_error_gadget = CommonErrorGadget::construct(cb, opcode.expr(), 0.expr()); Figure 15.3:

```
ErrorWriteProtectionGadget ( evm_circuit/execution/error_write_protection.rs#33-80 ) In
```

`ErrorInvalidOpcodeGadget`, there are in fact two total `RwLookups`; however, unintuitively, they occur inside

`CommonErrorGadget` itself, as shown in figure 15.4. Thus, any caller of `CommonErrorGadget`

effectively must add two to the value of `rw_counter_delta`.

```
pub (crate) fn construct_with_last callee_return_data (cb: &mut ConstraintBuilder<F>,
```

```
opcode: Expression<F>, rw_counter_delta: Expression<F>, return_data_offset: Expression<F>,
```

```
return_data_length: Expression<F>, ) -> Self { cb.opcode_lookup(opcode.expr(), 1.expr());
```

```
let rw_counter_end_of_reversion = cb.query_cell(); // current call must be failed.
```

```
cb.call_context_lookup(false.expr(), None, CallContextFieldTag::IsSuccess, 0.expr());
```

```
cb.call_context_lookup( false.expr(), None, CallContextFieldTag::RwCounterEndOfReversion,
```

```
rw_counter_end_of_reversion.expr(), ); Figure 15.4: Two RwLookups inside CommonErrorGadget ( zkevm-
```

```
circuits/src/evm_circuit/util/common_gadget.rs#1019-1038 ) The case of ErrorWriteProtectionGadget
```

is somewhat more complex but illustrates a useful alternative to manually counting lookups. `CommonErrorGadget`

is called with an `rw_counter_delta` value of 0. One might expect that this is incorrect; it should count the

two lookups inside `CommonErrorGadget`, plus one unconditional lookup and three

conditional lookups outside. However, upon closer inspection, `CommonErrorGadget` only uses

`rw_counter_delta` at all when `curr.state.is_root` is true. `ErrorWriteProtectionGadget`

can occur only inside of a static call, and the root call of a transaction is never a static call—

therefore, that case is never active and the value of `rw_counter_delta` can be set to a dummy value, in this case, 0

. Instead, `RestoreContextGadget` handles the `RwCounter` update, basing it on a call to

```
ConstraintBuilder::rw_counter_offset , as shown in figure 15.5. Since hashey
```

```
56ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
rw_counter_offset is updated automatically in each call to ConstraintBuilder::rw_lookup
```

```
, that count is correct by construction. let rw_counter_offset = cb.rw_counter_offset()
```

```
+ subsequent_rw_lookups + not::expr(is_success.expr()) * cb.curr.state.reversible_write_counter.expr();
```

```
// Do step state transition cb.require_step_state_transition(StepStateTransition {
```

```
rw_counter: Delta(rw_counter_offset), call_id: To(caller_id.expr()),
```

```
is_root: To(caller_is_root.expr()), is_create: To(caller_is_create.expr()),
```

```
code_hash: To(caller_code_hash.expr()), program_counter: To(caller_program_counter.expr()),
```

```
stack_pointer: To(caller_stack_pointer.expr()), gas_left: To(gas_left),
```

```
memory_word_size: To(caller_memory_word_size.expr()),
```

```
reversible_write_counter: To(reversible_write_counter), log_id: Same, }); Figure 15.5: The rw_counter
```

```
update in RestoreContextGadget ( zkevm-circuits/src/evm_circuit/util/common_gadget.rs#185-202 )
```

In general, the process of counting the number of `RwLookups` in any given gadget is both

subtle and tedious when done manually, but cases that use `ConstraintBuilder::rw_counter_offset`

to determine that offset are effectively trivial when checking for correctness. `CreateGadget`

```
hasverycomplexlogic,including threedifferentconditionalcalls to require_step_state_transition
and a large number of RW lookups. However, each StepStateTransition computes its RW counter
update automatically, as illustrated in figure 15.6. cb.condition(not::expr(is_precheck_ok.expr()),|cb|{
//Save caller's call state for field_tag in [ CallContextFieldTag::LastCalleeId,
CallContextFieldTag::LastCalleeReturnDataOffset, CallContextFieldTag::LastCalleeReturnDataLength, ]
{ cb.call_context_lookup(true.expr(),None,field_tag,0.expr()); }
cb.require_step_state_transition(StepStateTransition{ rw_counter:Delta(cb.rw_counter_offset()),
program_counter:Delta(1.expr()), stack_pointer:Delta(2.expr()+IS_CREATE2.expr()),
memory_word_size:To(memory_expansion.next_memory_word_size()), //-
(Reversible)WriteTxAccessListAccount(ContractAddress) hashey
57ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
reversible_write_counter:Delta(1.expr()), gas_left:Delta(-gas_cost.expr()),
..StepStateTransition::default() }); }; Figure 15.6: One possible rw_counter update in CreateGadget (
zkevm-circuits/src/evm_circuit/execution/create.rs#337-357 ) By contrast, CallOpGadget
has three different StepStateTransition s, each of which
```

```
has a manually constructed RW counter offset. This includes the StepStateTransition
shown in figure 15.7, which counts a grand total of 41 RW lookups plus four conditional
offsets, all of which need to be verified to be correct. let transfer_rwc_delta=
is_call.expr()*not::expr(transfer.value_is_zero.expr()*2.expr()); let rw_counter_delta=41.expr()
+is_call.expr()*1.expr() +transfer_rwc_delta.clone() +is_callcode.expr()
+is_delegatecall.expr()*2.expr(); cb.require_step_state_transition(StepStateTransition{
rw_counter:Delta(rw_counter_delta), call_id:To(callee_call_id.expr()), is_root:To(false.expr()),
is_create:To(false.expr()), code_hash:To(call_gadget.phase2_callee_code_hash.expr()),
gas_left:To(callee_gas_left), //For CALL opcode, `transfer` invocation has two account write if value is not
//zero. reversible_write_counter:To(transfer_rwc_delta), ..StepStateTransition::new_context() });
Figure 15.7: One rw_counter update in CallOpGadget ( zkevm-
circuits/src/evm_circuit/execution/callop.rs#440-458 ) Recommendations Short term, replace magic-
number RW counter updates with computed values, such as those provided by
```

```
ConstraintBuilder::rw_counter_offset . Long term, consider redesigning the API for building
StepStateTransition s. Since all RW lookups are performed via the ConstraintBuilder
API, update to simple counter-style
state variables, such as the stack pointer and the RW counter, can typically be computed
rather than manually specified. If the easy-to-calculate fields are always computed, that
core computation can be checked for correctness; as a result, all uses will be correct by construction. hashey
58ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
16. NativePCSAccumulationDecidersAcceptAnEmptyVector Severity:Medium Difficulty:Low
Type:DataValidationFindingID:TOB-SCROLL-16 Target: snark-verifier/src/pcs/{kzg,ipa}/decider.rs
Description Both the KZG and IPA native decide_all implementations accept an empty vector of
accumulators. This can allow an attacker to bypass verification by submitting an empty vector. fn decide_all(
dk:&Self::DecidingKey, accumulators:Vec<IpaAccumulator<C,NativeLoader>>, )->bool{ !accumulators
.into_iter().any(|accumulator|!Self::decide(dk,accumulator)) } Figure 16.1: snark-
verifier/src/pcs/kzg/decider.rs#L54-L69
```

```
This function contrasts with the EVM loader implementation that asserts that the accumulator vector is non-empty:
fn decide_all( dk:&Self::DecidingKey,
mut accumulators:Vec<KzgAccumulator<M::G1Affine,Rc<EvmLoader>>>, )->Result<(),Error>{ assert!
(!accumulators.is_empty()); Figure 16.2: snark-verifier/src/pcs/kzg/decider.rs#L120-L124
```

```
Exploit Scenario An attacker is able to control the arguments to decide_all and passes an empty vector,
causing the verification function to accept an invalid proof. Recommendations
Short term, add an assertion that validates that the vector is non-empty.
Long term, add negative tests for verification and validation functions, ensuring that wrong
or invalid arguments are not accepted. hashey 59ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

```
17. TheError00GSloadSstoreandtheError00GLoggadgetsHaveRedundantTableLookups
Severity:Informational Difficulty:N/A Type:DataValidationFindingID:TOB-SCROLL-17 Target: zkevm-
circuits/src/evm_circuit/execution/{error_oog_sload_sstore,error_oog_log}.rs Description Both the
Error00GSloadSstore and the Error00GLog gadgets do an RW table lookup to
check whether the current call is within a static context. However, the lookup result is not
used in any subsequent constraint, making the lookup redundant. //constraint in static call
let is_static_call=cb.call_context(None,CallContextFieldTag::IsStatic);
//cb.require_zero("is_static_call is false in LOGN",is_static_call.expr()); Figure 17.1:
evm_circuit/execution/error_oog_log.rs#L53-L55 The commented-
out constraint would provide a clear stated distinction between the Error00GLogGadget error case and the
ErrorWriteProtectionGadget , preventing an attacker from arbitrarily choosing one of the error states at will.
As far as we know, in this case, these two different error execution states do not translate
```

to diverging EVM states; thus, this finding's severity is informational. Recommendations Short term, decide whether to remove the RW table lookup or to uncommit the non-static environment constraint in both the Error006SloadStore and Error006Log gadgets. Investigate all commented-out constraints and remove them from the code base, or enable them if they are necessary. hashey 60ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

18. The State circuit does not enforce transaction receipt constraints Severity: Informational Difficulty: N/A Type: Data Validation Finding ID: TOB-SCROLL-18 Target: zkevm-

circuits/src/state_circuit/constraint_builder.rs Description The implementation of the State circuit does not enforce transaction receipt constraints. Currently, these have an unsatisfiable constraint ($1=0$), and the function that implements them, `build_tx_receipt_constraints`, is not called in the `ConstraintBuilder::build` function. `fn build_tx_receipt_constraints(&mut self, q: &Queries<F>){ // TODO: implement TxReceipt constraints self.require_equal("TxReceipt rows not implemented", 1.expr(), 0.expr()); self.require_equal("state_root is unchanged for TxReceipt", q.state_root(), q.state_root_prev(),); self.require_zero("value_prev_column is 0 for TxReceipt", q.value_prev_column(),); }` Figure 18.1:

state_circuit/constraint_builder.rs#L511-L524 Recommendations

Short term, implement the transaction receipt constraints and add them to the `constraint_builder.build` function. Long term, enable the `dead_code` compiler lint by removing the `#![allow(dead_code)]` line in `zkevm-circuits/src/lib.rs` and fix all warnings. hashey 61ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

20. The EXP opcode has an unused witness Severity: Informational Difficulty: N/A

Type: Data Validation Finding ID: TOB-SCROLL-20 Target: zkevm-circuits/src/evm_circuit/execution/exp.rs

Description The EXP opcode defines a witness that is used only in a constraint requiring its value to be zero. The constraint label suggests that it was used to validate the `base_sq` witness value at some point in the code development, but this is now done in the exponentiation table circuit.

`let zero_rlc = cb.query_word_rlc(); cb.require_zero("base*base+c==base^2(c==0)", sum::expr(&zero_rlc.cells),);` Figure 20.1: `evm_circuit/execution/exp.rs#L93-L97` Recommendations

Short term, remove the `zero_rlc` variable and its constraint from the EXP opcode gadget. hashey 62ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

21. The `bn_to_field` functions silently truncate big integers Severity: Low Difficulty: Low

Type: Data Validation Finding ID: TOB-SCROLL-21 Target: misc-precompiled-circuit/src/utils/mod.rs

Description The `bn_to_field` function converts arbitrary length integers into a field element.

However, if the byte representation of the integers is larger than 64 bytes, the big integer bytes will be silently truncated. This means that any two integers with the same 512 least significant bits will lead to the same field element. `pub fn bn_to_field<F: FieldExt>(bn: &BigUint) -> F { let mut bytes = bn.to_bytes_le(); bytes.resize(64, 0); F::from_bytes_wide(&bytes.try_into().unwrap()) }` Figure 21.1: `src/utils/mod.rs#L10-L15`

Instead, the functions should check whether the big integer fits into the field capacity by using the `F::capacity` constant. This would guarantee a faithful representation of the big integer into the field element and a successful reversion back to the `BigUint` type. Note that the `from_bytes_wide` function will also reduce the element modulo the field

orders so that it is represented as a field element. Exploit Scenario

An attacker provides two big integers with the same 512 least significant bits to the `bn_to_field` function, causing it to return the same element. When these elements are

used in future operations, they will lead to the same result, even though they were different. Recommendations

Short term, add documentation to the function explaining the intended behavior; add

checks that validate that the big integer is representable in the chosen field. hashey 63ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

22. The `field_to_bn` function depends on implementation-specific details of the underlying field

Severity: Low Difficulty: High Type: Data Validation Finding ID: TOB-SCROLL-22 Target: misc-precompiled-

circuit/src/utils/mod.rs Description The implementation of the `field_to_bn` function calls the `to_repr` function on the value of the input `f` when constructing the little-endian binary representation of the input `f`.

`pub fn field_to_bn<F: FieldExt>(f: &F) -> BigUint { let bytes = f.to_repr(); BigUint::from_bytes_le(bytes.as_ref()) }` Figure 22.1: The implementation of `field_to_bn` expects `to_repr` to return a little-endian representation of the value of `f`. (`src/utils/mod.rs#L5-L8`)

However, according to the documentation of the `PrimeField` trait, the `endianness` returned by

`PrimeField::to_repr` is implementation-dependent and may be different depending on the underlying field.

`/// Converts an element of the prime field into the standard byte representation for /// this field. ///`

`/// The endianness of the byte representation is implementation-specific. Generic`

`/// encodings of field elements should be treated as opaque. fn to_repr(&self) -> Self::Repr;`

Figure 22.2: The value returned by `to_repr` is implementation-dependent and should be

treated as opaque by the user. Exploit Scenario The `field_to_bn` function is reused with a scalar field `F` that uses a different internal representation of the elements of `F`

.Theresultingbigintegermightnotcorrespondtothesamefieldelement. Recommendations
Shortterm,implementafunctionthatassuredlyreturnsalittle-endianrepresentationofthefieldelement.
hasheyeye 64ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Longterm,reviewtheuseofthird-partyAPIstoensurethatthecodebasedoesnotdepend
ontheinternalrepresentationofdata. hasheyeye 65ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

23.Thevaluesofthebytecodetabletagcolumnarenotconstrainedtobe HEADERorBYTE
Severity:InformationalDifficulty:High Type:DataValidationFindingID:T0B-SCROLL-23 Target: zkevm-
circuits/src/bytecode_circuit/circuit.rs Description Thebytecodetablehasacolumnthatindicatesthe TAG
ofeachrow.Currently,the TAG cellsareassignedonlya HEADER ora BYTE
value.However,thecircuitdoesnotconstrain the TAG
valueofeachrowtoacceptonlythesevalues.Thismissingconstraintdoesnot
causeadirectsoundnessissuebecauseofotherindirectconstraintsandhowthebytecode
circuitisimplemented,butfuturecoderefactoringscouldcausetheissuebecome exploitable.
Thebytecodetablecontainsthesesetofbytecodesthatareexecutedinablock.Foreach bytecode,thetablecontainsa
HEADER row,followedby BYTE rowscorrespondingtoeach byteofthebytecode,andafinal HEADER row.
Thecircuitimposesconstraintsforeachtypeofrow(e.g.,figure23.1showshowa HEADER
rowisconstrainedtohaveanindexcolumnvalueof 0),andfortransitionsbetweentwo
rows(e.g.,transitioningfroma HEADER rowtoa BYTE row,thelengthcolumnmuststaythe same).
//Whenis_header→ //assertcur.index==0 //assertcur.value==cur.length
meta.create_gate("Headerrow",|meta|{ letmutcb=BaseConstraintBuilder::default(); cb.require_zero(
"cur.index==0", meta.query_advice(bytecode_table.index,Rotation::cur()),); cb.require_equal(
"cur.value==cur.length", meta.query_advice(bytecode_table.value,Rotation::cur()),
meta.query_advice(length,Rotation::cur()),); cb.gate(and::expr(vec![hasheyeye
66ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

meta.query_fixed(q_enable,Rotation::cur()), not::expr(meta.query_fixed(q_last,Rotation::cur())),
is_header(meta),)) }); Figure23.1: zkevm-circuits/src/bytecode_circuit/circuit.rs#L178-L200
Tocheckwhetherarowisa HEADER ora BYTE row,theimplementationperformssteps
thatrelyonparticularassumptions: • Itimplicitlyassumesthatthe enum valuecorrespondingto HEADER is 0
andtheone to BYTE is 1 .Thisassumptioncanbeprokeninthefutureifadeveloperaddsan extra enum
fieldonthe BytecodeFieldTagenum . • ItusesBooleanoperators and::expr , not::expr
onthebytecodetagvalues. TheseoperatorsmustoperateonlyonBooleanvalues;otherwise,theywillreturnan
unexpectedvalue. • Itgatestheconstraintswithconjunctionsresultingfromthe and::expr operator;if
alookupisguardedbytheconjunctionofnon-Booleanvalues,thethatvalueis
lookedupwillbeascaledversionoftheintendedvalue. letis_byte_to_byte=|meta:&mutVirtualCells<F>|{
and::expr(vec![meta.query_advice(bytecode_table.tag,Rotation::cur()),
meta.query_advice(bytecode_table.tag,Rotation::next()),]) };
letis_header=|meta:&mutVirtualCells<F>|{
not::expr(meta.query_advice(bytecode_table.tag,Rotation::cur())) }; letis_byte=
|meta:&mutVirtualCells<F>|meta.query_advice(bytecode_table.tag, Rotation::cur()); Figure23.2:
zkevm-circuits/src/bytecode_circuit/circuit.rs#L125-L137
Thesoundnessofallthesestepsandimplementationdetailsrelyonthebytecodetagvalue
beingBoolean.However,theimplementationdoesnothaveaconstraintvalidatingthatthe TAG
valuesare,infact,Boolean. Ifamaliciousproverweretoprovideanon-Booleanvalue,sincethe not::expr and
and::expr functionsoperateundertheassumptionthattheirinputvaluesareBoolean,
thecircuitwillhavesoundnessissues. hasheyeye 67ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Oneavenueofexploitationisonthe push_data_size_table_lookup onthebytecode
table:thislookupisgatedonthe is_byte(meta) constraint,causingittobeimplicitly
scaledtoadifferentlookupifthe is_byte(meta) resultisnon-Boolean.Thiswouldallow
amaliciousprovertoobtainthewrong push_data_size fromthe push_data_size_table_lookup
tableandprovideanincorrectbytecodewithrespectto the is_code column.Inotherwords,thedatapushedina
PUSH* opcodecouldbemarked ascode,thenallowingtheEVMexecutiontofollowanexecutionflowincompatiblewith
EVMsemantics. meta.lookup_any("push_data_size_table_lookup(cur.value,cur.push_data_size)", |meta|{
letenable=and::expr(vec![meta.query_fixed(q_enable,Rotation::cur()),
not::expr(meta.query_fixed(q_last,Rotation::cur())) , is_byte(meta),]); letlookup_columns=vec!
[value,push_data_size]; letmutconstraints=vec![]; foriin0..PUSH_TABLE_WIDTH{ constraints.push((
enable.clone()*meta.query_advice(lookup_columns[i], Rotation::cur()),
meta.query_fixed(push_table[i],Rotation::cur()),)) } constraints },); Figure23.3: zkevm-
circuits/src/bytecode_circuit/circuit.rs#L220-L241 However,arowwithanon-Boolean TAG
actuallysatisfiesboththe is_header(meta) and is_byte(meta)
constraints.Thus,foranattacker tobesuccessful,theywouldhaveto
satisfyanunsatisfiablesetofconstraintsonthe index column: • is_header :requires cur.index==0 •
is_header_to_byte :requires next.index==0 • is_byte_to_byte :requires next.index==cur.index+1

There exists another avenue of exploiting the missing constraint that requires the TAG value to equal BYTE or HEADER. If a malicious prover were able to inject rows in the table with a different TAG value, they would be able to disable the BYTE-TO-BYTE, BYTE-TO-HEADER, HEADER-TO-HEADER, and HEADER-TO-BYTE transition constraints. [hashey 68 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC](#)

As an example, if between two HEADER rows there existed a row different from HEADER or BYTE, the HEADER-TO-HEADER transition gate would always be false. This exploits a scenario is unexploitable for the same reason as the previous exploit. However, while correctly enforcing the `is_header`, `is_header_to_byte`, `is_byte_to_byte`, `is_header_to_header`, `is_byte`, and `is_byte_to_header` transition constraints to accept only the HEADER and BYTE values would prevent the `push_data_size_table_lookup` exploit, it would not prevent the row-to-row transition constraints from being broken. For a complete fix, it is necessary to constrain the TAG value to be one of HEADER or BYTE. [Exploit Scenario Anticipating a future addition to the TAG enum](#), a developer decides to implement the `is_header`, `is_byte`, and transition selectors by requiring that the TAG cell value equals the desired enum value. If they omit the check that the TAG value must be restricted to the enum's value set, a malicious prover would be able to break the transition constraints by inserting a row with a tag value different from HEADER or BYTE. [Recommendations](#) Short term, require that the TAG column is Boolean in the constraints system; make the `BytecodeFieldTag` values explicitly 0 and 1 by defining the enum as follows: `pub enum BytecodeFieldTag { ///Headerfield Header=0, ///Bytefield Byte=1, }` [Figure 23.4: Explicit enum definition](#) Document which constraints need to be changed in case the `BytecodeFieldTag` enum is extended. Long term, add strict type to the Boolean functions in `gadgets/src/util.rs`. These functions, as their documentation states, should operate only on Boolean values. Enforcing this in the type system would allow cases where this assumption is violated to be found and would prevent potential soundness issues. [hashey 69 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC](#)

24. Unconstrained columns on the bytecode HEADER rows Severity: Informational Difficulty: N/A
Type: Data Validation Finding ID: TOB-SCROLL-24 Target: `zkevm-circuits/src/bytecode_circuit/{circuit/to_poseidon_hash.rs, circuit.rs}` Description The bytecode table HEADER rows have two unconstrained columns, `is_code` and `field_input`, on the Poseidon bytecode extended columns. The lack of constraints on these columns does not seem to pose any soundness issue, but constraining these columns would serve as a defense-in-depth, preventing the circuit's flexibility from allowing a malicious prover to exploit a soundness issue if a vulnerability is introduced in the future. [Figure 24.1](#) shows the HEADER row constraints, and no constraint related to the `is_code` column. [meta.create_gate\("Header row", |meta| { let mut cb = BaseConstraintBuilder::default\(\); cb.require_zero\("cur.index==0", meta.query_advice\(bytecode_table.index, Rotation::cur\(\)\), \); cb.require_equal\("cur.value==cur.length", meta.query_advice\(bytecode_table.value, Rotation::cur\(\)\), meta.query_advice\(length, Rotation::cur\(\)\), \); cb.gate\(and::expr\(vec!\[meta.query_fixed\(q_enable, Rotation::cur\(\)\), not::expr\(meta.query_fixed\(q_last, Rotation::cur\(\)\)\) \], is_header\(meta\), \)\) }\);](#) [Figure 24.1: zkevm-circuits/src/bytecode_circuit/circuit.rs#L179-L201](#) [Recommendations](#) Short term, add constraints for the `is_code` and the `field_input` rows in the HEADER rows of the bytecode table. [hashey 70 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC](#)

Long term, document all table constraints and ensure that each type of row constrains all columns. [hashey 71 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC](#)

25. `decompose_limb` does not work as intended Severity: Informational Difficulty: N/A
Type: Data Validation Finding ID: TOB-SCROLL-25 Target: `misc-precompiled-circuit/src/circuits/modexp.rs` Description The for loop within `decompose_limb` requires `bool_limbs` to contain at least 31 elements to be correctly indexed from 0 to 30. Furthermore, if `limb_size` is large enough, then the `truncate` operation does not grow `bool_limbs` to the correct size, as `to_radix_le` produces a minimal `Vec` without any trailing zeroes. `let mut bool_limbs = field_to_bn(&limb.value).to_radix_le(2); bool_limbs.truncate(limb_size); bool_limbs.reverse(); let mut v = F::zero(); for i in 0..27 { let l0 = F::from_u128(bool_limbs[i] as u128); let l1 = F::from_u128(bool_limbs[i+1] as u128); let l2 = F::from_u128(bool_limbs[i+2] as u128); let l3 = F::from_u128(bool_limbs[i+3] as u128);` [Figure 25.1: misc-precompiled-circuit/src/circuits/modexp.rs#L514-L522](#) Additionally, the Boolean `limbs` are not properly constrained to be Boolean, but this is mentioned in a "TODO" comment. Overall, it can be concluded that the `decompose_limb` needs further development, but its intended purpose and usage within `mod_exp` is clear. [Recommendations](#) Short term, correctly implement `decompose_limb`. This will allow for proper testing of the `mod_exp` circuit. [hashey 72 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC](#)

26. Zero modulus will cause a panic Severity: Medium Difficulty: Low Type: Data Validation Finding ID: TOB-SCROLL-26 Target: `misc-precompiled-circuit/src/circuits/modexp.rs` Description According to EVM specifications, if the modulus is zero, then the result of `mod_exp` is zero

regardless of the input. The current codereferences successful to `mod_mult` with the passed-in modulus, but the `mod_mult` function computes a quotient that will panic.

```
let bn_quotient = bn_mult.clone().div(bn_modulus.clone()); // div_rem
```

Figure 26.1: `misc-precompiled-circuit/src/circuits/modexp.rs#L470` This results in differing behavior between the scroll `mod_exp` precompile and the standard EVM precompile, which may cause some existing systems that depend on this behavior to not work as intended. Recommendations Short term, correctly handle the zero modulus case of `mod_exp`. Add tests to the `mod_exp` circuit, including some that exercise its edge cases: the zero exponent case and the zero modulus case. `hasheyev33ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC`

27. The `ConstraintBuilder::condition` API is dangerous Severity: Informational Difficulty: High Type: Data Validation Finding ID: TOB-SCROLL-27 Target: Several files Description The `ConstraintBuilder` implements several useful ways of constructing constraints. One case is when a constraint should be added and conditioned by a particular value. If the value is true, the constraints must be satisfied; otherwise, they do not need to be satisfied. However, a problem arises if a developer forgets to consider that a new `ConstraintBuilder` function is called within the context of a `condition`. All functions in the `ConstraintBuilder` API must consider the case where they are being called from inside a conditioned scope. If these functions add constraints or change values irrespective of the condition value, they will lead to unintended results. As an example, the `opcode_lookup` function updates the `program_counter_offset` regardless of the current condition value.

```
pub(crate) fn opcode_lookup(&mut self, opcode: Expression<F>, is_code: Expression<F>){
    self.opcode_lookup_at( self.curr.state.program_counter.expr()+self.program_counter_offset.expr(),
    opcode, is_code, );
    self.program_counter_offset+=1; }
```

Figure 27.1: `evm_circuit/util/constraint_builder.rs#608-615` When used in a condition context, the `program_counter_offset` will be incremented irrespective of the condition value:

```
const NAME: &'static str = "STOP";
const EXECUTION_STATE: ExecutionState = ExecutionState::STOP;
fn configure(cb: &mut ConstraintBuilder<F>) -> Self {
    let code_length = cb.query_cell();
    cb.bytecode_length(cb.curr.state.code_hash.expr(), code_length.expr());
}
```

Figure 27.2: `src/evm_circuit/execution/stop.rs#33-45`

The provided argument to the `ConstraintBuilder::condition` function must also be ensured to be Boolean. Certain functions assume this, and they would have unexpected results otherwise:

```
pub(crate) fn stack_pop(&mut self, value: Expression<F>){
    self.stack_lookup(false.expr(), self.stack_pointer_offset.clone(), value);
    self.stack_pointer_offset = self.stack_pointer_offset.clone() + self.condition_expr();
}
pub(crate) fn stack_push(&mut self, value: Expression<F>){
    self.stack_pointer_offset = self.stack_pointer_offset.clone() - self.condition_expr();
    self.stack_lookup(true.expr(), self.stack_pointer_offset.expr(), value);
}
```

Figure 27.3: `evm_circuit/util/constraint_builder.rs#1160-1169` The `ConstraintBuilder::gate` function is another dangerous pattern that should be reconsidered and documented. In its current state, the function clones all constraints and gates them with the provided selector, returning these new gated constraints. It does not change the current constraints, which might be an interpretation that a new developer might have about the function. We have not seen incorrect usage of this particular pattern. One way of at least ensuring that the returning set of constraints is used is by adding the `#[must_use]` attribute to the function. Recommendations Short term, redesign the `ConstraintBuilder` API, especially with respect to the `condition` function. Add new Rust typesto ensure that the condition expression is Boolean. Add the `#[must_use]` attribute to the `ConstraintBuilder::gate` function. `hasheyev75ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC`

28. The `EXTCODECOPY` opcode implementation does not work when the account address does not exist Severity: Informational Difficulty: N/A Type: Data Validation Finding ID: TOB-SCROLL-28 Target: `zkevm-circuits/src/evm_circuit/execution/extcodecopy.rs` Description The current implementation of the `EXTCODECOPY` opcode does not consider the case where the account address does not exist. This is documented in a code comment, so the Scroll team should be aware of it. `// TODO: If external_address doesn't exist, we will get code_hash=0. With // this value, the bytecode_length lookup will not work, and the copy // from code_hash=0 will not work. We should use EMPTY_HASH when // code_hash=0.`

```
cb.bytecode_length(code_hash.expr(), code_size.expr());
```

Figure 28.1: `zkevm-circuits/src/evm_circuit/execution/extcodecopy.rs#84-88` Recommendations Short term, implement the missing functionality. Add tests to ensure its correctness. Long term, look for all "TODO" items in the code base and triage them into an organized

A. Vulnerability Categories

The following tables describe the vulnerability categories, severity levels, and difficulty levels used in this document.

Vulnerability Categories	Category Description
Access Controls	Insufficient authorization or assessment of rights
Auditing and Logging	Insufficient auditing of actions or logging of problems
Authentication	Improper identification of users
Configuration	Misconfigured servers, devices, or software components
Cryptography	Breach of system confidentiality or integrity
Data Exposure	Exposure of sensitive information
Data Validation	Improper reliance on the structure or values of data
Denial of Service	A system failure with an availability impact
Error Reporting	Insecure or insufficient reporting of error conditions
Patching	Use of an outdated software package or library
Session Management	Improper identification of authenticated users
Testing	Insufficient test methodology or test coverage
Timing	Race conditions or other order-of-operations flaws
Undefined Behavior	Undefined behavior triggered within the system

hashey
77ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Severity Levels

Severity Levels	Severity Description
Informational	The issue does not pose an immediate risk but is relevant to security best practices.
Undetermined	The extent of the risk was not determined during this engagement.
Low	The risk is small or is not one the client has indicated is important.
Medium	User information is at risk; exploitation could pose reputational, legal, or moderate financial risks.
High	The flaw could affect numerous users and have serious reputational, legal, or financial implications.

Difficulty Levels

Difficulty Levels	Difficulty Description
Undetermined	The difficulty of exploitation was not determined during this engagement.
Low	The flaw is well known; public tools for its exploitation exist or can be scripted.
Medium	An attacker must write an exploit or will need in-depth knowledge of the system.
High	An attacker must have privileged access to the system, may need to know complex technical details, or must discover other weaknesses to exploit this issue.

hashey
78ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

B. Code Maturity Categories

The following tables describe the code maturity categories and rating criteria used in this document.

Code Maturity Categories

Code Maturity Categories	Category Description
Arithmetic	The proper use of mathematical operations and semantics
Complexity Management	The presence of clear structures designed to manage system complexity, including the separation of system logic into clearly defined functions
Cryptography and Key Management	The safe use of cryptographic primitives and functions, along with the presence of robust mechanisms for key generation and distribution
Documentation	The presence of comprehensive and readable code-based documentation
Memory Safety and Error Handling	The presence of memory safety and robust error-handling mechanisms
Testing and Verification	The presence of robust testing procedures (e.g., unit tests, integration tests, and verification methods) and sufficient test coverage

Rating Criteria	Rating Description
Strong	No issues were found, and the system exceeds industry standards.
Satisfactory	Minor issues were found, but the system is compliant with best practices.
Moderate	Some issues that may affect systems safety were found.
Weak	Many issues that affect systems safety were found.
Missing	Are required component is missing, significantly affecting systems safety.
Not Applicable	The category is not applicable to this review.
Not Considered	The category was not considered in this review.
Further Investigation Required	Further investigation is required to reach a meaningful conclusion.

hashey
79ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

C. Code Quality Findings

We identified the following code quality issues through manual and automatic code review.

- Use constants instead of hard-coded values. Instead of 32, use the `N_BYTES_WORD` constant. `for idx in 0..32 { cb.add_lookup("Bitwise lookup", Lookup::Fixed{ tag: tag.clone(), values: [a.cells[idx].expr(), b.cells[idx].expr(), c.cells[idx].expr(),], },); }` **Figure C.1: zkevm-circuits/src/evm_circuit/execution/bitwise.rs#L47-L59**
- Use `Transition::Same` instead of `Delta(0.expr())`. The `require_step_state_transition` function will perform an extra addition when handling `Delta(0.expr())`. Also, the line referencing that field can be removed, since the type's default is `Same`. This issue is present in several files: `zkevm-circuits/src/evm_circuit/execution/{balance.rs, call_data_load.rs, extcodehash.rs, extcodesize.rs, is_zero.rs, not.rs}`. `// Stat transition let step_state_transition = StepStateTransition{ rw_counter: Delta(2.expr()),`

program_counter:Delta(1.expr()), stack_pointer:Delta(0.expr()), FigureC.2: zkevm-circuits/src/evm_circuit/execution/not.rs#L49-L53 • Removetheunusedfunction generate_lagrange_base_polynomial .The functionispresentat zkevm-circuits/src/evm_circuit/util/math_gadget.rs butitisnotusedin thecodebase. • Addextracheckswhendoingarithmeticontheopcode. Severalgadgetsdo arithmeticontheopcodetoextractarelevantvaluewhenmultiplereLATEDopcodes areadjacentinvalue. Thispatterniserror-prone,andthegadgetsdonotaddchecks hashey 80ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

toensuretheopcodeisinthecorrectrange. Extrachecksshouldbeincludedto preventmisuse,andtheseoperationsshouldbefactoredoutintoacommon moduletoaidreadability. letblockctx_tag=BlockContextFieldTag::Coinbase.expr() +(opcode.expr()- OpcodeId::COINBASE.as_u64().expr()); FigureC.3: evm_circuit/execution/block_ctx.rs#35-36 letswap_offset=opcode.expr()-(OpcodeId::SWAP1.as_u64()-1).expr(); FigureC.4: zkevm-circuits/src/evm_circuit/execution/swap.rs#35 letnum_additional_pushed=opcode.expr()- OpcodeId::PUSH1.as_u64().expr(); FigureC.5: zkevm-circuits/src/evm_circuit/execution/push.rs#85 lettopic_count=opcode.expr()-OpcodeId::LOG0.as_u8().expr(); FigureC.6: zkevm-circuits/src/evm_circuit/execution/logs.rs#102 letdup_offset=opcode.expr()-OpcodeId::DUP1.expr(); FigureC.7: zkevm-circuits/src/evm_circuit/execution/dup.rs#35 lettag= FixedTableTag::BitwiseAnd.expr()+(opcode.expr()- OpcodeId::AND.as_u64().expr()); FigureC.8: zkevm-circuits/src/evm_circuit/execution/bitwise.rs#45-46 • Thefollowingcommentisincorrect. The SSTORE gasrefundconstraintshave codecommentsdescribingeachconstraint. However,thecommentforthe delete_slot caseiswrong: //(value_prev≠value)&&(original_value≠value)&&(value== //Word::from(0)) letdelete_slot= not::expr(prev_eq_value.clone())*not::expr(original_is_zero.clone())* value_is_zero; FigureC.9: evm_circuit/execution/sstore.rs#L285-L288 • Theblanketmatchcasein require_step_state_transition couldleadto under-constrainedtransitions. The Transition::Any caseof require_step_state_transition ishandledimplicitlybyablanketmatch. Ifany newcaseisaddedtothe Transitionenum ,thedefaultbehaviorwillbetoleave thatfieldunconstrained. If,instead, Transition::Any isexplicitlyhandled, the Rustcompilerwillgenerateanincompletematcherror. hashey 81ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

```
matchstep_state_transition.$name{ Transition::Same⇒self.require_equal( concat(
("Statetransition(same)constraintof",stringify!($name)), self.next.state.$name.expr(),
self.curr.state.$name.expr(), ), Transition::Delta(delta)⇒self.require_equal( concat(
("Statetransition(delta)constraintof",stringify!($name)), self.next.state.$name.expr(),
self.curr.state.$name.expr()+delta, ), Transition::To(to)⇒self.require_equal( concat(
("Statetransition(to)constraintof",stringify!($name)), self.next.state.$name.expr(), to, ), _⇒{} }
FigureC.10: evm_circuit/util/constraint_builder.rs#538-555 • Somecommentsappeartobecopy-
pastedandrefertoothermodules. Commentsfor Error00GAccountAccessGadget and
ErrorInvalidCreationCodeGadget referto Error00GExpGadget and ErrorCodeStoreGadget ,respectively:
///Gadgettoimplementthecorrespondingoutofgaserrorsfor //[ `OpcodeId::EXP` ]. #[derive(Clone,Debug)]
pub(crate)structError00GAccountAccessGadget<F>{ FigureC.11:
```

```
evm_circuit/execution/error_oog_account_access.rs#21-24
///Gadgetforcodestoreoogandmaxcodesizeexceed #[derive(Clone,Debug)]
pub(crate)structErrorInvalidCreationCodeGadget<F>{ FigureC.12:
evm_circuit/execution/error_invalid_creation_code.rs#20-22 •
```

Thereisaredundantexpressionidentifiercomputationin store_expression . The store_expression functioncomputestheexpressionidentifiertwiceincase theexpressionisnotalreadystored:onceinthe find_stored_expression and againintheconstructionofthe StoredExpression . pub(crate)fnstore_expression(&mutself, name:&str, expr:Expression<F>, cell_type:CellType,)->Expression<F>{ hashey 82ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

```
//Checkifwealreadystoredtheexpressionsomewhere
letstored_expression=self.find_stored_expression(&expr,cell_type); matchstored_expression{
Some(stored_expression)⇒{ debug_assert!( !matches!(cell_type,CellType::Lookup(_)),
"Thesamelookupisdonemultipletimes", ); stored_expression.cell.expr() } None⇒{
//Evenifwe'rebuildingexpressionsforthenextstep, //theseintermediatevaluesneedtobestoredinthecurrent
step. letin_next_step=self.in_next_step; self.in_next_step=false;
letcell=self.query_cell_with_type(cell_type); self.in_next_step=in_next_step;
//Requirethestoredvaluetoequalthevalueoftheexpression letname=format!("{}",storedexpression),name);
self.push_constraint( Box::leak(name.clone().into_boxed_str()), cell.expr()-expr.clone(), );
self.stored_expressions.push(StoredExpression{ name, cell:cell.clone(), cell_type,
expr_id:expr.identifier(), expr, }); cell.expr() } } pub(crate)fnfind_stored_expression( &self,
expr:&Expression<F>, cell_type:CellType, )⇒Option<&StoredExpression<F>>{
letexpr_id=expr.identifier(); self.stored_expressions .iter()
.find(|&e|e.cell_type==cell_type&&e.expr_id==expr_id) } FigureC.13:
```

• Use `query_cell_phase2()` instead of `query_cell_with_type(CellType::StoragePhase2)` .
`let phase2_callee_code_hash = cb.query_cell_with_type(CellType::StoragePhase2);` Figure C.14: `zkvm-circuits/src/evm_circuit/util/common_gadget.rs#L711` • Use constants instead of hard-coded values.
`Ok(BaseFieldEccChip::<C>::variable_base_msm::<C>(self, ctx, &points, &scalars, C::Scalar::NUM_BITS as usize, 4, // empirically lump factor of 4 seem to be best)) } fn fixed_base_msm(&mut self, ctx:&mut Self::Context, pairs:&[(impl Deref<Target=Self::AssignedScalar>, C)],)->Result<Self::AssignedEcPoint, Error>{ let (scalars, points):(Vec<_>, Vec<_>)=pairs.iter().filter_map(|(scalar, point)|{ if point.is_identity().into(){ None } else{ Some((vec![scalar.deref().clone()], *point)) } }) .unzip(); Ok(BaseFieldEccChip::<C>::fixed_base_msm::<C>(self, ctx, &points, &scalars, C::Scalar::NUM_BITS as usize, 0, 4,)) }` Figure C.15: `snark-verifier/src/loader/halo2/shim.rs#L371-L406` • There are unnecessary type hints in `origin.rs` and `gasprice.rs` . The files contain several unnecessary type hints, such as `<N_BYTES_ACCOUNT_ADDRESS>` , `2u64` , `1u64` , and `-1i32` .
hashey 84ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

`fn configure(cb:&mut ConstraintBuilder<F>)->Self{ let origin = cb.query_word_rlc::<N_BYTES_ACCOUNT_ADDRESS>(); // Lookup in call_ctx the TxId
let tx_id = cb.call_context(None, CallContextFieldTag::TxId); // Lookup rw_table->call_context with tx origin address
cb.tx_context_lookup(tx_id.expr(), TxContextFieldTag::CallerAddress, None, // None because unrelated to call data
from_bytes::expr(&origin.cells),); // Push the value to the stack
cb.stack_push(origin.expr()); // State transition
let op_code = cb.query_cell(); let step_state_transition = StepStateTransition{ rw_counter:Delta(2u64.expr()), program_counter:Delta(1u64.expr()), stack_pointer:Delta((-1i32).expr()), } Figure C.16: evm_circuit/execution/origin.rs#L32-L53 • Unify the constraint builder APIs. There are several repeated functions in the ConstraintBuilder and BaseConstraintBuilder APIs in util/constraint_builder.rs . • There are functionally identical functions. The get_num_rows_required_no_padding and get_min_num_rows_required functions compute the same number of rows in a slightly different way.
pub fn get_num_rows_required_no_padding(block:&Block<F>)->usize{ // Start at 1 so we can be sure there is an unused `next` row available
let mut num_rows = 1; for transaction in &block.txs{ for step in &transaction.steps{ num_rows += step.execution_state.get_step_height(); } } num_rows += 1; // End Block num_rows } // ...
pub fn get_min_num_rows_required(block:&Block<F>)->usize{ let mut num_rows = 0; for transaction in &block.txs{ for step in &transaction.steps{ hashey 85ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC`

`num_rows += step.execution_state.get_step_height(); } } // It must have one row for EndBlock and at least one unused one
num_rows += 2 } Figure C.17: zkvm-circuits/src/evm_circuit.rs#L210-L242 • Consider renaming the offset_add function to set_offset . The offset_add function sets the offset as the argument instead of adding the argument to the offset , as the name suggests. /// Increment the step rw operation offset by `offset` .
pub (crate) fn offset_add(&mut self, offset:usize){ self.offset = offset } Figure C.18: zkvm-circuits/src/evm_circuit/util.rs#L659-L662 • There are incorrect comments in halo2-ecc . Several elliptic curve functions incorrectly say that they assume P.y is reduced, when they instead require Q.x to be reduced. These comments have been updated in version v0.3.0 of the upstream halo2-lib .
/// For optimization reasons, we assume that if you are using this with `is_strict=true`, then you have already called `chip.enforce_less_than_p` on both `P.x` and `P.y`
pub fn ec_add_unequal<F:PrimeField, FC:FieldChip<F>>(Figure C.19: halo2-lib/halo2-ecc/src/ecc/mod.rs#55-56 /// For optimization reasons, we assume that if you are using this with `is_strict=true`, then you have already called `chip.enforce_less_than_p` on both `P.x` and `P.y`
pub fn ec_sub_unequal<F:PrimeField, FC:FieldChip<F>>(Figure C.20: halo2-lib/halo2-ecc/src/ecc/mod.rs#97-98 • There are outdated documentation comments in halo2-ecc . The is_soft_zero and is_soft_nonzero methods of FieldChip have outdated comments that do not reflect the implementation. These comments have been updated in version v0.3.0 of the upstream halo2-lib . // Assume the witness for a is 0
// Constrain that the underlying big integer is 0 and <p. // For field extensions, check coordinate-wise.`

`fn is_soft_zero(&self, ctx:&mut Context<F>, a:&Self::FieldPoint)-> AssignedValue<F>; // Constrain that the underlying big integer is in [1, p-1].
// For field extensions, check coordinate-wise.
fn is_soft_nonzero(&self, ctx:&mut Context<F>, a:&Self::FieldPoint)-> AssignedValue<F>; Figure C.21: halo2-lib/halo2-ecc/src/fields/mod.rs#115-122 • Allow the dead_code lint and fix all issues. The dead_code lint is currently disabled; it should be enabled to allow developers to quickly detect unused`

functions and variables. • Reuse the CmpWordsGadget in the ComparatorGadget. The ComparatorGadget implementations should reuse the CmpWordsGadget instead of having the same constraints duplicated on both gadgets. • Simplify expression implementation. Add a comment explaining the deduction and simplify the expression. `let total_rws = not::expr(is_empty_block.expr()) * (cb.curr.state.rw_counter.clone().expr()-1.expr()+1.expr());` Figure C.22: `evm_circuit/execution/end_block.rs#L44-L45` • The logical and operator is used with a non-Boolean value. `cb.require_zero("value==0 when is_pad==1 for read", and::expr([meta.queryAdvice(is_pad, Rotation::cur()), meta.queryAdvice(value, Rotation::cur()),]),);` Figure C.23: `zkvm-circuits/src/copy_circuit.rs#L322-L328` hashey 87 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

D. Automated Analysis Tool Configuration

As part of this assessment, we used the tools described below to perform automated testing of the codebase.

D.1. Semgrep We used the static analyzer Semgrep to search for risky API patterns and weaknesses in the source code repository. For this purpose, we wrote rules specifically targeting the ConstraintBuilder APIs and the ExecutionGadget trait. `semgrep --metrics=off --sarif --config=custom_rule_path.yml` Figure D.1: The invocation command used to run Semgrep for each custom rule. `ImproperOpcodeEnforcementRule` The `ExecutionGadget::configure` implementations must check that the opcode being executed matches the execution state the machine is in. By using the `SameContextGadget`, the implementation implicitly enforces the correct opcode to execution state constraint. The following Semgrep rule finds `configure` functions that do not properly enforce opcode constraints by filtering the most common ways that this is validated. It results in 12 findings, some of which are the true positive issues reported in finding TOB-SCROLL-13, as well as some false positive results that can be dismissed. `rules: -id:improper-opcode-enforcement message:"configure function without proper opcode enforcement" languages:[rust] severity:ERROR patterns: -pattern:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... } -pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... let $V=SameContextGadget::construct(...); ... } -pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... let $V=BlockCtxGadget::construct(...); ... } hashey 88 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC`

`-pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... <... cb.require_equal(...,opcode.expr(),...)...>; ... } -pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... <... cb.require_in_set(...,opcode.expr(),...)...>; ... } -pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... <... CommonCallGadget::construct(...)...>; ... } -pattern-not:| fn configure(cb:&mut ConstraintBuilder<F>)->Self{ ... cb.add_lookup($LABEL,Lookup::Fixed{tag: FixedTableTag::ResponsibleOpcode.expr(),values:[..., opcode.expr(), ...],},...); }` Figure D.2: The `improper-opcode-enforcement.yml` rule. `OpcodeLookupWithinConditionRule` This rule aims to search for the `opcode_lookup` function called within a condition context, in search of instances of TOB-SCROLL-27. `rules: -id:opcode-lookup-in-condition message:"CB API calls do not take condition into account" languages:[rust] severity:ERROR patterns: -pattern-either: -pattern:$YY.opcode_lookup(...) -pattern-inside:| <... $XX.condition($COND,...)...>` Figure D.3: The `opcode-lookup-in-condition.yml` rule. `GateUsageOutsideofacreate_gateContextRule` This rule aims to search for uses of the `gate` issued described in TOB-SCROLL-27. hashey 89 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

`rules: -id:gate-usage message:".gate() cannot be used outside the create_gate() function" languages:[rust] severity:ERROR patterns: -pattern:$OBJ.gate(...) -pattern-not-inside:| $META.create_gate($LABEL,|$CB|{ ... });` Figure D.4: The `gate-usage.yml` rule. **D.2. Clippy** The Rust linter Clippy can be installed using `rustup` by running the command `rustup component add clippy`. Invoking `cargoclippy` in the root directory of the project runs the tool. Running Clippy with `cargoclippy --workspace --Wclippy::pedantic` will analyze the codebase with additional linters. hashey 90 ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

E. Fix Review Results When undertaking a fix review, Hashey reviews the fixes implemented for issues identified in the original report. This work involves a review of specific areas of the source code and system configuration, not comprehensive analysis of the system. From September 25 to October 2, 2023, Hashey reviewed the fixes and mitigations implemented by the Scroll team for the issues identified in this report. We reviewed each fix to determine its effectiveness in resolving the associated issue. Scroll provided PRs with fixes for all high-severity, medium-severity, and low-severity findings except for the low-severity finding TOB-SCROLL-21. Scroll also provided PRs with fixes for several of the informational-severity findings. In summary, of the 29 issues described in this report, Scroll has resolved 16 issues and has partially resolved four issues. Scroll has indicated that it does not intend to address two

issues, which have been labeled as resolved. No fix PRs were provided for the remaining six issues, so their fix statuses are undetermined. For additional information, please see the Detailed Fix Review Results below. ID Title Status

1 ModGadget is under constrained and allows incorrect MULMOD operation to be proven Resolved

2 The RlpU64Gadget is under constrained when is_lt_128 is false Resolved

3 The BLOCKHASH opcode is under constrained and allows the hash of any block to be computed Resolved 4 zkvm-circuits crate depends on an outdated version of halo2-ecc Partially Resolved

5 N_BYTES parameters are not checked to prevent overflow Partially Resolved

6 Differences in shared code between zkvm-circuits and halo2-lib Resolved hashey

91 ScrollzkEVM halo2Circuits Security Assessment PUBLIC

7 Under constrained warm status on CALL opcodes allows gas cost for gas cost forgery Resolved

8 RW table constants must match exactly when the verification key is created Undetermined

9 The CREATE and CREATE2 opcodes can be called within a static context Resolved

10 ResponsibleOp code table incorrectly handles CREATE and CREATE2 Resolved

11 Elliptic curve parameters omitted from Fiat-Shamir Unresolved

12 The gas cost for the CALL opcode is under constrained Resolved

13 Unconstrained opcodes allow nondeterministic execution Partially Resolved

14 Nondeterministic execution of ReturnDataCopyGadget and ErrorReturnDataOutOfBoundsGadget Resolved

15 Many RW counter updates are magic numbers Undetermined

16 Native PCS accumulation deciders accept an empty vector Resolved

17 The Error00GSloadStore and the Error00GLogGadget have redundant table lookups Undetermined

18 The State circuit does not enforce transaction receipt constraints Undetermined

20 The EXPop code has an unused witness Resolved

21 The bn_to_field functions silently truncate big integers Unresolved hashey

92 ScrollzkEVM halo2Circuits Security Assessment PUBLIC

22 The field_to_bn function depends on implementation-specific details of the underlying field Resolved

23 The values of the byte code table tag column are not constrained to be HEADER or BYTE Resolved

24 Unconstrained column on the byte code HEADER rows Undetermined

25 decompose_limb does not work as intended Resolved 26 Zero modulus will cause a panic Resolved

27 The ConstraintBuilder::condition API is dangerous Undetermined

28 The EXT CODECOPY opcode implementation does not work when the account address does not exist Resolved hashey

93 ScrollzkEVM halo2Circuits Security Assessment PUBLIC

Detailed Fix Review Results TOB-SCROLL-1: The ModGadget is under constrained and allows incorrect MULMOD operation to be proven Resolved in PR#512. The constraints for a_or_zero have been replaced with a select call that correctly forces a_or_zero to be 0 when n is 0. TOB-SCROLL-2: The RlpU64Gadget is under constrained when is_lt_128 is false Resolved in PR#615. The new constraints force is_lt_128 to be 1 in the case of a zero value. If is_lt_128 is 1, the circuit range-checks the original value. If is_lt_128 is 0, the circuit range-checks a value v, defined as follows: if byte[0] is the most significant byte, v equals byte[0]-128, and if not, v equals 0. If byte[0] is the most significant byte, this suffices to check that byte[0] is in the range [128, 256]. If byte[0] is not the most significant byte, other logic forces the result to be a non-zero limb after the first one, which means value = 256*x + y for some x > 0, y ≥ 0. Thus, is_lt_128 is 1 only if value is in [0, 128]. TOB-SCROLL-3: The BLOCKHASH opcode is under constrained and allows the hash of any block to be computed Resolved in PR#512. The commented-out lookup for current_block_number has been un-commented. The Scroll team should evaluate whether it is better to perform this lookup or to instead directly constrain current_block_number to equal cb.curr.state.block_number. TOB-SCROLL-4: zkvm-circuits crate depends on an outdated version of halo2-ecc Partially Resolved as of commit 7fe99fe4e3de14801f4d66f75bd35307de39b0a8 in zkvm-circuits and commit 70588177930400361c731659b15b2ab3f29f7784 in halo2-lib. The zkvm-circuits now crate depends on the v0.1.5 tag of the scroll-tech/halo2-lib repository, which includes a fix for the ECDSA implementation. However, we recommend at least updating to the upstream version 0.3.0, which includes many changes, including a different implementation of the scalar_multiply function used by the ECDSA implementation. The updated commit for halo2-lib contains all upstream changes we highlighted, but does not seem to be up to date with the upstream version of the library. Note that we did not perform a full security assessment of this commit, so we do not know which issues may still be present in it. TOB-SCROLL-5: N_BYTES parameters are not checked to prevent overflow Partially resolved in PR#512. assert!() calls have been added to constrain the N_BYTES parameter. The expression in constant_division.rs, shown in figure E.1, may overflow when compiled without overflow checks and may incorrectly allow extremely large values of N_BYTES. Note that the zkvm-circuits repository configures its release build to enable overflow checks. hashey

94 ScrollzkEVM halo2Circuits Security Assessment PUBLIC

```
assert!(N_BYTES*8+64-denominator.leading_zeros()asusize ≤ MAX_N_BYTES_INTEGER*8);
FigureE.1:Theexpressionthatmayoverflowwhencompiledwithoutoverflowchecks TOB-SCROLL-
6:Differencesinsharedcodebetweenzkevm-circuitsandhalo2-lib ResolvedinPR#709andPR#1001.Various
debug_assert!() callshavebeenreplaced with assert!() calls,andtheincorrect log2_ceil
functionhasbeenfixed. TOB-SCROLL-7:UnderconstrainedwarmstatusonCALLopcodesallows gascost forgery
ResolvedinPR#512,withsomeadditionalfixesaddedinPR#676.PR#512addsa constraintto CallOpGadget forcing
is_warm tobe true, buttheinitialvalueof is_warm is
notdirectlyconstrained.Theinitialvalueforallaccesslistreads isfalse,ashighlightedin
figureE.2,whiletheEthereumYellowPaperstates thatitshouldbetrueforprecompile
addresses,asshowninfigureE.3: fnbuild_tx_access_list_account_constraints(&mutself,q:&Queries<F>){
self.require_zero("field_tagis0forTxAccessListAccount",q.field_tag()); self.require_zero(
"storage_keyis0forTxAccessListAccount", q.rw_table.storage_key.clone(), );
self.require_boolean("TxAccessListAccountvalueisboolean",q.value()); self.require_zero(
"initialTxAccessListAccountvalueisfalse", q.initial_value(), ); self.require_equal(
"state_rootisunchangedforTxAccessListAccount", q.state_root(), q.state_root_prev(), );
self.condition(q.not_first_access.clone(),|cb|{ cb.require_equal(
"valuecolumnatRotation::prev()equalsvalue_prevatRotation::cur()", q.rw_table.value_prev.clone(),
q.value_prev_column(), ); }); } FigureE.2:Basicconstraintsfortheaccesslist hashey
95ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC
```

FigureE.3:AnexcerptfromtheEthereumYellowPaperspecifyingtheinitialaccesslist
TheinitialprecompileaccessvalueshavebeenfixedinPR#676byaddinganaccesslist
writeforeachprecompiletothe BeginTx state. TOB-SCROLL-
8:RWtableconstantsmustmatchexactlywhentheverificationkeyis created
Undetermined.No fixwas providedforthisissue,sowedonotknowwhetherthisissue hasbeenaddressed. TOB-
SCROLL-9:TheCREATEandCREATE2opcodescanbecalledwithinastatic context ResolvedinPR#512.The
CreateGadget::configure methodnowperformsa call_context lookup toretrieve the IsStatic
fieldandthenconstrainsthatresulttobe zero. TOB-SCROLL-
10:ResponsibleOpcodetableincorrectlyhandlesCREATEandCREATE2 ResolvedinPR#512.The responsible-
opcodetablehasbeenupdatedtomap ExecutionState::CREATE to OpcodeId::CREATE and
ExecutionState::CREATE2 to OpcodeId::CREATE2 . TOB-SCROLL-
11:EllipticurveparametersomittedfromFiat-Shamir
Unresolved.The Scrollteamhasindicatedthatitdoesnotintendtofixthisissue. TOB-SCROLL-
12:The gascostfortheCALLopcodeisunderconstrained
ResolvedinPR#774.Previously,whencallinganemptyaddressoraprecompile,the unconstrainedcell
step_gas_cost wasusedtodeterminethe gascostofthe CALL opcode.Now,inallcases,the gas_left fieldof
StepStateTransition valuesisderived fromthefullyconstrainedcells callee_gas_left , gas_cost ,and
call_gadget.has_value .Wedidnotevaluatewhethertheoverallgas calculationis
correctinthisfixreview, butitisconstrained. TOB-SCROLL-
13:Unconstrainedopcodesallow nondeterministic execution
PartiallyresolvedinPR#633andPR#736.InPR#633,allmentionedgadgetsnowconstrain
theopcode.Scrollshouldinspectandtest ErrorPrecompileFailed asdevelopment hashey
96ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

continuestoensurethatonlyprecompilecalls can trigger it andonlywhentheyfail.PR #736addsconstraintstothe
ReturnRevertGadget componenttoensurethat REVERT
opcodesareaccompaniedbyareversionintheRWtable.Wewerenotabletodetermine
fromthesediffswhetheritispossibletohaveareversionwhileexecutinga RETURN
opcode.Scrollshouldinspectandtestthiscomponenttoensurethatamaliciousprover
cannotmanipulatethebehaviorofthe RETURN opcode. TOB-SCROLL-
14:NondeterministicexecutionofReturnDataCopyGadgetand ErrorReturnDataOutOfBoundGadget
ResolvedinPR#661.Theoverflow-relatedvalidationlogicof ReturnDataCopyGadget and
ErrorReturnDataOutOfBoundGadget hasbeenfactoredoutintoacommoncomponent, CommonReturnDataCopyGadget
.Thiscommongadgetforcesthevalidationtosucceedor failbasedonthe is_overflow parameter.Notethat
is_overflow issetto 1.expr() in onecaseand false.expr()
intheother.Whiletheseexpressionsdobehavecorrectly,we
recommendthatthe Scrollteam makethesesymmetrical;thatis,either 1.expr() and 0.expr() ,or true.expr()
and false.expr() . TOB-SCROLL-15:ManyRWcounterupdatesaremagicnumbers
Undetermined.No fixwas providedforthisissue,sowedonotknowwhetherthisissue hasbeenaddressed. TOB-
SCROLL-16:NativePCSaccumulationdecidersacceptanemptyvector
ResolvedinPR#17.Thecodenowpanicswithanassertionfailureifpassedanempty vector. TOB-SCROLL-
17:TheError00GSloadSstoreandtheError00GLoggadgetshave redundanttablelookups
Undetermined.No fixwas providedforthisissue,sowedonotknowwhetherthisissue hasbeenaddressed. TOB-
SCROLL-18:TheStatecircuitdoesnotenforcetransactionreceiptconstraints
Undetermined.No fixwas providedforthisissue,sowedonotknowwhetherthisissue hasbeenaddressed. TOB-

SCROLL-20: The EXPopcode has an unused witness Resolved in PR#838. The zero_rlc field has been removed. TOB-SCROLL-21: The bn_to_field functions silently truncates big integers Unresolved. The Scroll team has indicated that it does not intend to fix this issue. TOB-SCROLL-22: The field_to_bn function depends on implementation-specific details of the underlying field hasheyeye 97ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

Resolved in PR#15. A test has been added to confirm that the data representation matches the expectations of the implementation. Scroll should ensure that this test runs before any deployment. TOB-SCROLL-23: The values of the bytecode data table tag column are not constrained to be HEADER or BYTE Resolved in PR#681. A Boolean constraint has been added to the tag column. TOB-SCROLL-

24: Unconstrained columns on the bytecode HEADER rows

Undetermined. No fix was provided for this issue, so we do not know whether this issue has been addressed. TOB-SCROLL-25: decompose_limb does not work as intended Resolved as of commit

483feb2e4554fcab58878d7c8e6a6f8be792e2f2 . The decompose_limb

method has been more completely implemented and appears to now

iterate correctly through the limbs. We did not fully review the correctness of this

implementation, and there do not appear to be any direct tests for this implementation.

Scroll should add tests to ensure its correct behavior. TOB-SCROLL-26: Zero modulus will cause a panic

Partially resolved in PR#4, then fully resolved in PR#12. The mod_mult method has been modified to use the

number_is_zero method to return 0 when the modulus parameter is 0 . However, the original fix to the

number_is_zero method, shown in figure E.4, introduces a new error by incorrectly checking number.limbs[0]

three times. //return 0 if not zero, 1 if zero for number pub fn number_is_zero(&self, region: &mut Region<F>,

range_check_chip: &mut RangeCheckChip<F>, offset: &mut usize, number: &Number<F>,)-

> Result<Limb<F>, Error> { let zero = F::zero(); let three = F::from(3u64); //limb0_zero is 0 if not zero, 1 if zero

let limb0_zero = self.config.eq_constant(region, range_check_chip, offset, &number.limbs[0], &zero)?;

let limb1_zero = self.config.eq_constant(region, range_check_chip, offset, &number.limbs[0], &zero)?;

let limb2_zero = hasheyeye 98ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

self.config.eq_constant(region, range_check_chip, offset, &number.limbs[0], &zero)?;

Figure E.4: The incorrect indexing into number.limbs (misc-precompiled-

circuit/src/circuits/modexp.rs#525-545)

The Scroll team has fixed this newly introduced issue in PR#12, and this prevents a zero

modulus causing a panic by replacing it with the value 1 when dividing.

Although this particular issue has been resolved, the Scroll team should inspect and test

this component carefully, as it has had several correctness problems.

In particular, we recommend that Scroll investigate whether these methods behave

correctly in the presence of malformed values of the Number type. The Number type, shown

in figure E.5, appears to contain a "CRT-style" representation of a large number, similar to the CRTInteger

type in halo2-ecc . The first three entries of its limbs array contain the

"truncation" part, and the fourth entry contains the "native" part. This representation is not

documented in the code, and we have not determined whether it is possible to create a malformed Number

value within the modexp circuit. #[derive(Clone, Debug)] pub struct Number<F: FieldExt> { limbs:

[Limb<F>; 4], } Figure E.5: The Number type (misc-precompiled-circuit/src/circuits/modexp.rs#23-26)

TOB-SCROLL-27: The ConstraintBuilder::condition API is dangerous

Undetermined. No fix was provided for this issue, so we do not know whether this issue has been addressed. TOB-

SCROLL-28: The EXT_CODE_COPY opcode implementation does not work when the account address does not exist

Resolved in PR#846. Conditional constraints have been added to explicitly handle accounts

with no code, as indicated by a zero value in their CodeHash field. When an account has no

code, the code size is forced to be zero, and when an account has code, the code size is determined by a

bytecode_lookup . We did not find any tests for this behavior. We

recommend that the Scroll team add tests to ensure that this behavior is correct and

remains correct during ongoing development. hasheyeye 99ScrollzkEVMhalo2CircuitsSecurityAssessment PUBLIC

F. Fix Review Status Categories

The following table describes the statuses used to indicate whether an issue has been sufficiently addressed.

FixStatus StatusDescription Undetermined The status of the issue was not determined during this engagement.

Unresolved The issue persists and has not been resolved.

Partially Resolved The issue persists but has been partially resolved.

Resolved The issue has been sufficiently resolved. hasheyeye 100ScrollzkEVMhalo2CircuitsSecurityAssessment

PUBLIC