

# Branch Flipper: Unlocking Fuzz Blockers with Coverage-Grounded LLMs

Youngjoong Kim  
*Theori Inc.*

Team U.S.  
*Theori Inc.*

Team AI for Offensive Security  
*Theori Inc.*

## Abstract

Fuzzing is a software testing method that assesses the safety and functionality of the software by feeding it with randomly generated inputs. Among various fuzzing techniques, coverage-guided fuzzing and greybox fuzzing aim to expand the portion of tested code, referred to as coverage, by using it as a signal to select seed inputs for effective mutation and exploration. However, despite this feedback mechanism, the number of explorable code paths diminishes exponentially as the branching depth increases, making it challenging to cover deeper execution paths solely through random mutations. In such cases, a branch blocker refers to a specific branch that is not taken before, thereby blocking the execution of subsequent code paths. Besides, with the significant advancements in large language models (LLMs), fuzzing strategies that leverage LLMs have been explored recently, particularly for seed generation and fuzz harness generation. In this paper, we propose a system *Branch Flipper* that localizes branch blockers and generates inputs to unlock them using LLMs. To mitigate LLM hallucinations and improve reliability, the system incorporates a grounding process using oracles and coverage-based feedback. Rather than solely relying on randomness, the system addresses branch blockers semantically, utilizing the reasoning capabilities of LLMs. We conduct experiments to test the validity of the system, and they demonstrate that the branch flipper effectively addresses fuzz blockers. Moreover, its capability of generating inputs that reach suspicious or potentially vulnerable functions also indicates its practical applicability to penetration testing contexts.

## 1 Introduction

Fuzzing is a software testing methodology that assesses software robustness by generating random program inputs to identify potential ill-behaviors. Among various approaches, greybox fuzzing utilizes code coverage to expand the testing scope by applying random mutations to seeds that previously increased coverage. However, random mutation alone

poses limitations even under the coverage feedback. Particularly, randomly mutated inputs struggle to take the unhit branches of nested or complex predicates. In practice, grey-box fuzzers used in projects like OSS-Fuzz [23] often exhibit coverage saturation, where further code paths are no longer being discovered despite continued testing. To address this, grammar-based fuzzing and similar techniques have been studied, but these typically require hand-crafted features and domain knowledge, making it difficult to scale fuzzing to a wide range of software targets.

Large Language Models, also referred to as LLMs [2], have shown a dramatic improvement in their natural language generation capabilities. With an instruction tuning [20], it has become possible to produce user-intended outputs even with zero or few samples [28]. By leveraging these capabilities, there has been a growing interest in applying LLMs to fuzzing. Recent studies have explored using LLMs to generate fuzz harnesses(drivers) [16, 23, 29], and to generate input seeds [7, 12, 18, 24, 25]. In the case of harness generation, however, evaluating whether the generated harnesses truly reflect the actual use cases remains a significant challenge. Even if a crash is triggered by the generated harness, it is often difficult to convincingly argue that the resulting crash is meaningful or to maintain consensus on its validity.

On the other hand, leveraging LLMs for seed generation also suffers from challenges. While random mutation can be performed and evaluated at a very high frequency, LLMs require high-performance hardware and thus suffer from the inevitable latency, making it impractical to completely replace random mutation with them. Even when using LLM APIs, the process can be costly in terms of both time and resources. As a result, there is a growing need for strategies that selectively incorporate LLMs.

Recently, there have been efforts to fully automate penetration testing using LLM-based agentic systems, such as ReAct [32]. These systems are capable of planning actions, sending queries, interpreting results, and managing history throughout the process [8, 13]. Furthermore, U.S. DARPA, Defense Advanced Research Projects Agency, launched the

AI Cyber Challenge (AIXCC), a competition aimed at developing AI agents that automatically discover vulnerabilities in open-source software projects, which are widely and actively used, but their security and reliability are often treated as secondary concerns. Given the diverse scope of open-source projects and the broad range of potential vulnerabilities, the competition requires a general-purpose testing system that can effectively handle such variability.

To support greybox fuzzing in various situations and complement its random mutation strategy, we propose a system, *Branch Flipper*, which aims to generate *blocker-unlocking* inputs when fuzz blockers are encountered during fuzz testing. Our system consists of a blocker localizer, a LLM-based agentic framework capable of generating seeds with feedback-driven refinement. Our system extends beyond traditional LLM-based input generation approaches, which narrow the scope to focus on LLM-generatable formats, such as JSON, XML, and source code. To enable the generation of various types of inputs, including binary files, we incorporate a Python-based and Kaitai Struct [21]-based seed generator. Additionally, we provide the LLM with reference seeds to facilitate the generation of inputs. To mitigate LLM hallucination, we employ coverage-based feedback, along with various auxiliary tools, enabling the system to reliably flip specific branches in the target function. These findings demonstrate the potential of our approach to enhance grey-box fuzzing by supporting the execution of previously unreachable paths in a semantic manner.

This study is part of the AIXCC competition; this capability is not only valuable for resolving fuzz blockers but also contributes to the broader AIXCC objective by enabling the generation of seed inputs that can reach suspicious code regions. If the existing seed pool is unable to reach a target function, our system can analyze the call chain, identify the blocking condition, and generate inputs to unlock the blocked path.

This report presents the development and technical validation of a system designed to identify and overcome fuzz blockers, facilitating deeper and more targeted software testing in large-scale, diverse open-source codebases.

## 2 Related Works

Fuzzing open-source software, specifically, OSS-Fuzz [23] has conducted cluster-level fuzzing on open-source projects and reported discovered crashes. The range of testing includes not only command-line programs, but also more complex systems such as HTTP servers [22], GUI software, and even device drivers. To fuzz such software using existing fuzzers [3, 15, 33], a corresponding fuzz driver or fuzz harness is required to forward random byte strings to the proper APIs. However, despite the novelty of greybox fuzzing, its heavy reliance on random mutations poses limitations, particularly in exploring code paths beyond a certain depth. As a result, deeply nested

or complex execution paths often remain untested.

Recently, large language models (LLMs) have been developed and utilized in various areas. Large Language Models (LLMs) are a class of deep learning models constructed with carefully designed neural network architectures [26] and a vast number of parameters, enabling them to store and represent a large amount of information. These models are capable of identifying patterns from user-provided input prompts and generating appropriate responses [2]. Moreover, through Reinforcement Learning from Human Feedback (RLHF) [20], they are further trained to produce helpful and user-aligned responses. LLMs are typically trained to predict the next token given a sequence of input tokens. During generation, they follow an autoregressive decoding scheme: the model predicts the next token, appends it to the end of the sequence, and then uses it as input for the next token prediction. This process continues iteratively, producing coherent and context-aware outputs. Depending on the provided prompt, LLMs can interact with users in a conversational manner. Recently, advanced prompting techniques beyond simple dialogue have also been explored [1, 4, 27, 28, 31].

Moreover, LLMs have evolved in a "thinking mode", where they internally generate intermediate thoughts before producing a final answer [5, 6, 9, 10, 19]. These models often begin by generating internal representations that they consider useful for solving the given task, and use them as a precondition to produce more accurate and reliable responses. This approach typically utilizes Monte Carlo Tree Search, where the model is trained to select token sequences with higher expected accuracy. By simulating and evaluating multiple reasoning paths before answering, these models demonstrate improved performance on complex tasks requiring multistep reasoning.

From this background, there has been a growing interest in applying LLMs to fuzzing. These approaches may involve generating fuzz harnesses(drivers) or producing a seed corpus. The generation of fuzz drivers typically requires substantial domain-specific knowledge about the target library or API, making the process human-intensive. To alleviate this, several studies have proposed automating harness generation using LLMs [16, 23, 29]. OSS-Fuzz-Gen [23] attempts to generate a harness for each individual function, while PromptFuzz [16] explores interactions across multiple APIs. However, the harness generation presents a notable challenge: evaluating the correctness of the generated harness. When crashes are triggered, it is often difficult to reach a consensus on whether the crash is meaningful or realistically exploitable in real-world scenarios.

There have been numerous attempts to apply LLMs to fuzz corpus generation [7, 12, 17, 18, 25]. Most of these efforts have focused on LLM-generatable inputs, such as XML, JSON, or source code, often targeting specific domains. However, applying these methods to binary-level software remains challenging. To address these limitations, some approaches generate Python scripts instead, which are then executed to

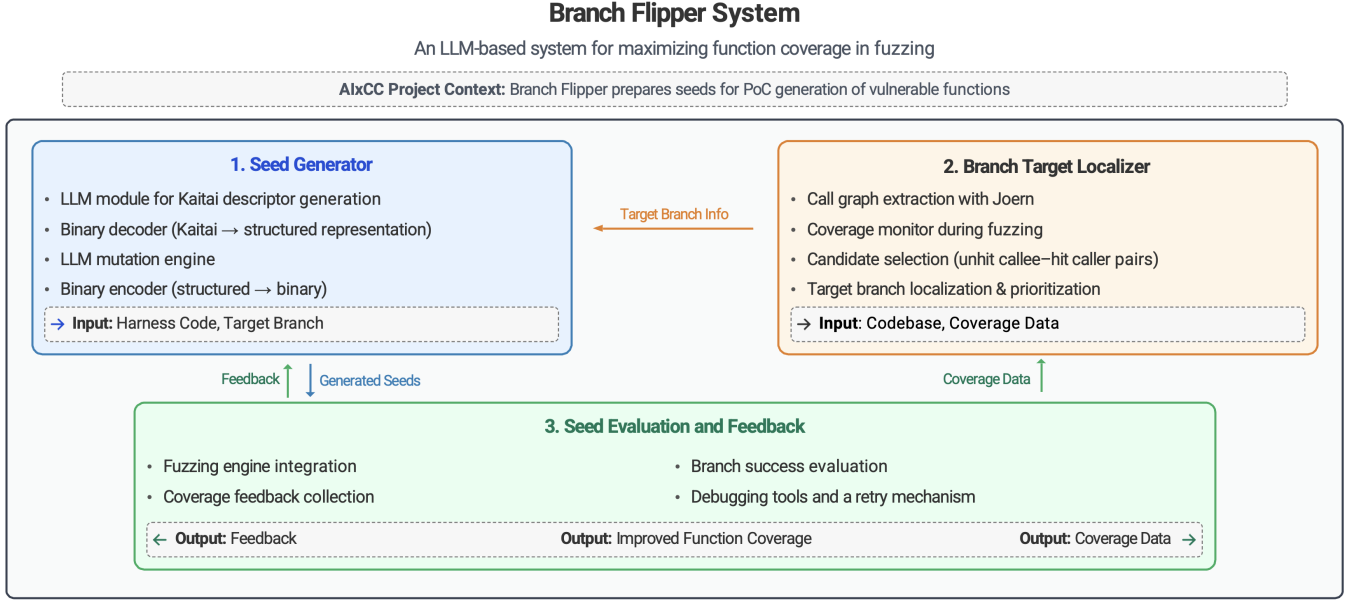


Figure 1: Overall architecture of the Branch Flipper. The system consists of three components: (1) a seed generator, (2) a branch target localizer, and (3) a seed evaluation and feedback module. Each component runs in a peer-to-peer manner, where each agent is self-directed, performing local tasks and initiating tool calls or requesting other agents via a message queue.

produce binary inputs [24]. Nonetheless, due to the latency of LLM, fully replacing traditional mutation engines is infeasible in terms of throughput. As a result, a common strategy has augmented existing fuzzing workflows by generating a limited number of high-quality seed corpora, which can then be used to enhance or bootstrap the mutation-based fuzzing process [17].

Some studies have tried to fully automate penetration testing using LLM agents [32], which can plan, execute, and self-correct from feedback in a fully end-to-end manner [8, 13]. The AI Cyber Challenge (AIxCC), organized by the U.S. Defense Advanced Research Projects Agency (DARPA), is a gamified autonomous hacking competition. In the competition, teams are provided with a selection of open-source projects of various programming languages and domains in the OSS-Fuzz format. Teams then deploy AI agents that identify vulnerabilities and generate proof-of-concept (PoC) seed inputs without any human intervention, and submit patches for discovered bugs. Points are awarded when a submitted seed triggers a sanitizer when run through the predefined harness, indicating a successful vulnerability discovery. However, duplicate submissions that stem from the same root cause are penalized. Projects are presented in two modes: diff mode, which focuses on identifying vulnerabilities introduced in recent Git commits, or full mode, which analyzes the entire codebase without considering the editing history. After discovering vulnerabilities, teams can also submit patches, and points are granted if the patch prevents the vulnerability while maintaining its functionality.

### 3 Method

In this section, we introduce the BranchFlipper, an LLM-based seed generation for unlocking fuzz blockers. The BranchFlipper comprises several components: LLM-based seed generation using Kaitai Struct [21], a Joern [14]-based fuzz blocker localizer, LLM-based seed generation, and a coverage-based feedback loop.

#### 3.1 Overview

This work, *Branch Flipper*, is part of the AIxCC project. Our AIxCC system runs fuzzers in the background with the BranchFlipper to prepare a collection of seeds, which serve as a baseline input when the system attempts to generate a PoC of corresponding vulnerable functions. Thus, to acquire at least one seed for each function, this work aims to maximize function coverage by generating seeds that unlock fuzz blockers and support the discovery of new functions.

Our proposed Branch Flipper system consists of three major components: (1) a seed generator, (2) a branch target localizer, and (3) a seed evaluation and feedback.

LLM-based seed generation has been an active area of research [7, 12, 17, 18, 24, 25], with most prior works focusing on natural language-compatible formats, and directly generating binary inputs remains challenging. Some approaches [24] guide LLMs to generate Python scripts, which are then executed to produce binary files.

To support more flexible and LLM-friendly binary han-

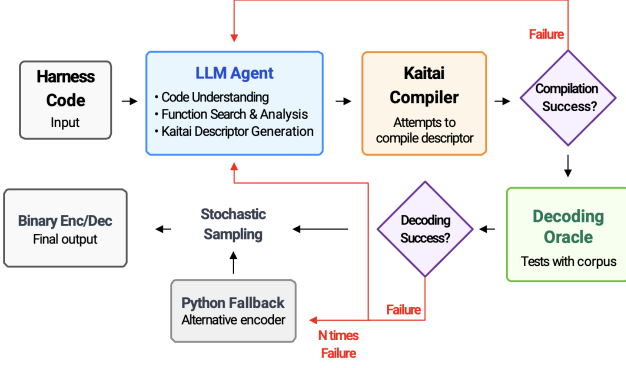


Figure 2: Workflow of a Kaitai descriptor generation.

dling, we utilize Kaitai Struct [21], a structured YAML-like language for describing binary formats. Modern LLMs are typically familiar with this format, making it a practical choice for bridging the gap between natural language capability and binary manipulation. In our system, given a harness code, the LLM attempts to generate a corresponding Kaitai Struct descriptor. From a pool of Kaitai descriptors, the system selects one to decode the binary, yielding a structured representation, and then feeds it to an LLM. The LLM then mutates this structured representation based on the analysis of the targeted branch blocker, re-encodes it into a new binary input. If the targeted branch is successfully taken by the new input, the system proceeds to the next candidate; otherwise, it requests a retry with the aid of debugging tools.

To localize the fuzz blocker, we initially extract the function call graph using Joern [14] from the whole codebase. During fuzzing, we monitor function coverage, and if a function is covered but its callee remains unhit, these unhit callees and caller pairs are collected as candidates. The most impactful candidate then becomes the next target. Concurrently, we collect the seeds that reach each function and use them as a baseline for seed generation.

Through this iterative loop, we aim to generate blocker-unlocker with enhanced binary generation capabilities and grounding LLM hallucinations. Ultimately, this enables the LLM to contribute meaningfully to branch blockers within the greybox fuzzing pipeline.

### 3.2 Kaitai Struct Generation

Previous studies have primarily focused on natural language-compatible seed formats, those used in JSON/XML parsers or compilers, allowing them to directly generate fuzzing seeds and use them as-is during fuzzing. However, due to the inherent limitations of LLMs in directly constructing low-level byte streams, [24] has attempted to generate Python code that creates a binary file. This approach is particularly useful when fuzzing harnesses rely on FuzzedDataProvider-style APIs, which require a precise understanding and crafting of

byte-level inputs.

Kaitai Struct [21] offers a YAML-based structured language for describing binary formats. Given a description, the Kaitai compiler generates Python modules that provide both parsers and generators for the corresponding byte stream. In our study, we support two methods for generating byte stream inputs:

1. The traditional approach of using LLMs to generate Python-based parsers and generators (decoders and encoders).
2. A new approach that leverages LLM-generated Kaitai Struct descriptors, which are then compiled and used to parse and generate binary inputs.

The system provides a fuzzing harness code to the LLM and instructs it to first understand the code. The LLM is capable of searching function definitions, references, and reading the code to understand the context. It then attempts to generate a Kaitai struct descriptor based on the understanding. If the Kaitai compiler fails to compile the descriptor, the system returns the compiler error to the LLM, prompting a revision until a valid and compilable Kaitai description is produced.

Modern LLMs such as GPT-4o, o4-mini, o3, and Claude 3.5 Sonnet already know and understand Kaitai Struct and often demonstrate notable generation capabilities for near-correct descriptors. However, syntactic and semantic errors also remain common. Hence, we construct an iterative feedback loop involving compilation success/failure signals to guide the LLM toward producing compilable Kaitai descriptors.

```

meta:
  id: png
  title: PNG (Portable Network Graphics) file
  endian: be
  xref:
    forensicswiki: portable_network_graphics_(png)
    iso: 15948:2004
    loc: fdd000153
  # Structure definition sequence
  seq:
    # PNG file signature
    - id: magic
      contents: [137, 80, 78, 71, 13, 10, 26, 10]
    # IHDR chunk
    - id: ihdr_len
      type: u4
      valid: 13
    - id: ihdr_type
      contents: "IHDR"
    - id: ihdr
      type: ihdr_chunk
    - id: ihdr_crc
      size: 4
    # Remaining chunks
    - id: chunks
      type: chunk
      repeat: until
      repeat-until: _.type == "IEND" or _.io.eof

chunk:
  ... SHORTENED
  ihdr_chunk:
    doc-ref: https://www.w3.org/TR/png/#11IHDR
    seq:
      - id: width
        type: u4
      - id: height
        type: u4
      - id: bit_depth
        type: u1
      - id: color_type
        type: u1
      enum: color_type
      - id: compression_method
        type: u1
      - id: filter_method
        type: u1
      - id: interlace_method
        type: u1
  plte_chunk:
    doc-ref: https://www.w3.org/TR/png/#11PLTE
    ... SHORTENED

```

Figure 3: A sample descriptor generated by the Claude 3.5 Sonnet when the harness code of the project *libpng* is given.

However, the performance of LLMs tends to degrade with longer context lengths [11]. Empirically, we observed that



when the provided source code or conversation becomes lengthy, the compile success rate decreases. We also noticed that LLMs sometimes oversimplify Kaitai descriptions, omitting crucial structural details to minimize compile errors. Despite these issues, Kaitai-based representations can sometimes provide a more precise description of binary inputs than their Python-only counterparts.

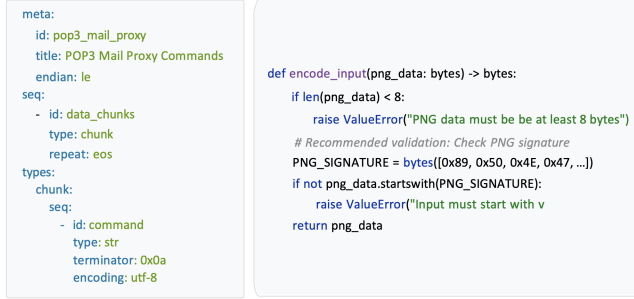


Figure 4: Some cases of oversimplified generated encoders of Kaitai and Python settings. Left one is generated Kaitai descriptor when the Nginx harness is given, and the right one is generated Python-based encoder when the libpng harness is given (the same harness as Figure 3.) Both are oversimplified and simply forward inputs to outputs as is.

In practice, to maximize the system’s effectiveness, we adopt a stochastic strategy: for each seed generation task, the system stochastically chooses between the Python-based encoder and the Kaitai struct encoder. And also in cases where Kaitai compilation fails, the Python-based encoder serves as a fallback.

To be effective, the Kaitai descriptor must not only be compilable but also correctly decode actual seeds. Since our generated decoders operate on the corpus generated by the fuzzer, they must be robust enough to handle a wide range of inputs. Therefore, we incorporate a decoding test (or oracle) after compilation: given a fuzzing corpus sample, the compiled decoder tries to decode it, and decoding failures are reported back to the LLM. This additional feedback helps guide the LLM to produce more resilient and generalizable Kaitai structures.

### 3.3 LLM-based Seed Generation

Once a suitable encoder has been constructed (either Python-based or Kaitai-based), the LLM is now able to generate concrete binary seeds.

Our goal is to analyze a specific branch predicate and generate a suitable seed. However, reaching this particular predicate often requires passing through multiple prior control flows, and each imposes its own constraints on the inputs. As the depth of nested predicates increases, attempting to satisfy the target predicate from scratch (i.e., without initial guidance or baseline) while preserving all prior becomes exponentially

more complex, especially given the context-length limitations of LLMs.

To address this, we utilize the collected seeds from the fuzzer run. We assume that we already possess a seed that has successfully reached the function containing the target predicate. Using this as a starting point allows us to focus only on solving the target predicate, which has already satisfied the earlier branches, and we can effectively reduce the problem complexity.

In our system, given a seed that reaches the function containing the target branch predicate, we perform the following sequence:

1. The LLM is provided with the line number of the target branch, and the corresponding seed previously reached the function.
2. The LLM analyzes the branch predicate at the target location and uses the given seed as a starting point.
3. The LLM tries to generate a new seed that seems to take the desired branch.

Since raw binary seeds are uninformative to LLMs, we decode the seed into a structured JSON format using the previously generated decoders. This JSON representation allows the LLM to understand and reason about the input fields in a semantically meaningful way.

Based on this plan, the LLM generates a Python script that produces a binary seed aiming to satisfy the branch condition. This script is executed within a sandboxed environment to produce the final seed file. The resulting binary input is then executed within the fuzzing harness to verify whether the target branch was taken.

## 3.4 Joern-based Fuzz Blocker Localizer

Our second challenge was how to localize fuzz blockers and feed them into the branch flipper. The current requirement is that unlocking a predicate within a given function must allow the discovery of previously unexplored functions. We addressed this in two stages: the first defines the frontier, and the second localizes the targeting branch within that frontier.

### 3.4.1 Call Graph and the Frontiers

The first stage involves defining the *frontier*. A frontier refers to a function candidate that has been hit at least once during fuzzing and has at least one callee function that has never been invoked. In other words, if a certain predicate in this function is triggered, it could lead to the exploration of a new function that had not been visited before.

There are various ways to build a call graph. We employ Joern [14], a general-purpose static analyzer known for its speed and stability. Joern supports the extraction of a Code Property

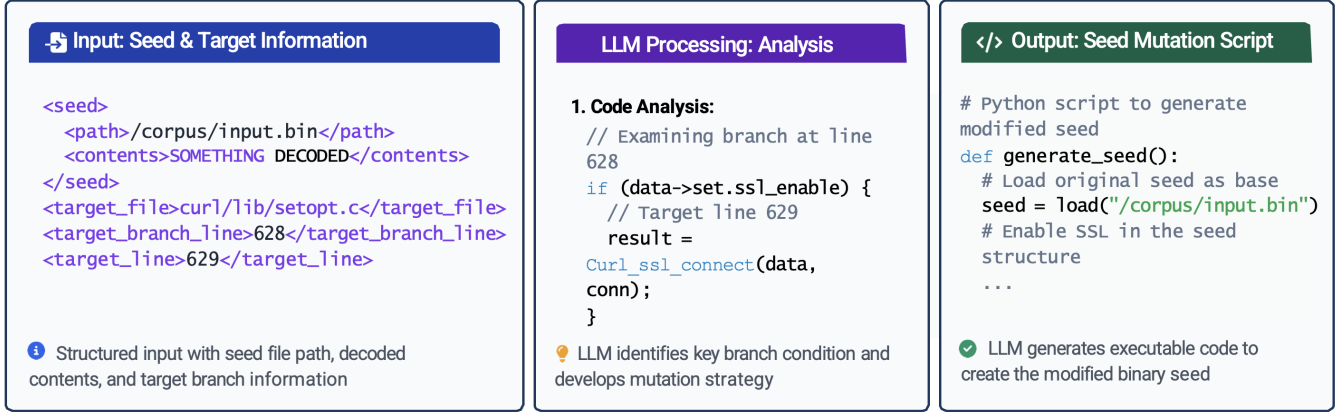


Figure 5: Workflow of the blocker-unlocker generation. The LLM is provided with structured inputs, including a decoded baseline seed, target file, and branch. Then, the LLM starts to analyze the predicate using certain tools, finding function definitions and references, reading them, and attempting to generate a new seed.

Graph (CPG), including call relationships. In C/C++ open-source projects, macro variables and compiler preprocessor often conditionally enable or disable certain code paths. As a result, call relationships visible in source code may not exist in the final compiled binary, creating false positives. Joern mitigates this by explicitly handling macros and automatically including relevant header files, yielding more reliable results than tools like GNU Global [30].

In parallel to the LLM Agent, we continuously run fuzzers in the background. Whenever the LLM generates a new seed, it is added to the fuzzer’s seed corpus. This setup allows the system to continually integrate new seeds and expand function coverage through basic mutation. When a new seed is generated, a background process monitors this event, checks coverage, and updates the metadata from the call graph accordingly.

Each node in the call graph stores metadata, including its shallowest-known caller, the call depth, and a visit flag. When a new seed expands function coverage, we traverse the call graph, updating the caller metadata and propagating to its callee nodes. If a function node has at least one unvisited callee, it is labeled as a *frontier* and queued for branch flipping.

However, the number of such frontiers can be massive, often exceeding 10,000 per hour during early fuzzing. Invoking the LLM Agent for each of these is impractical in terms of computational budgets. In practice, some of these frontiers are unlocked by random mutators before being conveyed to the branch flipper agent. Therefore, we implement a hard waiting time for each queued frontier. Only after this waiting period, we check whether the callee remains unvisited, and only then proceed to branch flipping.

### 3.4.2 Branch Localization

Once a frontier is determined, we must identify a branch predicate to target. Since our goal is to trigger an unvisited function, we target branches that may call such functions.

Joern supports queries related to control structures. Given a call expression to an unvisited callee within a hit function, we can use Joern’s CPG to locate the branch predicate enclosing the call. The first such branch is guaranteed to be un-taken. However, in nested control structures, this may not be the closest former branch predicate that has been evaluated, and identifying the correct one may require line-level coverage analysis. This can be costly due to complications like comments and parsing behavior. Given AIXCC’s multi-language scope, we conduct a simplified strategy: we target only the immediate branch predicate preceding the call expression. Despite this simplification, we still observe a significant number of viable cases, and we believe flipping such branches remains highly meaningful. We left a more sophisticated localization of the earliest evaluated branch predicate as future work.

Even with precise localization, some branches may still be unreachable due to environmental constraints. For example, when a harness hardcodes certain variables, such as enabling logging, callee functions cannot be invoked by such variables.

To address this, we add a single-stage query to the LLM, issuing a one-sentence prompt to assess whether the branch is likely to be hit. We only attempt flipping for those deemed feasible. Although this may yield false negatives, this filtering is acceptable since the number of candidate frontiers still exceeds our processing budget. Additionally, we prioritize candidate frontiers based on the estimated number of reachable functions or lines of code that can be explored upon unlocking the blocker, thereby guiding the flipper toward high-impact branches, which we previously referred to as the *impactful frontier*.

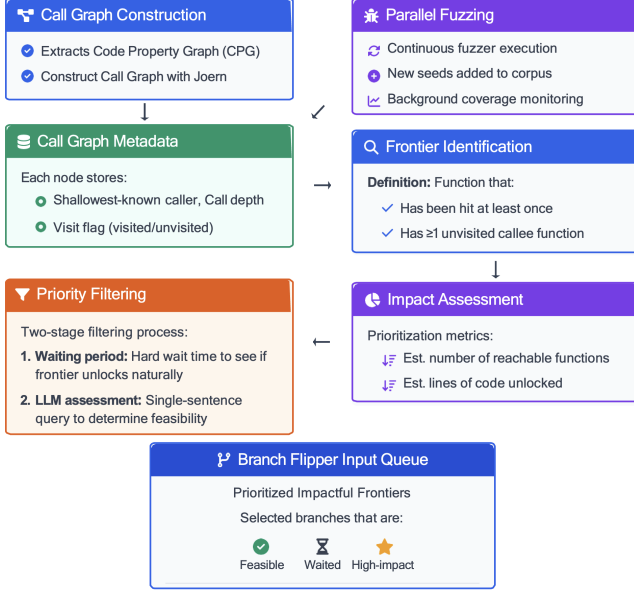


Figure 6: Workflow from collecting frontiers to transporting to the branch flipper agent. Frontiers are collected by the background coverage monitor with a call graph and assessed for their impact through the estimated reachable functions that have not been explored before.

### 3.5 Grounding LLM with Coverage Feedback

After the target branch predicate has been localized and the LLM generates a corresponding mutated seed, we perform a grounding step to evaluate the effectiveness of the attempt.

Once a new seed is generated, we execute the fuzz harness with it and collect coverage information. If the target function is reached or the unhit target branch is taken, the seed is marked as a success and proceeds to the next branch candidates. However, if coverage does not reflect success, it is crucial to inform the LLM about the progress made in execution.

Specifically, we compute the nearest hit line before and after the target branch and feed this information back to the LLM. This allows the agent to estimate which portions of the program were executed and where progress was stuck. The LLM then re-analyzes the code and attempts a new seed generation along with some debugging trials.

However, we observe that even after multiple failures, the LLM may hallucinate success, claiming that the predicate was likely satisfied and that coverage measurement is faulty. In these cases, the LLM often continues to generate seeds under false assumptions.

To counteract this, we enforce a strict grounding. If coverage data indicates that the target was not reached, the system requests at least one additional retry, regardless of the LLM’s confidence. By introducing retry policies, we improve the reliability and stability.

### 3.6 Further Details

Unlike centralized architectures where a single orchestrator governs all tool calls, our system adopts a peer-to-peer agentic model, in which LLM agents with distinct objectives communicate through message queue interfaces. Each agent is self-directed, performing localized reasoning tasks and initiating tool calls or requests to other agents as necessary.

The background fuzzer continuously monitors the evolution of the seed corpus. Upon detecting new frontiers, it triggers a localization routine and stacks it as a candidate. The seed generator pulls a candidate and tries to generate a new seed. It can request coverage verification from the background fuzzer, forming a mutually dependent workflow. This decentralized architecture removes the bottleneck of a single orchestrator and supports a more scalable, fault-tolerant execution flow.

The other challenge is managing context length, especially when operating on low-level codebases, such as those written in C. In these settings, individual functions often span dozens of lines, and naively injecting full source code can rapidly exhaust the context length, degrading performance.

To address this, we implement a dedicated Source Review Agent. When an LLM agent requires a detailed understanding of specific code regions (e.g., to infer predicate conditions or data structures), it queries the Source Review Agent instead of embedding raw code into its own prompt. This subagent handles code lookup and semantic summarization. Thus, the parent agent can operate within a compact and structured representation of the source, retaining high-fidelity analysis.

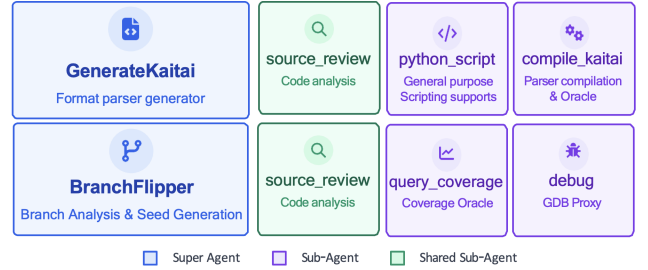


Figure 7: LLM Agents and their corresponding tools.

In practice, we also observe LLM’s tool usage biases, where only a few tools are consistently invoked. To make the system more consistent and compact, initially, we added all potentially useful tools and observed the tool call tendencies. We ensure that frequently used tools execute before invoking the agent and inject the execution result into the agent prompt, if possible. For example, initially, we gave a path to the harness, and we observed that the agent always read the harness first. We added the harness code to the prompt and removed the harness reading tools. On the other hand, rarely used or optional tools are either moved to subagents or completely removed. This strategy reduces tool overload while ensuring

necessary data is available without active tool invocation.

We structure the overall execution into a tree of LLM sub-agents, each responsible for a narrow task domain. Upon the system initialization, a target project is loaded, and Kaitai and Python-based encoders and decoders are generated and stored in a centralized database. During fuzzing, when a frontier is detected, the localizer agent identifies the corresponding branches and saves them to a candidate database. The branch flipper pulls the most impactful candidate and tries to generate a seed. The flipper receives the harness code, a stochastically chosen encoder, decoder, and the definition of the target function, as well as the branch location and a visualized version of the input seed. It solves subtasks using three primary tools: the Source Review Agent, the Debugger Agent (which receives high-level debugging instructions and operates GDB to gather information), and the Coverage Query Agent.

## 4 Experiments

We conducted experiments on several projects, targeting C/C++-based applications such as Curl, Nginx, libpng, and Zstd, as well as the Java-based application Tomcat. These targets include both benchmark applications provided by AIXCC and external open-source projects.

**Generating Kaitai Descriptor:** The first step is generating Kaitai descriptors. The LLM Agent is composed of tools such as `source_question`, `python_script`, and `compile_kaitai` to produce the Kaitai descriptor. For the Python-based encoder, the `compile_kaitai` tool is replaced with a `sanity_check` tool, which also performs oracle-like checks using the existing corpus. For each project, we selected harnesses and attempted to generate both encoders and decoders.

As seen in the Figure 3, 4, while the Python-based encoders often missed details like the PNG header and just formed an uninformative identity function, such details are clearly represented in the generated Kaitai descriptors. It seems that the LLMs already know the published descriptors and utilize them. For targets like Tomcat, where `FuzzedDataProvider` is used, we found that byte stream issues, such as endianness, were more accurately handled after integrating Kaitai. Especially in the case of Zstd, which has a highly complex `FuzzedDataProvider` logic, the original Python encoder is often reduced to trivial identity mappings. In contrast, the LLM was encouraged to generate more sophisticated byte descriptions when utilizing Kaitai-struct, even if partially.

We then used these encoders to attempt PoC seed generation, in the Table 1, as part of the AIXCC competition. During seed generation, we observed that it successfully triggers vulnerabilities with the generated PoC when using Kaitai descriptors. Since these results were complementary to the ones generated with Python-based encoders, this provides indirect evidence that using both in a dual-path setup is a meaningful strategy for branch flippers.

ID	Baseline	Kaitai
Curl#48	Pass (0.170\$, 7min)	Pass (0.697\$, 8min)
Curl#49	Pass (0.415\$, 13min)	Pass (1.372\$, 12min)
libpng#37	Pass (0.247\$, 4min)	Pass (0.392\$, 4min)
libpng#38	Pass (0.217\$, 3min)	Pass (0.326\$, 3min)
libpng#39	Pass (0.267\$, 4min)	Pass (0.897\$, 8min)
nginx#0	Fail (2.664\$, 24min)	Pass (0.637\$, 11min)
nginx#11	Pass (0.238\$, 3min)	Pass (0.735\$, 6min)
nginx#21	Pass (0.302\$, 3min)	Pass (0.989\$, 12min)
nginx#34	Pass (0.266\$, 2min)	Fail (2.577\$, 19min)
nginx#44	Pass (1.611\$, 16min)	Pass (0.869\$, 7min)
nginx#74	Fail (2.477\$, 26min)	Fail (1.962\$, 15min)
nginx#88	Pass (0.257\$, 2min)	Pass (0.532\$, 5min)
tomcat#38	Fail (2.012\$, 30min)	Pass (0.900\$, 8min)
tomcat#39	Fail (0.843\$, 17min)	Fail (0.724\$, 17min)
tomcat#40	Fail (2.068\$, 30min)	Fail (1.331\$, 20min)
zstd#47	Pass (0.144\$, 2min)	Pass (0.516\$, 5min)
zstd#48	Fail (1.895\$, 20min)	Fail (1.397\$, 11min)
zstd#49	Pass (0.310\$, 4min)	Fail (1.985\$, 20min)

Table 1: Trials of PoC generation for given vulnerability with several variants. *ID* represents each vulnerabilities, *Baseline* use Python encoder only when generate a binary seed, *Kaitai* utilizes Kaitai-struct only.

**Branch Flipper:** Next, we moved on to exploring branch blockers and executing the flipper. Before we built an algorithmic localizer, we had the LLM directly localize the targeted branches. However, case studies revealed that more than half of the attempts failed to correctly localize the branches. Algorithmic localization, on the other hand, significantly reduced such mislocalizations.

Project	#Trials	#Success
Curl	13	7 (53.8%)
Nginx	11	3 (27.3%)
Tomcat	10	7 (70.0%)
PoC Generation	13	9 (69.2%)

Table 2: The success rates of precollected branch blockers.

We collected branch blocker cases that were ensured to be solvable for each harness and attempted to flip them. Among the branch blockers, some were semantically trivial, and the branch flipper was able to solve those cases without difficulty. However, in real cases, there were also instances where hard-coded configurations or arguments in the harness prevented child function invocation, or where flipping failed in some cases despite being considered solvable. Nevertheless, given that each project accumulates over 10,000 branch blockers per hour, solving even just the not-that-difficult cases makes a reasonable and effective complement to random mutation-based greybox fuzzing.

We also investigated the applicability of our method to PoC



ID	Baseline	+ Branch Flipper
Curl#48	Pass (0.170\$, 7min)	Pass (0.126\$, 11min)
Curl#49	Pass (0.415\$, 13min)	Pass (0.436\$, 14min)
libpng#37	Pass (0.247\$, 4min)	Pass (0.171\$, 5min)
libpng#38	Pass (0.217\$, 3min)	Pass (0.158\$, 3min)
libpng#39	Pass (0.267\$, 4min)	Pass (0.348\$, 9min)
nginx#0	Fail (2.664\$, 24min)	Pass (0.208\$, 4min)
nginx#11	Pass (0.238\$, 3min)	Pass (0.245\$, 6min)
nginx#21	Pass (0.302\$, 3min)	Pass (0.305\$, 8min)
nginx#34	Pass (0.266\$, 2min)	Fail (1.496\$, 15min)
nginx#44	Pass (1.611\$, 16min)	Pass (0.266\$, 4min)
nginx#74	Fail (2.477\$, 26min)	Pass (0.795\$, 11min)
nginx#88	Pass (0.257\$, 2min)	Pass (0.216\$, 4min)
tomcat#38	Fail (2.012\$, 30min)	Pass (0.359\$, 15min)
tomcat#39	Fail (0.843\$, 17min)	Pass (1.641\$, 46min)
tomcat#40	Fail (2.068\$, 30min)	Fail (1.354\$, 23min)
zstd#47	Pass (0.144\$, 2min)	Pass (0.102\$, 3min)
zstd#48	Fail (1.895\$, 20min)	Fail (1.342\$, 13min)
zstd#49	Pass (0.310\$, 4min)	Pass (0.804\$, 8min)

Table 3: Trials of PoC generation for given vulnerability with and without Branch Flipper. *ID* represents each vulnerabilities, *Baseline* without branch flipper and *+ Branch Flipper* with branch flipper when generating PoC.

generation. In practice, we observed that while PoC generators could often identify suspicious or vulnerable code regions, they failed to produce seeds that could successfully exploit them. A common cause was the inability to reach the target functions. To address this, we integrated our branch flipper tool into the PoC generation pipeline. Experimental results 3 showed that our tool enabled successful PoC generation in cases where the original system had previously failed due to unreachable functions. Moreover, beyond simply reaching the target functions, the branch flipper enabled the process to be decoupled into two stages: generating inputs that reach the vulnerable function and then modifying them to become exploitable, resulting in improved overall performance in some cases.

## 5 Conclusion

We propose a cooperative agentic system designed to unlock fuzz blockers in a greybox fuzzing environment. The system operates with the goal of achieving coverage over previously unhit functions by localizing blockers and generating targeted seed inputs that can take the obstructing branches.

To support this, we suggest using an LLM to generate a binary encoder and decoder, utilizing a Python sandbox and the Kaitai-struct compiler. The LLM leverages subagents to analyze the execution path, such as a source code review agent and a GDB-based debugging agent. Based on reference seeds that partially reach the target region, the LLM generates new

candidate seeds and evaluates them iteratively, incorporating coverage-based grounding to mitigate hallucinations.

Using this system, we successfully unlocked real-world fuzz blockers, contributing as one of the components in the AIXCC challenge pipeline. However, we observed that many blockers remain unsolvable, and in such cases, continued operation of the LLM-based agent can lead to inefficient use of computational resources and budget. This highlights a next step of future work: developing a method to predict the solvability of fuzz blockers, thereby improving resource efficiency in autonomous fuzzing workflows.

## References

- [1] BESTA, M., BLACH, N., KUBICEK, A., GERSTENBERGER, R., PODSTAWSKI, M., GIANINAZZI, L., GAJDA, J., LEHMANN, T., NIEWIADOMSKI, H., NYCZYK, P., AND HOEFER, T. Graph of thoughts: Solving elaborate problems with large language models. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 16 (Mar. 2024), 17682–17690.
- [2] BROWN, T., MANN, B., RYDER, N., SUBBIAH, M., KAPLAN, J. D., DHARIWAL, P., NEELAKANTAN, A., SHYAM, P., SASTRY, G., ASKELL, A., AGARWAL, S., HERBERT-VOSS, A., KRUEGER, G., HENIGHAN, T., CHILD, R., RAMESH, A., ZIEGLER, D., WU, J., WINTER, C., HESSE, C., CHEN, M., SIGLER, E., LITWIN, M., GRAY, S., CHESSE, B., CLARK, J., BERNER, C., MCCANDLISH, S., RADFORD, A., SUTSKEVER, I., AND AMODEI, D. Language models are few-shot learners. In *Advances in Neural Information Processing Systems* (2020), H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33, Curran Associates, Inc., pp. 1877–1901.
- [3] CHEN, P., XIE, Y., LYU, Y., WANG, Y., AND CHEN, H. Hopper: Interpretative fuzzing for libraries. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2023), CCS ’23, Association for Computing Machinery, p. 1600–1614.
- [4] CHEN, X., PUN, C.-M., AND WANG, S. Medprompt: Cross-modal prompting for multi-task medical image translation, 2023.
- [5] CLAUDE. Claude 3.7 sonnet and claude code.
- [6] DEEPSEEK-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.
- [7] DENG, Y., XIA, C. S., PENG, H., YANG, C., AND ZHANG, L. Large language models are zero-shot fuzzers: Fuzzing deep-learning libraries via large language models. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis* (New York, NY, USA, 2023), ISSTA 2023, Association for Computing Machinery, p. 423–435.
- [8] FANG, R., BINDU, R., GUPTA, A., ZHAN, Q., AND KANG, D. Llm agents can autonomously hack websites, 2024.
- [9] GOOGLE. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025.
- [10] GUAN, X., ZHANG, L. L., LIU, Y., SHANG, N., SUN, Y., ZHU, Y., YANG, F., AND YANG, M. rstar-math: Small llms can master math reasoning with self-evolved deep thinking, 2025.
- [11] HSIEH, C.-P., SUN, S., KRIMAN, S., ACHARYA, S., REKESH, D., JIA, F., ZHANG, Y., AND GINSBURG, B. Ruler: What’s the real context size of your long-context language models? *arXiv preprint arXiv:2404.06654* (2024).
- [12] HU, J., ZHANG, Q., AND YIN, H. Augmenting greybox fuzzing with generative ai, 2023.
- [13] ISOZAKI, I., SHRESTHA, M., CONSOLE, R., AND KIM, E. Towards automated penetration testing: Introducing llm benchmark, analysis,

- and improvements. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization* (New York, NY, USA, 2025), UMAP Adjunct '25, Association for Computing Machinery, p. 404–419.
- [14] JOERN.IO. Joern: The Bug Hunter’s Workbench, Jan. 2024.
  - [15] LLVM. libfuzzer – a library for coverage-guided fuzz testing.
  - [16] LYU, Y., XIE, Y., CHEN, P., AND CHEN, H. Prompt fuzzing for fuzz driver generation, 2024.
  - [17] MENG, R., DUCK, G. J., AND ROYCHOUDHURY, A. Large language model assisted hybrid fuzzing, 2024.
  - [18] MENG, R., MIRCHEV, M., BÖHME, M., AND ROYCHOUDHURY, A. Large language model guided protocol fuzzing. In *Proceedings of the 31st Annual Network and Distributed System Security Symposium (NDSS)* (2024).
  - [19] OPENAI. Introducing o3 and o4-mini.
  - [20] OUYANG, L., WU, J., JIANG, X., ALMEIDA, D., WAINWRIGHT, C., MISHKIN, P., ZHANG, C., AGARWAL, S., SLAMA, K., RAY, A., SCHULMAN, J., HILTON, J., KELTON, F., MILLER, L., SIMENS, M., ASKELL, A., WELINDER, P., CHRISTIANO, P. F., LEIKE, J., AND LOWE, R. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems* (2022), S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35, Curran Associates, Inc., pp. 27730–27744.
  - [21] PROJECT, K. Kaitai struct: declarative binary format parsing language, 2025.
  - [22] REESE, W. Nginx: the high-performance web server and reverse proxy. *Linux J.* 2008, 173 (Sept. 2008).
  - [23] SEREBRYANY, K. OSS-Fuzz - google’s continuous fuzzing service for open source software. USENIX Association.
  - [24] SHI, W., ZHANG, Y., XING, X., AND XU, J. Harnessing large language models for seed generation in greybox fuzzing, 11 2024.
  - [25] SHOU, C., LIU, J., LU, D., AND SEN, K. Llm4fuzz: Guided fuzzing of smart contracts with large language models, 2024.
  - [26] VASWANI, A., SHAZEER, N., PARMAR, N., USZKOREIT, J., JONES, L., GOMEZ, A. N., KAISER, L. U., AND POLOSUKHIN, I. Attention is all you need. In *Advances in Neural Information Processing Systems* (2017), I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30, Curran Associates, Inc.
  - [27] WANG, X., WEI, J., SCHUURMANS, D., LE, Q. V., CHI, E. H., NARANG, S., CHOWDHURY, A., AND ZHOU, D. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations* (2023).
  - [28] WEI, J., WANG, X., SCHUURMANS, D., BOSMA, M., ICHTER, B., XIA, F., CHI, E., LE, Q. V., AND ZHOU, D. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems* (2022), S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35, Curran Associates, Inc., pp. 24824–24837.
  - [29] XIA, C. S., PALTENGHI, M., LE TIAN, J., PRADEL, M., AND ZHANG, L. Fuzz4all: Universal fuzzing with large language models. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering* (Apr. 2024), ICSE ’24, ACM, p. 1–13.
  - [30] YAMAGUCHI, S. Gnu global.
  - [31] YAO, S., YU, D., ZHAO, J., SHAFRAN, I., GRIFFITHS, T., CAO, Y., AND NARASIMHAN, K. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems* (2023), A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36, Curran Associates, Inc., pp. 11809–11822.
  - [32] YAO, S., ZHAO, J., YU, D., DU, N., SHAFRAN, I., NARASIMHAN, K. R., AND CAO, Y. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations* (2023).
  - [33] ZALEWSKI, M. American fuzzy lop.