

Visualizing The Intermediate Representation of Just-in-Time Compilers

HeuiChan Lim and Stephen Kobourov

Department of Computer Science, University of Arizona

Abstract. Just-in-time (JIT) compilers are used by many modern programming systems in order to improve performance. Bugs in JIT compilers provide exploitable security vulnerabilities and debugging them is difficult as they are large, complex, and dynamic. Current debugging and visualization tools deal with static code and are not suitable in this domain. We describe a new approach for simplifying the large and complex intermediate representation, generated by a JIT compiler and visualize it with a metro map metaphor to aid developers in debugging.

Experiments using our prototype implementation on Google’s V8 JavaScript interpreter and TurboFan JIT compiler demonstrate that it can help identify and localize buggy code.

1 Introduction

Many modern programming systems, such as JavaScript engines that are running our web browsers, use just-in-time (JIT) compilers to improve performance. The most well-known web browsers that we use every day are Google Chrome’s V8 [22], which has 2.65 billion users [5], Microsoft Edge’s ChakraCore [2], which has 600 million users [16], Mozilla Firefox’s SpiderMonkey [17], which has 222 million users [4], and Apple Safari’s WebKit [14], which has 446 million users [7], etc., Additionally, programming languages such as Java, C#, and Ruby, also use JIT compilers to improve performance.

Meanwhile, the JIT compiler bugs can lead to exploitable security vulnerabilities [1, 9–11, 13]. For example, a JIT compiler bug in Google V8 was found by the Microsoft Offensive Security Research team (CVE-2017-5121 [1]). Fortunately, this was reported quickly and fixed by V8 team, but hackers could have used it to hijack Chrome user data (i.e., passwords), and to navigate to other sites and execute malicious programs. Therefore, it is important to analyze and localize the bug quickly. However, existing work and available tools focus on static code [23, 24, 32], so they are not suitable for developers in debugging the JIT compiler, which generates code at run-time. Additionally, the size and complexity of JIT-based systems [19] combined with the dynamic nature of JIT compiler optimizations, make it challenging to analyze and locate bugs quickly. For example, Google V8 has more than 2,000 source files and more than 1 million lines of code, and Apple’s Webkit has approximately 22,000 source files and 500,000 lines of code.

Traditional debuggers rely on text even though the main feature of a JIT compiler is building a graph-like structure to translate bytecode into optimized machine code. With this in mind, we propose a new debugging aid tool, which visualizes the JIT compiler’s intermediate representation (IR) that it generates and optimizes. Our approach uses IR identification and generation techniques described by Lim and Debray [35], which discusses the compiler-related half of the visualization tool’s pipeline. This paper focuses on the visualization half, which includes: (1) merging multiple IR graphs into a single graph; (2) simplifying the merged graph; (3) converting the simplified graph into a hypergraph; (4) simplifying the hypergraph; and (5) visualizing the hypergraph using a metro map metaphor [34].

The resulting visualization of the intermediate representation of the JIT compiler allows developers to answer questions such as:

1. What optimizations took place to generate the machine code?
2. What is the relationship among the optimization phases?
3. Which optimization phase was most active?
4. With a specific node that represents some operation, which optimization(s) took place (e.g., what phase optimized an arithmetic operator node for a subtraction)?
5. Which optimization phases are likely to be buggy?

2 Related Work

Although there is a considerable body of work on debugging approaches for static code compilers and optimized code, there is very little work on using the intermediate representation and visualizing it to show the explicit information about the compilation and optimization processes.

Google V8’s Turbolizer [8, 20] is one of very few IR visualization tools. It shows the final IR graph after each optimization process and provides interactive features to view the control flow graphs for each optimization phase. Although Turbolizer provides some information about the IR nodes and their relationships, it does not provide enough information about the optimization process and cannot answer several of our initial set of questions.

Dux *et al.* [30] introduced an approach to visualize dynamically modifying code at run-time in call graph and control flow graphs. Their tool shows the changes in graph with an animation of the underlying graph, allowing end-to-end play, pause, and forward/backward step-by-step animation. Colberg *et al.* [25] proposed GEVOL: a system that visualizes the evaluation of software. It uses the CVS version control system to get the information about how a program has been changed or evolved over time and visualizes it as a dynamic graph. CFGExplorer [29] visualizes the control flow graph of a program to represent the program structure for dynamic binary analysis. It provides interactive features where users can find specific memory addresses, loops, and functions that they are interested to analyze the system. CcNav [28] analyzes and visualizes a C++

compiler’s optimization process with a call graph, control flow graph, and loop hierarchies.

Control flow graph and call graphs are popular in program analysis, especially, when analyzing static code, with a code-to-compiler-to-assembler pipeline. However, these graphs are different from the intermediate representation generated dynamically from bytecode by JIT compilers. Tools for visualizing and interacting with control flow graph and call graphs (such as those above) are not sufficient for visualizing the IR graph as, e.g., they cannot capture the optimization phases.

3 Background

Here we briefly introduce several concepts relevant to JIT compilers and to our visualization.

Interpreter: An interpreter is a computer program that converts input source code into bytecode and executes it without compiling it into a machine code [31].

Just-in-Time (JIT) compiler: A just-in-time (JIT) compiler is a program that turns bytecode into instructions that can be sent directly to a computer’s processor, which is widely used to improve performance [33]. Google V8 is an example of a JavaScript engine that has a JIT compiler in the pipeline; see Fig. 1 (a). Other systems, such as ChakraCore, JavaScriptCore, and Java Runtime Environment, have similar pipelines.

Bytecode: Bytecode is an instruction generated from input source code by an interpreter. Bytecode is portable, unlike compiled programs, so they are widely used in modern languages and systems, such as JavaScript, Python, and Java, etc. [26].

Optimized code: Optimized code is an instruction generated from bytecode by a JIT compiler. Optimized code is machine code that can be directly executed by a processor. Unlike bytecode, optimized code is system-specific, and much faster in execution.

Intermediate Representation (IR): The IR is type of representation known as sea-of-nodes [18, 21, 27]. Unlike other graphs used in the program analysis field, such as control-flow or data-flow graphs which have specific types of nodes, nodes in the sea-of-nodes graph represent different types of nodes: from scalar values and arithmetic operators to variables and control-flow nodes and function entry nodes. Consider, for example, the expression `var myVar = 100;` optimized at phase α . This generates 3 different types of nodes: a variable node for `myVar`, an arithmetic operator node for `Equal`, and a scalar value node for constant `100`. Other types of nodes in the IR include control-flow nodes (e.g., `call`, `branch`, `switch`) and function entry nodes (e.g., `Start`, `Return`, `Checkpoint`).

Optimization: In the context of JIT compilers, optimization means adding, removing, and merging nodes and edges in the graph during the execution. In a single JIT compilation, the JIT compiler executes several different phases of

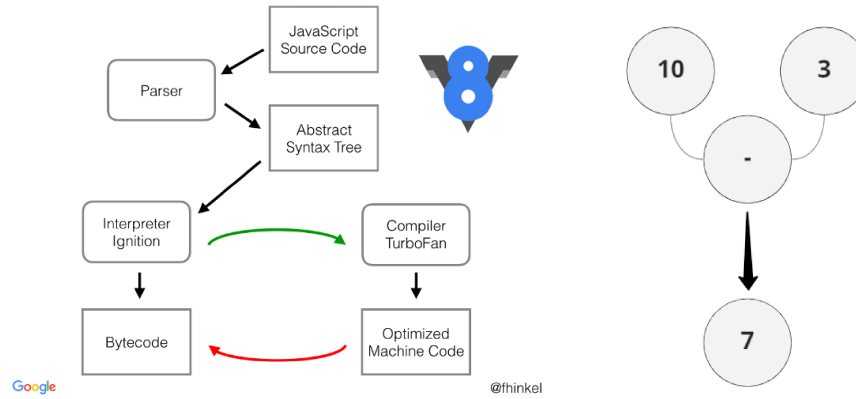


Fig. 1. (a) V8 Pipeline [19] (b) Example of constant folding optimization: if the compiler determines that the value of $10-3$ will not change during the code execution, it is more efficient to replace the 3-node subgraphs with a single node.

optimization (i.e., inlining, loop peeling, constant propagation) to generate efficient machine code; see Fig. 1(b) for a constant propagation example. Each optimization phase modifies the IR graph and corresponds to a new hyperedge (all the nodes generated or optimized in this phase); we store all of them.

Proof-of-Concept Program: A “Proof-of-Concept” (PoC) program is an input program that is used to trigger the buggy behavior in the JIT compiler, i.e., a valid program (without any bugs) which when run can reveal bugs in the JIT compiler. In our experiment, we are targeting JavaScript engine V8, so the PoC that we are using is a JavaScript program.

Instruction-level trace: An instruction-level trace is a file that holds all the instructions that a programming system, such as a JIT compiler, has generated and executed at run-time. The instructions are in a machine-level code with symbol information (i.e., function names). Such traces are used for analyzing the executed system’s status, such as performance and behavior, as well as for debugging.

4 Visualizing the Intermediate Representation

Our approach to capturing and visualizing the intermediate representation of a JIT compiler, starting from the input JavaScript source code, consists of the following steps:

1. Start by automatically modifying the original input JavaScript program P_0 , where P is program, to create a set of N similar input programs $\{P_1, \dots, P_N\}$. The likelihood of identifying the suspicious buggy phase increases as N increases (by default we use $N = 20$). The program variants are created by

first generating the abstract syntax tree for the original input code and then randomly modifying nodes in the tree, with a set of allowable modifications as described in [35]. Consider, for example, an expression that adds two numbers $2 + -1$. We can modify 2 or 1 to any other integer. We can change the operation $+$ to another arithmetic operation. All such allowable modifications must pass semantic and syntax checks.

2. Run each program P_i and collect the instruction-level traces.
3. Analyze the traces to determine whether or not P_i manifested the bug and to identify P_i 's IR within the JIT compiler, together with the optimization phases executed while optimizing P_i .
4. Select candidate hyperedges, suspected to be buggy, from the information gathered in step 3.
5. Merge all selected candidate hyperedges into the original IR from P_0 .
6. Simplify the merged IR as a graph by reducing the number of nodes and edges.
7. Convert the simplified graph into a hypergraph by extracting the hyperedges from step 4 and analyzing each node's optimization status. Consider, for example, a node A generated in hyperedge α and optimized in hyperedge ϕ and γ ; in this step A (which is in α from the start) is added to ϕ and γ .
8. Simplify the hypergraph by reducing the number of hyperedges and nodes.
9. Visualize the simplified hypergraph with MetroSets [34].

The compiler-related steps 1-4 above are covered in detail in [35] and details for steps 5-9 are in the next section.

4.1 Intermediate Representation

Recall that the intermediate representation (IR) of a JIT compiler is a sea-of-nodes graph that the compiler generates at the beginning of its execution by parsing the bytecode and optimizing it with several optimization phases.

Formally, the IR is a simple, undirected graph $G = (V, E)$, where V represents the nodes optimized by the JIT compiler and E contains pairs of nodes connected by different relationships (e.g., semantic and syntax relationships, such as math expressions). By keeping track of the optimization information for each node we construct the hypergraph $H = (V, S)$ from G , where V is a set of nodes optimized by the JIT compiler and each hyperedge in S represents an optimization phase.

Two important node features are phases and opcodes. Phases are the optimization phases where a node was generated and optimized (and which later correspond to hyperedges). Opcodes represent node operations (e.g., add, sub, return). A node also has two different attribute groups: (1) *basic*, such as a node id, address, list of neighbors, opcode, and IR ID (the ID of IR that the node belongs to); and (2) *optimization*, such as hyperedge (phase) ID, hyperedge name, replace/kill/remove/append phase.

Recall that given one JavaScript code we generate N similar versions to see if any of them trigger bugs. We generate the IRs for all of these versions (typically about 20). In the real-world examples we work with, each such IR graph has about 300-500 nodes and 30-40 optimization phase executions.

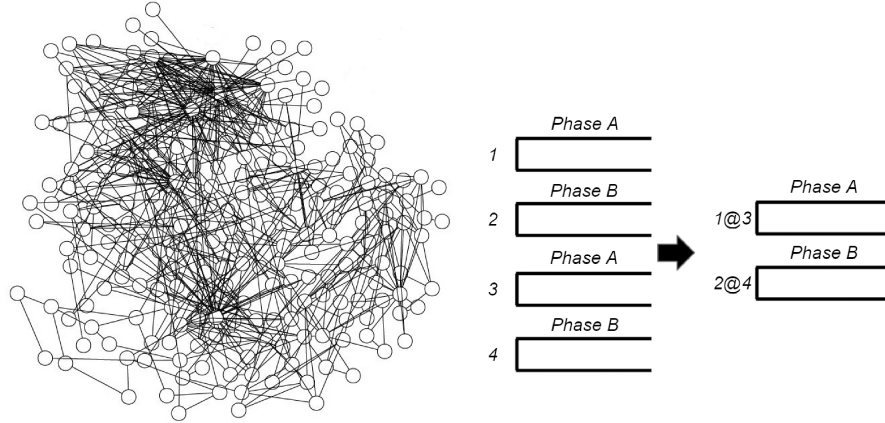


Fig. 2. (a) Example of an IR graph; (b) example of hypergraph simplification.

4.2 Merging Intermediate Representation Hyperedges

Recall that we now have N intermediate representations from our original JavaScript program. They are similar, but different, and in this step of our pipeline we merge them all into one single graph by merging hyperedges. There are two main reasons for this IR merging. First, we would like to visualize the differences among the graphs in one single view. Second, as shown in [35], by comparing hyperedges from a buggy program IR to hyperedges from a non-buggy program IR, we can find differences in some hyperedges due to different optimizations, and thus find the bug. Consider, for example, a hyperedge α in both buggy and non-buggy program IRs and suppose that an additional node (the result of incorrect optimization) makes a buggy program's α different from the non-buggy program's α . A merged hyperedge will show this additional node and its attributes will identify the buggy IR. A developer can now see that there was an optimization difference in α between the buggy and non-buggy programs, check whether this difference causes the buggy behavior.

We use the following procedure to merge IRs. Let R_0 be the IR from the original program and $\{R'_1, \dots, R'_N\}$ the IRs from the modified programs. Let $\{r'_1, \dots, r'_n\}$ be sub-IRs, where r'_i is a subgraph of R'_i when $R'_i \neq R_0$; i.e., n is the number of IRs different from R_0 ($n \leq N$). Each r'_i holds candidate hyperedges: R'_i hyperedges different from R_0 's hyperedges. We traverse all sub-IRs, comparing each to R_0 , and update the merged IR. If $s'_i \in r'_i$ and $s'_i \notin R_0$, where s'_i is a hyperedge, we add s'_i to R_0 . If $s'_i \in r'_i$ and $s_i \in R_0$, where s_i has same hyperedge name and execution order with s'_i , but $s'_i \neq s_i$, we merge s'_i with s_i ; see Algorithm 1 in Supplementary Materials for details.

4.3 Intermediate Representation Simplification

Although a graph with all the nodes and all the edges may be useful for debugging, the complexity, as we can see in Fig. 2 (a) makes it difficult for developers

to use. Therefore, we simplify the graph (hopefully without losing much information), convert it into a hypergraph, and simplify the hypergraph (again, hopefully without losing much information). The main goal is to end up with an interactive visualization that allows developers to debug. Our IR simplification has two parts: (1) Reduce IR as a graph, (2) Reduce IR as a hypergraph.

Reducing the IR Graph We remove dead nodes (nodes with no adjacent edges) as they are not translated into machine code and do not affect other nodes. We then identify nodes that can be merged (reducing the total number of nodes and edge) without losing important information. A pair of nodes is merged if they have the same opcode, the same optimization information, belong to the same IR (which can be identified by the IR id attribute), and share the same neighbors; see Fig. 3. We first traverse each node in the graph and check for dead nodes and removes them. We then re-traverse the graph and compare each node to rest of the nodes, looking for pairs that can be merged. Algorithm 2 shows a pseudo-code for how the graph simplification is implemented.

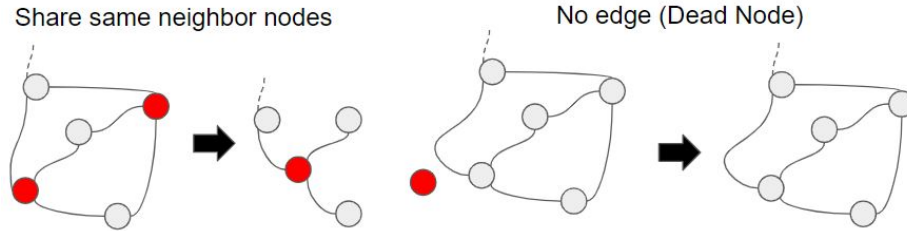


Fig. 3. Reducing the IR Graph: two red nodes that share the same neighbors are merged (left) and a red dead node is removed (right).

Reducing the IR Hypergraph We first convert the simplified graph $G = (V, E)$ into a hypergraph $H = (V, S)$, by extracting hyperedges based on the optimization phases. Recall that a node A generated in phase/hyperedge α and optimized in phase/hyperedge ϕ and γ now belongs to all three hyperedges. We reduce the hypergraph by merging suitable pairs of hyperedges. Different nodes can have the same hyperedge names as attribute, but different hyperedge id, as this id is assigned based on the execution order. Therefore, we merge hyperedges with the same name into a single-hyperedge while assigning a new unique identifier generated from the original IDs. We use ID concatenation to obtain unique identifiers. Fig. 2(b) illustrates hyperedge merging. Consider two hyperedges A and B , which were executed twice in the order shown. The order is used to create the unique IDs. We merge these 4 hyperedges into two larger hyperedges and assign new IDs that were generated by concatenating two IDs

delimited with a special character '@'. Algorithm 4 provides an overview of this procedure.

This method significantly reduces the number of hyperedges that we have to deal with but increases the number of nodes in each hyperedge. Next, we traverse each hyperedge s in S , and we use the nodes opcodes to see if they can be merged. Algorithm 5 shows the hypergraph simplification.

5 MetroSets

MetroSets [34] is a tool for visualizing hypergraphs using the metro map metaphor. It provides a clean way to show the relationships among the hyperedges, which in our case represent the relationships among the optimizations. Further, MetroSets provides simple and intuitive interactions that make it possible to quickly identify hyperedges that contain a suspicious (buggy) node, or hyperedges that intersect with a particular suspicious (buggy) hyperedge. MetroSets is designed for small-to-medium set systems with no more than 500 nodes and no more than 30 hyperedges. Additional internal requirements limit the number of nodes in each hyperedge to at least 2 and require that all hyperedges must intersect at least 1 other hyperedge. Since some of our hyperedges can be singletons, we add a dummy node when needed (which is later ignored). If one or more of our hyperedges do not intersect any others, they can be treated separately and shown next to the main hypergraph.

Each node in the MetroSet map is labeled with its unique id (representing the timeline of node generation). The attributes shown when hovering over a node are phase, opcode, address, graph id, and phase id. A phase attribute tells the user where the node was generated, and it is useful for the node belongs to multiple sets. The user can distinguish the phase that it was generated and the phases where it was optimized. The phase id helps us analyze the order in which nodes along a given line were generated.

6 Evaluation

The following experiments were done with our prototype system using a machine with 32 cores (@ 3.30 GHz) and 1TB of RAM, running Ubuntu 20.04.1 LTS. We used a dynamic analysis tool built on top of Intel’s Pin software [36] for program instrumentation and collecting instruction-level execution traces, XED [6] for instruction decoding [6]. Additionally, we used esprima-python [15] to generate the syntax-tree for JavaScript code, and escodegen [3] to regenerate the JavaScript code from the syntax-tree.

Our prototype targets Google’s JavaScript engine V8, focusing in particular on TurboFan, V8’s JIT compiler.

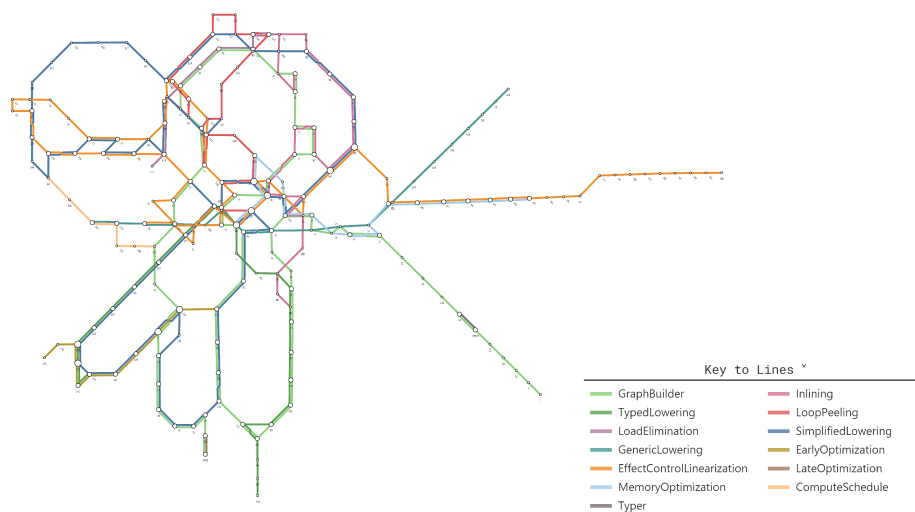
We use data from the Chromium bug report site; details about the data, bug descriptions, etc. can be found here [35]. We are able to localize the bugs in all of the listed bug reports.

Table 1 shows the result of the simplification step on the bug report data.

Table 1. Graph simplification result table

Report#	Original Set#	Reduced Set#	Original Element#	Reduced Element#
5129	38	13 (-59.52%)	1695	355 (-79.06%)
8056	36	13 (-57.50%)	371	176 (-52.56%)
791245	31	13 (-51.43%)	433	198 (-54.27%)
961237	36	13 (-57.50%)	478	191 (-60.04%)
1072171	40	18 (-50.00%)	701	326 (-53.50%)
880207	33	13 (-54.05%)	555	235 (-57.66%)

On average 50-60% reductions in the number of nodes and hyperedges is achieved, resulting in manageable inputs sizes for MetroSets.

**Fig. 4.** Metro map of IR graph from bug report 5129

Case Study: Here we briefly go over one example showing how our visualization can assist developers in debugging.

The data comes from JavaScript source code (chromium bug report number 5129). This version of V8's JIT compiler has a problem in the *EarlyOptimization* phase. We generate 18 additional modified JavaScript codes from the original source code and ran all 19 programs with V8's local executable program called *d8*. The resulting individual instruction trace files are used to generate the IR

graph shown in Fig. 2(a). Now, we apply our 7-step pipeline to obtain the final metro map representation which shows 335 nodes and 18 sets; see Fig. 4.

We can now attempt to answer some of the questions from Section 1.

“What optimizations took place to generate the machine code?” By looking at the set names in “Key to Lines” legend, we can identify which optimization phases took place.

“What’s the relationships among the optimization phases?” and *“With a specific node that represents some operation, which optimizations took place?”* We can answer these two questions by looking at the lines sharing the nodes. Hovering over the nodes of interest highlights the names of phases in the “Key to Lines” legend. We can also select the interactive intersection/union/etc. (exploration modes) to see the relationships among the optimization phases. Fig. 10 shows an example of the relationships between TypedLowering, SimplifiedLowering, and EarlyOptimization, with 3 nodes in their intersection. A developer can investigate the relationship by examining the attributes of each node. The generation order and the hyperedge responsible for the generation can help identify which hyperedge impacted others in the optimization.

“Which optimization phase was most active?” We can answer this question by hovering over each line (which reports the number of nodes in the line). The graph builder phase always has the largest number of nodes but it is not an optimization phase. Other lines with many nodes are the most active optimization phases. Fig. 9 shows an example of a most active optimization phase.

“With a specific node that represents some operation, which optimization(s) took place?” We can answer this question by hovering over the node of interest. The lines that don’t contain the node are grayed out and the displayed node attributes include the opcode, which represents the operation of a node.

“Which optimization phase is likely to be buggy?” One natural way to do this is to find parts that differ in the IR graphs with the bug and those without [35]. In other words, a program is buggy because either it has additional optimizations or missing optimizations, and this information is captured in the IRs. For example, let’s say there is program *A*, which is buggy and program *B*, which is non-buggy. And, both programs were optimized at optimization phase γ . However, the JIT compiler falsely evaluated the code in program *A* and omitted some optimization resulting in missing nodes to be translated into machine codes. These nodes are visible in the non-buggy program *B*’s IR as the JIT compiler has correctly performed optimization. Therefore, any line in the metro map that has a high density of nodes from non-original IRs represents that there was a significant difference between the IRs from the buggy and non-buggy programs. In this case study, we have found that the majority of nodes (9 out of 11) in the EarlyOptimization phase line are from different IRs, which are all from non-buggy IRs. This indicates that there was a significant difference in optimization between buggy and non-buggy programs, so the developers are recommended to begin debugging from the EarlyOptimization phase. Additionally, investigating individual nodes in the suspicious line can be helpful. Figure 11 shows an example

of visualized node details that the user can use to learn about the node, such as which bytecode the node corresponds to and which optimization took place, etc.

7 Discussion, Limitations, and Future Work

We described a new approach for visualizing the intermediate representation of just-in-time compilers using the metro map metaphor and showed how it could be used for debugging. The visualization approach described here is a functional early prototype (available on <https://github.com/hlim1/JITCompilerIRViz>) that provides some useful functionality for debugging JIT compilers. There are many limitations and missing features.

To start with, we target a specific JIT compiler. Generalizing this approach to other JIT compilers (e.g., ChakraCore, SpiderMonkey) can be done by modifying the IR generation process to handle different IR graphs.

After a series of simplifications (of the IR graphs, the merged graph, and the hypergraph) some useful information might be lost. A two-level visualization which shows the simplified hypergraph as an overview but also provides all details on demand will likely be useful.

Currently, we need to hover over each line to identify suspicious phases. This can be improved by applying techniques to localize suspicious phases as described in [35] and automatically highlight/bold such lines.

Along those lines, we might want to provide more detailed information about each node: “why is this node connected to another?”, “what optimization (i.e., removed, added) created this node?”, etc. The answers to several such questions can be collected and provided via appropriate interactions.

8 Conclusion

In this paper, we have introduced a new implementation for visualizing the intermediate representation of just-in-time compiler generates and optimizes at run-time in a metro map. We have introduced algorithms for graph merging and simplification to effectively reduce the size of a large and complex graph into simpler sets while maintaining the important information to aid the developers in debugging.

References

1. Browser security beyond sandboxing (2017). <https://www.microsoft.com/security/blog/2017/10/18/browser-security-beyond-sandboxing/>, accessed 2021-01-22
2. Chakra-core (2016). <https://github.com/chakra-core>, accessed 2021-01-22
3. Estools/escodegen (2012). <https://github.com/estools/escodegen>, accessed 2021-02-03
4. Firefox public data report (2021). <https://data.firefox.com/dashboard/user-activity>, accessed 2021-05-23

5. Google chrome statistics for 2021 (2021).
<https://backlinko.com/chrome-users>, accessed 2021-05-23
6. Intel xed (2019). <https://intelxed.github.io/>, accessed 2021-02-03
7. Internet browser market share 2012-2021 (2021).
<https://www.statista.com/statistics/268254/market-share-of-internet-browsers-worldwide-since-2009/>, accessed 2021-05-23
8. Intro to chrome's v8 from an exploit development angle (2020).
<https://sensepost.com/blog/2020/intro-to-chromes-v8-from-an-exploit-development-angle/>, accessed 2021-02-21
9. Issue 1072171: Security: missing the -0 case when intersecting and computing the type::range in numbermax (2020).
<https://bugs.chromium.org/p/chromium/issues/detail?id=1072171>, accessed 2021-02-01
10. Issue 5129: Turbofan changes $x - y \downarrow 0$ to $x \downarrow y$ which is not equivalent when $(x - y)$ overflows (2016). <https://bugs.chromium.org/p/v8/issues/detail?id=5129>, accessed 2021-02-01
11. Issue 8056: [turbofan] optimized array indexof and array includes ignore a prototype that is not initial (2018).
<https://bugs.chromium.org/p/v8/issues/detail?id=8056>, accessed 2021-02-01
12. Issue 880207: Security: incorrect type information on math.expm1 (2018).
<https://bugs.chromium.org/p/chromium/issues/detail?id=880207>, accessed 2021-02-01
13. Issue 961237: Security: jit difference on comparison in d8 (2019).
<https://bugs.chromium.org/p/chromium/issues/detail?id=961237>, accessed 2021-02-01
14. Javascriptcore (2014). <https://trac.webkit.org/wiki/JavaScriptCore>, accessed 2021-01-22
15. Kronuz/esprima-python (2017). <https://github.com/Kronuz/esprima-python>, accessed 2021-02-03
16. Microsoft edge now has 600 million users (2021).
<https://www.newsbreak.com/news/2138865000496/after-expanding-support-microsoft-edge-now-has-600-million-users>, accessed 2021-05-23
17. Spidermonkey. <https://spidermonkey.dev/>, accessed 2021-02-22
18. Turbofan ir (2016). https://docs.google.com/presentation/d/1Z9iIHojKDrXvZ27gRX51UxHD-bKf1QcPzSijntpMJBm/edit?usp=embed_facebook, accessed 2021-01-21
19. Understanding v8's bytecode (2017).
<https://medium.com/dailyjs/understanding-v8s-bytecode-317d46c94775>, accessed 2021-01-21
20. Using turbolizer to inspect the v8 jit compiler (2019).
<https://lukeolney.me/posts/v8-turbolier/>, accessed 2021-02-22
21. V8: Behind the scenes (2016). <https://benediktmeurer.de/2016/11/25/v8-behind-the-scenes-november-edition>, accessed 2021-01-21
22. V8 javascript engine. <https://v8.dev/>, accessed 2021-01-21

23. Adl-Tabatabai, A., Gross, T.R.: Source-level debugging of scalar optimized code. In: Fischer, C.N. (ed.) Proceedings of the ACM SIGPLAN'96 Conference on Programming Language Design and Implementation (PLDI), Philadelphia, Pennsylvania, USA, May 21-24, 1996. pp. 33–43. ACM (1996). <https://doi.org/10.1145/231379.231388>
24. Brooks, G., Hansen, G.J., Simmons, S.: A new approach to debugging optimized code. In: Feldman, S.I., Wexelblat, R.L. (eds.) Proceedings of the ACM SIGPLAN'92 Conference on Programming Language Design and Implementation (PLDI), San Francisco, California, USA, June 17-19, 1992. pp. 1–11. ACM (1992). <https://doi.org/10.1145/143095.143108>
25. Collberg, C.S., Kobourov, S.G., Nagra, J., Pitts, J., Wampler, K.: A system for graph-based visualization of the evolution of software. In: Diehl, S., Stasko, J.T., Spencer, S.N. (eds.) Proceedings ACM 2003 Symposium on Software Visualization, San Diego, California, USA, June 11-13, 2003. pp. 77–86. ACM (2003). <https://doi.org/10.1145/774833.774844>
26. Dahm, M.: Byte code engineering. In: Cap, C.H. (ed.) JIT '99, Java-Informations-Tage 1999, Düsseldorf 20./21. September 1999. pp. 267–277. Informatik Aktuell, Springer (1999). https://doi.org/10.1007/978-3-642-60247-4_25
27. Demange, D., Fernández de Retana, Y., Pichardie, D.: Semantic reasoning about the sea of nodes. In: Proceedings of the 27th International Conference on Compiler Construction. p. 163–173. CC 2018, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3178372.3179503>
28. Devkota, S., Aschwanden, P., Kunen, A., LeGendre, M.P., Isaacs, K.E.: Ccnv: Understanding compiler optimizations in binary code. *IEEE Trans. Vis. Comput. Graph.* **27**(2), 667–677 (2021). <https://doi.org/10.1109/TVCG.2020.3030357>
29. Devkota, S., Isaacs, K.E.: Cfgexplorer: Designing a visual control flow analytics system around basic program analysis operations. *Comput. Graph. Forum* **37**(3), 453–464 (2018). <https://doi.org/10.1111/cgf.13433>
30. Dux, B., Iyer, A., Debray, S.K., Forrester, D., Kobourov, S.G.: Visualizing the behavior of dynamically modifiable code. In: 13th International Workshop on Program Comprehension (IWPC 2005), 15-16 May 2005, St. Louis, MO, USA. pp. 337–340. IEEE Computer Society (2005). <https://doi.org/10.1109/WPC.2005.45>
31. Gregg, D., Ertl, M.A., Krall, A.: Implementing an efficient java interpreter. In: Hertzberger, L.O., Hoekstra, A.G., Williams, R. (eds.) High-Performance Computing and Networking, 9th International Conference, HPCN Europe 2001, Amsterdam, The Netherlands, June 25-27, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2110, pp. 613–620. Springer (2001). https://doi.org/10.1007/3-540-48228-8_70
32. Hölzle, U., Chambers, C., Ungar, D.M.: Debugging optimized code with dynamic deoptimization. In: Feldman, S.I., Wexelblat, R.L. (eds.) Proceedings of the ACM SIGPLAN'92 Conference on Programming Language Design and Implementation (PLDI), San Francisco, California, USA, June 17-19, 1992. pp. 32–43. ACM (1992). <https://doi.org/10.1145/143095.143114>
33. Ishizaki, K., Kawahito, M., Yasue, T., Takeuchi, M., Ogasawara, T., Suganuma, T., Onodera, T., Komatsu, H., Nakatani, T.: Design, implementation, and evaluation of optimizations in a javatm just-in-time compiler. *Concurr. Pract. Exp.* **12**(6), 457–475 (2000). [https://doi.org/10.1002/1096-9128\(200005\)12:6<457::AID-CPE485>3.0.CO;2-0](https://doi.org/10.1002/1096-9128(200005)12:6<457::AID-CPE485>3.0.CO;2-0)

34. Jacobsen, B., Wallinger, M., Kobourov, S.G., Nöllenburg, M.: Metrosets: Visualizing sets as metro maps. *IEEE Trans. Vis. Comput. Graph.* **27**(2), 1257–1267 (2021). <https://doi.org/10.1109/TVCG.2020.3030475>
35. Lim, H., Debray, S.: Automated bug localization in JIT compilers. In: Titzer, B.L., Xu, H., Zhang, I. (eds.) *VEE '21: 17th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, Virtual USA*, April 16, 2021. pp. 153–164. ACM (2021). <https://doi.org/10.1145/3453933.3454021>
36. Luk, C., Cohn, R.S., Muth, R., Patil, H., Klauser, A., Lowney, P.G., Wallace, S., Reddi, V.J., Hazelwood, K.M.: Pin: building customized program analysis tools with dynamic instrumentation. In: Sarkar, V., Hall, M.W. (eds.) *Proceedings of the ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation, Chicago, IL, USA, June 12-15, 2005*. pp. 190–200. ACM (2005). <https://doi.org/10.1145/1065010.1065034>

Supplementary Materials

In this section, we provide the pseudocode for several of the non-trivial steps in our pipeline, as well several additional examples of buggy code visualized with our system.

1 Pseudocode

Algorithm 1: Intermediate Representation merging

Input: Original IR $R_0 = (V, E)$, Sub-IRs $r' = \{r'_1, \dots, r'_n\}$

Result: Merged IR R^*

```

1 function merge_nodes_to_original( $R^*$ ,  $r'_i$ ):
2   for  $j = 0$  to  $size(r'_i) - 1$  do
3     if  $v'_j.hyperedge$  exists in  $R^*.hyperedges$  then
4        $\lfloor$  add  $v'_j$  to  $R^*.hyperedges\_to\_nodes$ 
5     else
6        $\lfloor$  add  $v'_j.hyperedge$  to  $R^*.hyperedges$ 
7        $\lfloor$  add  $v'_j$  to  $R^*.hyperedges\_to\_nodes$ 
8   return  $R^*$ 

9 begin
10   $R^* = R_0$ 
11  for  $i = 0$  to  $size(r') - 1$  do
12     $\lfloor$   $R^* = merge\_nodes\_to\_original(R^*, r'_i)$ 

```

Algorithm 1 shows how to merge several sub-IRs (each of which contains candidate hyperedges) into the IR of the original JavaScript program. While traversing the sub-IRs, we call a *merge_nodes_to_original* function to update the copy (R^*) of original IR R_0 . The function updates *hyperedges*, which holds the list of hyperedge names, and *hyperedges_to_nodes*, which is a map between hyperedges and the list of nodes.

Algorithm 2: IR Simplification as Graph

Input: Merged IR $R^* = (V^*, E^*)$
Result: Simplified IR R^{**}

```

1 function Remove_Dead_Nodes( $R^*$ ):
2   for  $i = 0$  to  $\text{size}(V^*) - 1$  do
3     if  $v_i^*$  has no edge or has only self looping edge then
4       remove  $v_i^*$  from  $R^*$ 
5   return  $R^*$ 
6 function Remove_Node_Edges( $R^*$ ):
7    $R^{**} = \text{copy of } R^*$ 
8    $\text{removed\_nodes} = \emptyset$ 
9   for  $i = 0$  to  $\text{size}(V^*) - 1$  do
10    if  $v_i^*$  not in  $\text{removed\_nodes}$  then
11      for  $j = i + 1$  to  $\text{size}(V^*) - 1$  do
12        if  $v_j^*$  not in  $\text{removed\_nodes}$  and
13           $v_i^*.properties$  equal  $v_j^*.properties$  then
14            add  $v_j^*$  to  $\text{removed\_nodes}$ 
15            remove  $v_j^*$  from  $R^{**}$ 
16  for  $i = 0$  to  $\text{size}(V^{**}) - 1$  do
17    for  $j = 0$  to  $\text{size}(v^{**}.edges) - 1$  do
18      if  $v^{**}.edges_i$  in  $\text{removed\_nodes}$  then
19        remove  $v^{**}.edges_i$  from  $v^{**}.edges$ 
20  return  $R^{**}$ 
21 begin
22    $R^* = \text{Remove\_Dead\_Nodes}(R^*)$ 
23    $R^{**} = \text{Remove\_Node\_Edges}(R^*)$ 

```

Algorithm 2 shows how we simplify the merged IR as a graph, by removing nodes and edges. *Remove_Dead_Nodes* seeks nodes with no adjacent edges (or only with self-loop edges) and removes them from the IR. Then, *Remove_Node_Edge* merges pairs of nodes if they have the same opcode, the same optimization information, belong to the same IR (which can be identified by the IR id attribute), and share the same neighbors.

Algorithm 3: Hypergraph Construction

Input: Simplified graph $R^{**} = (V^{**}, E^{**})$
Result: Hypergraph $H = (V, S)$

```

1 function Construct_HyperGraph( $R^{**}$ ,  $R^{**}.hyperedges$ ):
2    $V = \emptyset$ 
3    $S = \{s_1, s_2, \dots, s_N\}$  // Where N is a total number hyperedges in
    $R^{**}.hyperedges$  and each  $s$  represents hyperedges.
4    $H = (V, S)$ 
5   for  $i = 0$  to  $size(V^{**}) - 1$  do
6     if  $v_i^*.hyperedge$  in  $R^{**}.hyperedges$  then
7        $v = v_i^*$  //  $v$  stands for element.
8       identify the appropriate set  $s$  in  $S$  and add  $v$  to  $s$ 
9       add  $v$  to  $V$ 
10      while  $v_i^*.optimized\_hyperedges$  do
11         $w = v_i^*$ 
12        identify the appropriate set  $s$  in  $S$  and add  $v$  to  $s$ 
13        add  $v$  to  $V$ 
14    return  $H$ 
15 begin
16    $H = Construct\_HyperGraph(R^{**}, G^{**}.phases)$ 

```

Algorithm 3 shows how we construct a hypergraph $H = (V, S)$ from the simplified graph $G = (V, E)$. First, we extract hyperedges from the simplified graph, which group nodes based on their generation, and add the nodes to V . Then we add nodes to all their corresponding hyperedges.

Algorithm 4: Reducing the Hypergraph

Input: Hypergraph $H = (V, S)$ **Result:** Hyperedge reduced IR $H^* = (V^*, S^*)$

```

1 begin
2    $V^* = V$ 
3    $S^* = \emptyset$ 
4    $H^* = (V^*, S^*)$ 
5    $merged\_hyperedges = \emptyset$ 
6   for  $i = 0$  to  $size(S) - 1$  do
7     if  $s_i$  is not in  $merged\_hyperedges$  then
8        $concatenated\_id = s_i.id$ 
9        $merged\_node\_list[s_i.id] = s_i.nodes$ 
10      add  $s_i.name$  to  $merged\_hyperedges$ 
11      for  $j = i + 1$  to  $size(S) - 1$  do
12        if  $s_i.name == s_j.name$  then
13           $j_i d = s_j.id$ 
14          add  $j_i d$  with @ symbol to  $concatenated\_id$ 
15          add  $s_j.nodes$  to  $merged\_node\_list[s_i.id]$ 
16      add  $merged\_node\_list[s_i.id]$  to  $S^*[concatenated\_id]$ 

```

Algorithm 4 shows the steps for reducing the hypergraph $H = (V, S)$. This is accomplished by merging hyperedges with the same name. Specifically, when merging hyperedges, we generate a new ID by concatenating the IDs delimited with a special character '@' for the merged hyperedge. The corresponding nodes in the hyperedges also get merged.

Algorithm 5: Hypergraph Simplification

Input: Hypergraph $H^* = (V^*, S^*)$
Result: Simplified hypergraph $H^{**} = (V^{**}, S^{**})$

```

1 function Simplify_Hyperedge( $s^*$ ):
2    $s^{**} = \text{copy of } s^*$ 
3    $\text{simplified\_nodes} = \emptyset$ 
4   for  $i = 0$  to  $\text{size}(s^*.nodes) - 1$  do
5     if  $\text{node}_i$  not in  $\text{simplified\_nodes}$  then
6       for  $j = i + 1$  to  $\text{size}(s^*.nodes) - 1$  do
7         if  $\text{node}_j$  not in  $\text{simplified\_nodes}$  then
8           if  $\text{node}_i.\text{properties}$  equal  $\text{node}_j.\text{properties}$  and
            $\text{node}_i.\text{hyperedges}$  equal  $\text{node}_j.\text{hyperedges}$ 
           then
9             add  $\text{node}_j$  to  $\text{simplified\_nodes}$ 
10            remove  $\text{node}_j$  from  $s^{**}$ 
11  return  $s^{**}$ 

12 begin
13    $V^{**} = \emptyset$ 
14    $S^{**} = \emptyset$ 
15    $H^{**} = (V^{**}, S^{**})$ 
16   for  $i = 0$  to  $\text{size}(S^*) - 1$  do
17      $s^{**} = \text{Simplify\_Hyperedge}(s_i^*)$ 
18     add all elements in  $s^{**}$  to  $V^{**}$ 
19     add  $s^{**}$  to  $S^{**}$ 

```

Algorithm 5 shows the steps for node reduction in the hypergraph. While traversing all nodes in the hyperedge-simplified H^* , we compare each node to the other nodes that belong to the same hyperedges. We compare the properties (i.e., opcode) of the nodes, and we merge any nodes that are found to be the same.

2 More Results and Examples

The bug report 1072171 (Figure 5) is an example with a bug in the Typer hyperedge, where JavaScript code "Math.max" and "Math.min" generates a wrong type due to incorrectly removing one of the typing properties during optimization [9, 35].

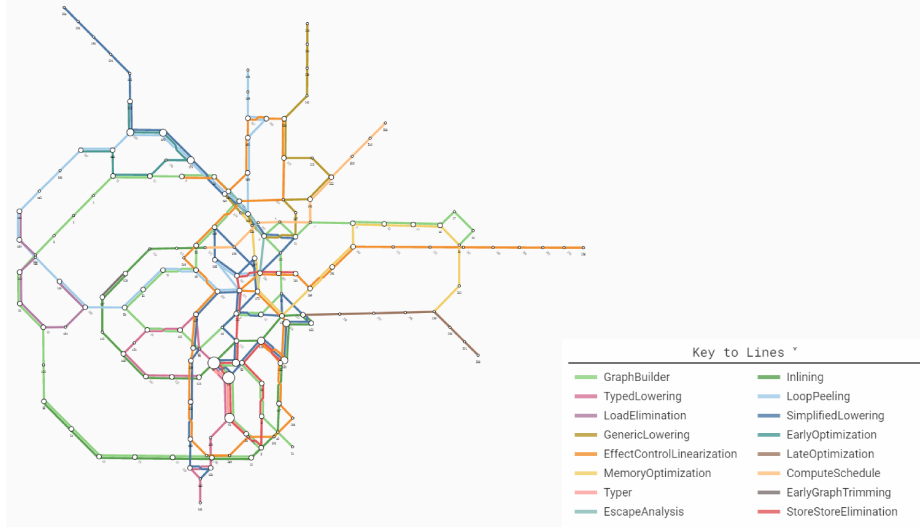


Fig. 5. Bug Report 1072171

When we compare the buggy and non-buggy IRs, we can see that the buggy IR is missing a node generation at the Typer hyperedge. In other words, when we compare the node generation hyperedges, the Typer hyperedge only exists in the IRs from non-buggy programs, but not in the buggy program IR. Therefore, we find the node that was added to the line because of the generation status (and not the optimization status). Specifically, there is only one node, among the four in the Typer line, that was generated by Typer. The other nodes were generated in different phases/lines. Checking the IR ID attribute for this particular node shows that it is from another IR that is non-buggy and all the other nodes on the line are generated from the buggy programs.

The bug report 961237 (Figure 6) is an example with a bug in the SimplifiedLowering hyperedge where *null* is truncated to *+0* even in contexts such as $-0 == null$, due to missing type checking [13, 35].

If a type checking is missing during the evaluation for the optimization, it skips the optimization. Thus, we see a line in which the majority of the nodes belong to other IRs. In this case, the SimplifiedLowering line has 16 nodes, but the majority of the nodes (13 nodes) are from other IRs. This indicates that there was a significant difference between the IR from the original program, which is

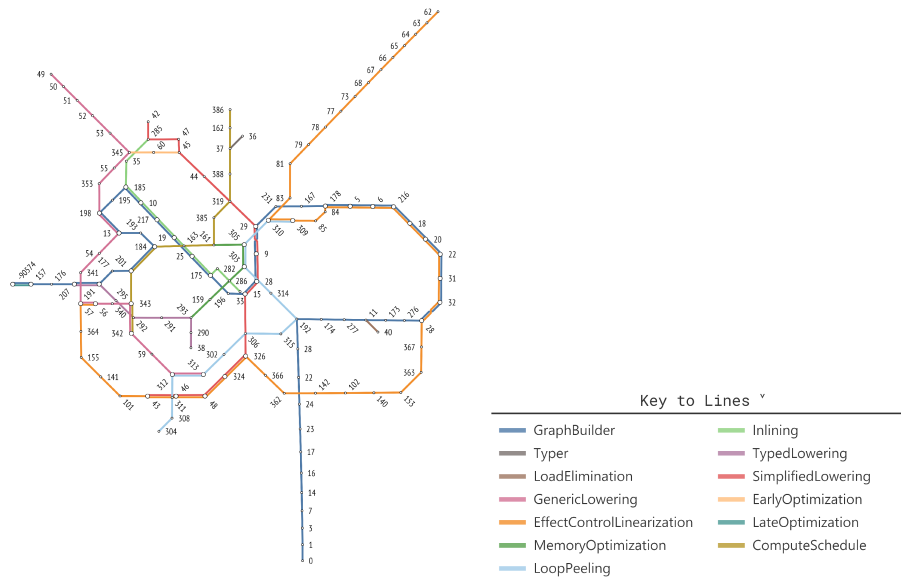


Fig. 6. Bug Report 961237

buggy, and the other IRs from non-buggy programs at SimplifiedLowering hyperedge.

The bug report 880207 (Figure 7) details an issue similar to the previous one: there is a bug in the SimplifiedLowering hyperedge where it incorrectly checks the type of "Math.expm1" [12].

Again, we can find a line in which the majority of nodes come from the other IRs. That means that the buggy program is missing the appropriate optimization, resulting in missing nodes in the IR compared to the original IR.

3 MetroMap Features

Here we briefly illustrate how different interaction with MetroSets are used to answer our list of desired questions.

Figure 8 shows an example of line hovering: this results in graying out all other lines. The line name (GraphBuilder, in this example) in the "Key to Lines" is also bolded.

Figure 9 shows an example of hovering over the most active hyperedge. The example from bug report 5129 and shows that the SimplifiedLowering operation was the most active one.

Figure 10 shown an example of the intersection exploration mode (on the left side panel). This is particularly useful when we want to find nodes that are optimized in multiple phases. The example shows that there are 3 nodes

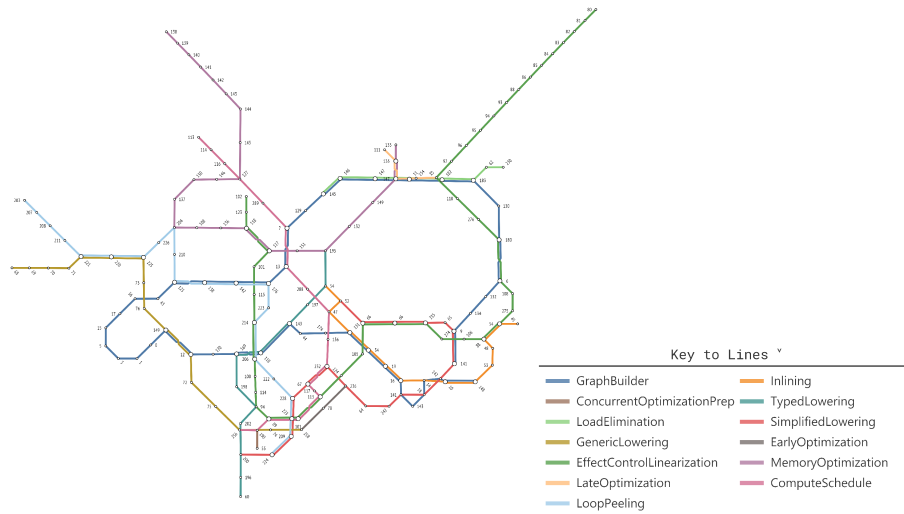


Fig. 7. Bug Report 880207

intersecting the selected lines: TypedLowering, SimplifiedLowering, and EarlyOptimization.

Figure 11 shows an example of node hovering, which provides the attribute information associated with a node of interest. In this example, the node was generated for the arithmetic operator subtract, which can be found by opcode `0x007c`, was generated at the TypedLowering phase and optimized at SimplifiedLowering and EarlyOptimization phases. Moreover, we can tell that this node generation and optimizations happened in late phases of the JIT compilation, identified by the phase IDs, which represent the execution order of phases.



Fig. 8. Hovering on a line.

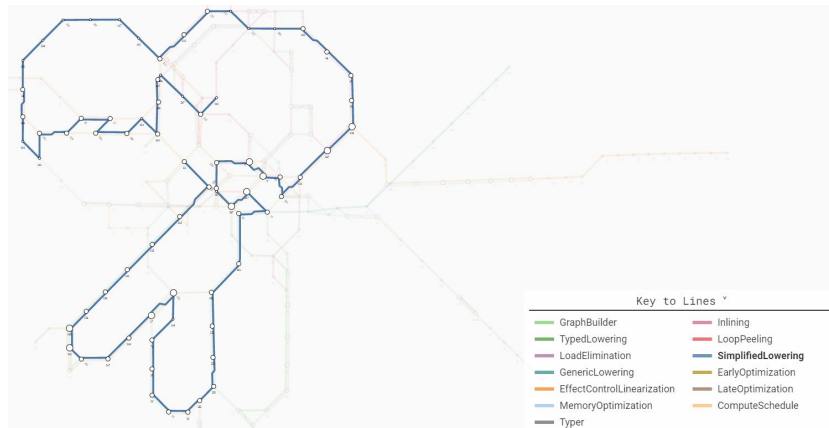


Fig. 9. Most active hyperedge.

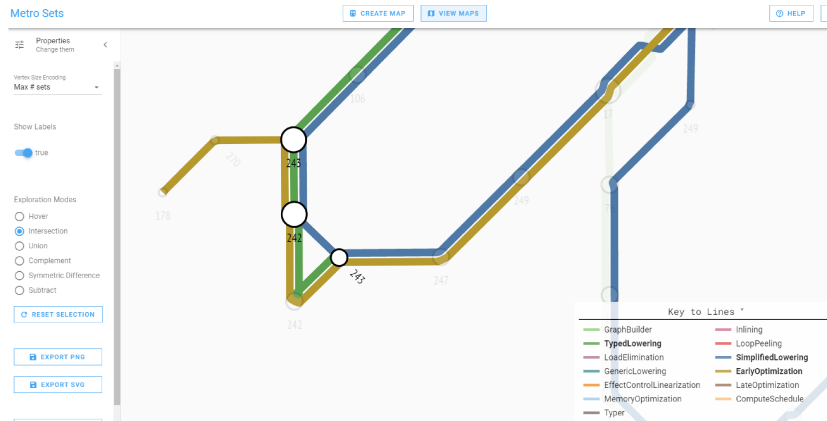


Fig. 10. Intersection.

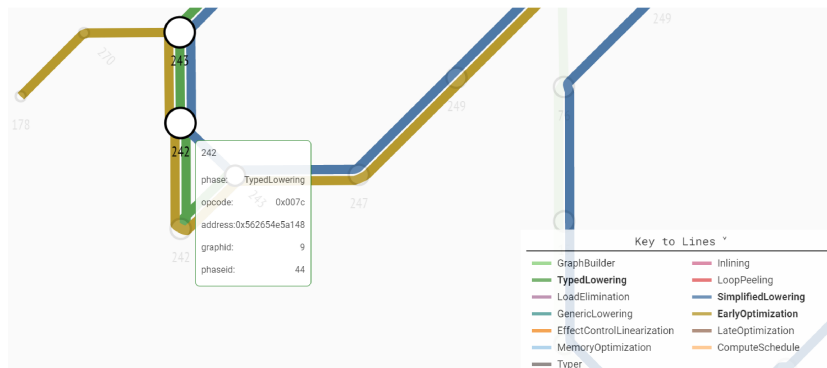


Fig. 11. Node hovering