Demanded Abstract Interpretation

Benno Stein
University of Colorado Boulder
USA
benno.stein@colorado.edu

Bor-Yuh Evan Chang*
University of Colorado Boulder
Amazon
USA
evan.chang@colorado.edu

Manu Sridharan
University of California, Riverside
USA
manu@cs.ucr.edu

Abstract

We consider the problem of making expressive static analyzers *interactive*. Formal static analysis is seeing increasingly widespread adoption as a tool for verification and bug-finding, but even with powerful cloud infrastructure it can take minutes or hours to get *batch* analysis results after a code change. While existing techniques offer some demand-driven or incremental aspects for certain classes of analysis, the fundamental challenge we tackle is doing *both* for arbitrary abstract interpreters.

Our technique, *demanded abstract interpretation*, lifts program syntax and analysis state to a dynamically evolving graph structure, in which program edits, client-issued queries, and evaluation of abstract semantics are all treated uniformly. The key difficulty addressed by our approach is the application of general incremental computation techniques to the complex, cyclic dependency structure induced by abstract interpretation of loops with widening operators. We prove that desirable abstract interpretation meta-properties, including soundness and termination, are preserved in our approach, and that demanded analysis results are equal to those computed by a batch abstract interpretation. Experimental results suggest promise for a prototype demanded abstract interpretation framework: by combining incremental and demand-driven techniques, our framework consistently delivers analysis results at interactive speeds, answering 95% of queries within 1.2 seconds.

CCS Concepts: • Theory of computation → Program analysis; • Software and its engineering → Formal software verification.

*Bor-Yuh Evan Chang holds concurrent appointments at the University of Colorado Boulder and as an Amazon Scholar. This paper describes work performed at CU Boulder and is not associated with Amazon.

This work is licensed under a Creative Commons Attribution International 4.0 License.

PLDI '21, June 20–25, 2021, Virtual, Canada
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8391-2/21/06.
<https://doi.org/10.1145/3453483.3454044>

Keywords: Abstract interpretation, Incremental computation, Demand-driven query evaluation, Demanded fixed-points

ACM Reference Format:

Benno Stein, Bor-Yuh Evan Chang, and Manu Sridharan. 2021. Demanded Abstract Interpretation. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*, June 20–25, 2021, Virtual, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3453483.3454044>

1 Introduction

Static analysis is seeing increasing real-world adoption for verification and bug finding, particularly as part of continuous integration (CI) and code review processes [9, 36]. However, a pain point with these deployments is that developers cannot get quick local analysis results for code they are editing; ideally, updated results would appear nearly instantly in their IDE. In this paper, we present an *interactive* analysis engine designed to handle local queries and edits efficiently, which can complement a *batch* engine that exhaustively analyzes a fixed program, for example by quickly verifying whether a local change silences an alarm raised in CI.

The well-known techniques of demand-driven analysis and incremental analysis help address this challenge. Demand-driven analyses compute only those results needed to answer a set of extrinsically-provided *queries*, while incremental analyses speed up re-analysis of an *edited* program by re-using as many previously-computed results as possible.

Powerful frameworks for incremental (e.g., [5, 48]) or demand-driven (e.g., [24]) static analysis do exist, but nearly all such frameworks target restricted analysis domains (e.g., finite or finite-height domains), whereas well-known analyses like octagon and shape analysis require an infinite-height abstract domain. There are also approaches to adapt summary-based analyses to offer coarse-grained method- or file-level incrementality (e.g., [9, 14, 18]). Though these approaches effectively scale to industrial codebases in CI pipelines, they are not intended to achieve real-time interactivity during the development process.

In contrast, our aim is to support fine-grained incremental and demand-driven analysis over *arbitrary* abstract domains expressed in general-purpose languages, thus enabling the reuse of existing optimized abstract domain implementations at interactive speeds. To our best knowledge, no general

technique exists to automatically compute a demand-driven or incremental version of an arbitrary abstract interpretation with arbitrary widening operators, a key requirement to ensure termination in more complex analyses.

This paper presents *demanded abstract interpretation*, an abstract interpretation framework with first-class support for *both* demand-driven and incremental analysis. We build on abstract interpretation [11], which provides a methodology for expressing static analyses and guaranteeing their correctness, and take inspiration from work on general incremental computation using “demanded computation graphs” [22].

Our framework reifies the abstract interpretation (AI) of a program as a dynamically evolving *demanded abstract interpretation graph* (DAIG), which explicitly represents program statements, abstract states, and the dependency structure of analysis computations. In this representation, program edits, client-issued queries, and the evaluation of abstract semantics can all be treated uniformly. Cyclic control flow is a key difficulty for this approach, since cyclic dependencies lead to unclear evaluation semantics. We define an operational semantics for DAIGs that preserves an acyclic invariant while modifying and extending the graph on demand, thus soundly analyzing loops with guaranteed termination, assuming termination of the underlying abstract interpretation.

The DAIG encoding and evaluation enables efficient abstract interpretation in an *interactive* mode, analyzing a minimal number of statements to respond to queries with maximal reuse of previously-computed results. In particular, this paper makes the following key contributions:

- We introduce a framework for *demanded abstract interpretation* in which program syntax and analysis computation structure are reified into *demanded abstract interpretation graphs* (DAIGs) (Section 4).
- We specify an operational semantics for DAIGs that realizes incremental updates and demand-driven evaluation via demanded unrolling of abstract interpretation fixed-point computations (Section 5).
- We prove that demanded abstract interpretation preserves soundness and termination, and that its results are from-scratch consistent with classical abstract interpretation by global fixed-point iteration (Section 6).
- We provide evidence for the expressivity and efficacy of demanded abstract interpretation using a prototype framework instantiated with interval, octagon, and shape domains (Section 7). In our experiments, DAIGs support context-sensitive interprocedural analysis at interactive speeds, answering 95% of queries within 1.2 seconds.

2 Overview

Fig. 1 shows a simple imperative program that appends two linked lists. Given well-formed (i.e., null-terminated and acyclic) input lists p and q , `append` must return a well-formed list and not dereference `null` in order to be correct. These

```

function append(p: List, q: List): List {
 $\ell_0$    if (p == null) {
 $\ell_1$        return q; }
 $\ell_2$    var r: List = p;
 $\ell_3$    while (r.next != null) {
 $\ell_4$        r = r.next; }
 $\ell_5$    r.next = q;
 $\ell_6$    return p;
 $\ell_{ret}$  }
```

Figure 1. A procedure to append two linked lists. The labels ℓ_i mark program locations in its control flow.

properties can be verified using a separation logic, abstract interpretation-based shape analysis [7, 15, 27], tracking facts like $\text{lseg}(p, \text{null})$ to represent well-formedness of list p . Our goal is to enable interactive performance for arbitrary abstract interpretations, including such analyses, in response to a user’s edits and queries.

In this section, we illustrate our approach to *demanded abstract interpretation* by example. We demonstrate how abstract interpretation of forwards control flow is reified in a DAIG (Section 2.1), and then show how the DAIG supports both demand-driven and incremental interactions with the underlying abstract interpretation (Section 2.2). Finally, we highlight key difficulties introduced by cyclic control flow, show how analysis thereof can be encoded acyclically, and demonstrate how our operational semantics evolve the DAIG on-demand to soundly compute fixed points (Section 2.3).

2.1 Reifying Abstract Interpretation in DAIGs

The `append` procedure from Fig. 1 may equivalently be represented as a control-flow graph (CFG) as shown in Fig. 2, with vertices for program locations and edges labelled by atomic program statements.¹ A classical abstract interpreter analyzes such a program by starting with some initial abstract state at the entry location and applying an abstract transfer function $\llbracket \cdot \rrbracket^\#$ to interpret statements, a join operator $\sqcup$ at nodes with multiple predecessors, and a widen operator ∇ at cycles as needed until a fixed point is reached.

The demanded abstract interpretation graph (DAIG) shown in Fig. 3 reifies the computational structure of such an abstract interpretation of the Fig. 2 CFG. Its vertices are uniquely-named mutable reference cells containing program syntax or abstract state, and its edges fully specify the computations of an abstract interpretation. Names identify values for reuse across edits and queries, and hence must *uniquely* identify the inputs and intermediate results of the abstract interpretation. In Fig. 3 and throughout this paper, *underlined* symbols denote a name derived from that symbol: hashes, essentially.

¹As is standard, we have simply broken down the guard conditions from `if` and `while` into `assume` guard statements for each side of a branch.

Figure 2. The control-flow graph (CFG) of the append procedure from Fig. 1.

Figure 3. A demanded abstract interpretation graph (DAIG) for the program given in Fig. 1 before any queries are issued. The elided loop encoding is shown in Fig. 4c.

To encode abstract interpretation computations, DAIG edges are labelled by a symbol for an abstract interpretation function and connect cells storing the function inputs to the cell storing the output, capturing the dependency structure of the analysis computation.² For example, the computation of the abstract transfer function over the CFG edge $\ell_0 \rightarrow \ell_1$ is encoded in Fig. 3 as a DAIG edge with input cells ℓ_0 and $\ell_0 \cdot \ell_1$ (respectively containing the fixed-point state at ℓ_0

²More precisely, DAIGs have *hyper*-edges, since they connect multiple sources (function inputs) to one destination (function output).

and the corresponding statement $s_0: \text{assume}(p == \text{null})$, labelled by the abstract transfer function symbol $\llbracket \cdot \rrbracket^\#$.

2.2 Demand-Driven and Incremental Analysis

Next, we demonstrate how a DAIG encoding naturally supports demand-driven and incremental analysis. We use the aforementioned shape-analysis domain for our example, a separation logic-based domain with a “list segment” primitive $\text{lseg}(x, y)$ that abstracts the heaplet containing a list segment from x to y .³ This domain is of infinite height, absent a best abstraction function, and with complex widening operators, and therefore incompatible with previous frameworks that restrict the domain form.

In Fig. 4a, we show the result of evaluating a demand query on our example DAIG. Suppose a client issues a query for the $\underline{1} \cdot \ell_{\text{ret}}$ cell in Fig. 3, the abstract state corresponding to the **return** q statement at ℓ_1 in Fig. 1. Since the $\underline{1} \cdot \ell_{\text{ret}}$ cell has predecessors $\underline{\ell_1} \cdot \ell_{\text{ret}}$ and $\underline{\ell_1}$, we issue requests for the values of those cells. Cell $\underline{\ell_1}$ is empty, but depends on $\underline{\ell_0} \cdot \ell_1$ and $\underline{\ell_0}$, so more requests are issued. Both of those cells hold values, so we can compute and store the value of $\underline{\ell_1}$. Now, having satisfied its dependencies, we can compute the value of $\underline{1} \cdot \ell_{\text{ret}}$, as shown in Fig. 4a. Note that DAIGs are always acyclic, so this recursive traversal of dependencies is well-founded.

Crucially, these results are now *memoized* for future incremental reuse; a subsequent query for ℓ_{ret} , for example, will memo match on $\underline{1} \cdot \ell_{\text{ret}}$ and only need to compute $\underline{2} \cdot \ell_{\text{ret}}$ and its dependencies from scratch. This fine-grained reuse of intermediate abstract interpretation results is a key feature of the DAIG encoding for demand-driven analysis.

To handle developer edits to code, DAIGs are also naturally *incremental*, efficiently recomputing and reusing analysis results across multiple program versions, following the incremental computation with names approach [21].

Consider a program edit which adds a logging statement `print("p is null")` just before the **return** at ℓ_1 in Fig. 1. Intuitively, program behaviors are unchanged at those locations unreachable from the added statement, so an incremental analysis should only need to re-analyze the sub-DAIG reachable from the new statement.

Fig. 4b illustrates this program edit’s effect on the DAIG. The **green** nodes correspond to the added statement cell $\underline{\ell_1} \cdot \ell_7$ and its corresponding abstract state cell ℓ_7 . Nodes forward-reachable from the green nodes — those marked in **red** — are invalidated (a.k.a. “dirtied”) by our incremental computation engine. In particular, cells $\underline{\ell_1} \cdot \ell_{\text{ret}}$ and ℓ_{ret} containing abstract states ϕ' and ϕ'' , respectively, are dirtied.

Crucially, while nodes are dirtied *eagerly*, they are recomputed from up-to-date inputs *lazily*, only when demanded. That is, the DAIG encoding allows our analysis to avoid constant recomputation of an analysis as a program is edited, instead computing results on demand while soundly keeping

³That is, a sequence of iterated next pointer dereferences from x to y .

Figure 4. Demanded abstract interpretation: (a) Demanding a value for $1 \cdot \underline{\ell}_{\text{ret}}$ recursively triggers demand for its dependencies and is resolved by computing its value from the statements in $\ell_0 \cdot \underline{\ell}_1$ and $\underline{\ell}_1 \cdot \underline{\ell}_{\text{ret}}$ and the initial state φ_0 in ℓ_0 . We show only the relevant subgraph here, but this operation occurs in the full DAIG of Fig. 3. (b) DAIG from Fig. 3 updated to reflect nodes added ($\underline{\ell}_1 \cdot \underline{\ell}_7$ and $\underline{\ell}_7$) and potentially affected ($1 \cdot \underline{\ell}_{\text{ret}}$ and $\underline{\ell}_{\text{ret}}$) by the edit in Section 2.2. All other nodes are unchanged. (c) DAIG for the ℓ_3 -to- ℓ_4 -to- ℓ_3 loop of Fig. 2 after one demanded unrolling, with the new DAIG region shown in red and the (removed) pre-unrolling fix edge shown in dotted grey. Note that cells containing program syntax are not duplicated. The DAIG with the black vertices and edges along with the grey edge (but not the red ones) is the initial sub-DAIG for the dashed ellipse in Fig. 3.

track of which intermediate results — possibly from a previous program version — are available for reuse. For example, assuming that $\underline{\ell}_1$ and $2 \cdot \underline{\ell}_{\text{ret}}$ were both computed before the edit, a query for $\underline{\ell}_{\text{ret}}$ must execute only two transfers and one join: the red and green edges of Fig. 4b. This represents a significant savings over recomputing the entire analysis — including the loop fixed point — as would be necessary without incremental analysis.

If desired, an auxiliary memoization (memo) table can also be used to cache computations independent of program locations to enable further incrementalization, with names based on the input values (e.g., memoizing $\llbracket s_0 \rrbracket^\#(\varphi_0)$ in a cell named $\llbracket \cdot \rrbracket^\# \cdot s_0 \cdot \varphi_0$). As with batch analysis, it is sound to drop cached results from the DAIG and/or memo table and later recompute those results if needed, trading efficiency of reuse for a lower memory footprint.

2.3 Cyclic Control Flow and Demanded Fixed Points

As shown in Section 2.1, encoding program structure and analysis data-flow into a dependency graph is relatively straightforward when the control-flow graph is acyclic. However, when handling loops or recursion, like lines ℓ_3 and ℓ_4 in Fig. 1, an abstract interpreter’s fixed-point computation is inherently cyclic. Properly handling these cyclic control-flow and data-flow dependency structures is the crux of realizing demanded abstract interpretation. For instance, introducing cyclic dependencies in the DAIG yields an unclear evaluation semantics. The key insight we leverage is that we can instead enrich the demand-driven query evaluation and the incremental edit semantics to dynamically evolve DAIGs such that each step preserves the acyclic dependency structure

invariant. To do so, we use a distinguished edge label (fix) to indicate a dependency on the fixed-point of a given region of the DAIG, which is then dynamically unrolled on-demand by query evaluation and rolled by incremental edits.

The details are formalized in Section 5; we proceed here by example on the demanded abstract interpretation of the append program of Fig. 2. Consider Fig. 4c, and focus on the black and grey vertices and edges, ignoring the red for the moment. Rather than encoding the CFG back edge from ℓ_4 to ℓ_3 directly into the DAIG (violating acyclicity in the process), the ℓ_3 -to- ℓ_4 -to- ℓ_3 loop is initially encoded with separate reference cells for the 0th and 1st abstract iterates at the loop head ℓ_3 (named $\underline{\ell}_3^{(0)}$ and $\underline{\ell}_3^{(1)}$ respectively), along with a fix edge from those two cells to ℓ_3 ’s fixed-point cell $\underline{\ell}_3$, as seen in the grey dotted edge. Crucially, this initial DAIG is acyclic.

Query Evaluation. For query evaluation, we can compute fixed points on demand by “unrolling” the *abstract interpretation* of loop bodies in the DAIG one *abstract iteration* at a time until a fixed point is reached, preserving the acyclic DAIG invariant at each step. That is, our key observation is to unroll at the semantic level of the abstract interpretation rather than the syntactic level of the control-flow graph.

From the initial DAIG in Fig. 4c, when the fixed-point $\underline{\ell}_3$ is demanded, its dependencies — the 0th and 1st abstract iterates — are computed. If their values are equal, then a fixed-point has been reached and may be written to $\underline{\ell}_3$.⁴ If their values are *not* equal, then the abstract interpretation

⁴We describe here the widening strategy of applying ∇ every iteration until a fixed-point is reached for simplicity, but the same general idea applies for other widening strategies or checking convergence with $\sqsubseteq$ instead of $=$.

of the loop body — but not the loop body statements — is unrolled one step further and the `fix` edge slides forward to now depend on the 1st and 2nd abstract iterates, as seen in the red cells and edges of Fig. 4c. And crucially, this one-step unrolled DAIG is also acyclic.

From here, the process continues, and termination is guaranteed by leveraging the standard argument of abstract interpretation meta-theory: the sequence of abstract iterates in the cells $\underline{\ell}_3^{(0)}, \underline{\ell}_3^{(1)}, \underline{\ell}_3^{(2)}, \dots$ converges because it is produced by widening a monotonically increasing sequence of abstract states, so this demanded unrolling of `fix` occurs only finitely — but unboundedly — many times. We see that in essence, the sequence of abstract interpretation iterates $\underline{\ell}_3^{(0)}, \underline{\ell}_3^{(1)}, \underline{\ell}_3^{(2)}, \dots$ are encoded into the DAIG on demand during query evaluation.

In classical abstract interpretation, a widen ∇ is a join that enforces convergence during interpretation and thus is only strictly needed if the abstract domain has infinite height. Our approach can be seen as an application of this widening principle to demanded computation. For an abstract domain of finite height k , it would have been sufficient to encode the unrolling of `fix` eagerly into an acyclic DAIG by inlining the abstract iteration k times to k iterate cells $\underline{\ell}_3^{(0)}, \dots, \underline{\ell}_3^{(k)}$. However, many expressive, real-world abstract domains — including the shape analysis domain of our example and most numerical domains — are of infinite height.

Incremental Edits. Since the acyclic DAIG invariant is always preserved, invalidating on incremental edits still only requires eager dirtying forwards in the DAIG, with some special semantics for `fix` edges. When dirtying along a `fix` edge, the `fix` edge is rolled back to a non-dirty cell (i.e., the 0th and 1st iterate). In Fig. 4c, if the statement cell $\ell_4 \cdot \underline{\ell}_3$ is edited, then dirtying will happen along the red solid `fix` edge at which point it will slide back to be the grey dotted one.

Interprocedural Demand. This demanded unrolling of `fix` also suggests an approach to interprocedural demanded analysis parameterized by a context-sensitivity policy. To analyze a call when evaluating a query, we construct a DAIG for the callee procedure on demand, indexed by a context determined opaquely by the context-sensitivity policy.

The “functional approach” to interprocedural analysis of Sharir and Pnueli [39] could also potentially be adapted to our framework by constructing disjoint DAIGs for each phase and inserting dependencies from phase-2 callsites to corresponding phase-1 summaries.

These techniques both rely on a static call graph, which can be computed soundly using either abstract interpretation (which may itself be expressed in a DAIG) or type-/constraint-based approaches.

statements	$s \in Stmt$	
locations	$\ell \in Loc$	
control-flow edges	$e \in Edge$	$::= \ell \cdot [s] \rightarrow \ell'$
programs	$\langle L, E, \ell_0 \rangle$	$: \mathcal{P}(Loc) \times \mathcal{P}(Edge) \times Loc$
concrete states	$\sigma \in \Sigma$	(with initial state σ_0)
concrete semantics	$\llbracket \cdot \rrbracket$	$: Stmt \rightarrow \Sigma \rightarrow \Sigma_{\perp}$
collecting semantics	$\llbracket \cdot \rrbracket_{\langle L, E, \ell_0 \rangle}^*$	$: Loc \rightarrow \mathcal{P}(\Sigma)$

Figure 5. A generic programming language of control-flow graphs edge-labelled by an unspecified statement language.

3 Preliminary Definitions

Our technique lifts a program and an abstract interpreter *together* into a demanded abstract interpretation graph (DAIG), a representation that is amenable for sound incremental and demand-driven program analysis. By design, this construction is *generic* in the underlying programming language and concrete semantics as well as the abstract domain and abstract semantics. In this section, we fix a generic programming language and an abstract interpreter interface that serve as inputs to DAIG construction, both to define syntax and to make explicit our assumptions about their semantic properties. Selected instantiations of the framework for real-world analysis problems are given in Section 7.

Programs under analysis are given as control-flow graphs, edge-labelled by an unspecified statement language and interpreted by a denotational concrete semantics as shown in Fig. 5. A program $\langle L, E, \ell_0 \rangle$ is a 3-tuple composed of a set L of control locations, a set E of directed, statement-labelled control-flow edges between locations, and an initial location ℓ_0 . We say that a program $\langle L, E, \ell_0 \rangle$ is *well-formed* when (1) ℓ_0 and all locations in E are drawn from L , and (2) L and E form a reducible control-flow graph. These conditions ensure that we avoid degenerate edge cases and only consider control flow graphs which correspond to realistic programs [4].

Statements are interpreted by the concrete denotational semantics $\llbracket \cdot \rrbracket$ as partial functions over concrete program states. As is standard, we can also lift this statement semantics to a collecting $\llbracket \cdot \rrbracket_{\langle L, E, \ell_0 \rangle}^*$ of full programs, by computing the transitive closure of the statement semantics over a flow graph. That is, $\llbracket \ell \rrbracket_{\langle L, E, \ell_0 \rangle}^*$ is the set of all concrete states that can be witnessed at program location ℓ in a valid program execution. We elide the subscript when it is clear from context. Such a collecting semantics is uncomputable in general, but is an important tool for reasoning about analysis soundness.

Now, we define the interface of a generic abstract interpreter over this control-flow graph language. These definitions are intended simply to fix notation and minimize ambiguity and are as standard as possible.

An abstract interpreter is a 6-tuple $\langle \Sigma^{\#}, \varphi_0, \llbracket \cdot \rrbracket^{\#}, \sqsubseteq, \sqcup, \nabla \rangle$ composed of:

- An *abstract domain* $\Sigma^\#$ (elements of which are referred to as *abstract states*) which forms a semi-lattice under
 - a *partial order* $\sqsubseteq \in \mathcal{P}(\Sigma^\# \times \Sigma^\#)$ with a bottom $\perp \in \Sigma^\#$
 - an *upper bound* (a.k.a. *join*) $\sqcup : \Sigma^\# \rightarrow \Sigma^\# \rightarrow \Sigma^\#$
- An *initial abstract state* $\varphi_0 \in \Sigma^\#$
- An *abstract semantics* $\llbracket \cdot \rrbracket^\# : Stmt \rightarrow \Sigma^\# \rightarrow \Sigma^\#$ that interprets program statements as monotone functions over abstract states.
- A *widening operator* $\nabla : \Sigma^\# \rightarrow \Sigma^\# \rightarrow \Sigma^\#$ that is an upper bound operator (i.e., $(\varphi \sqcup \varphi') \sqsubseteq (\varphi \nabla \varphi')$ for all φ, φ') and enforces convergence (i.e., for all increasing sequences of abstract states $\varphi_0 \sqsubseteq \varphi_1 \sqsubseteq \varphi_2 \sqsubseteq \dots$, the sequence $\varphi_0, \varphi_0 \nabla \varphi_1, (\varphi_0 \nabla \varphi_1) \nabla \varphi_2, \dots$ converges).

Furthermore, a concretization function $\gamma : \Sigma^\# \rightarrow \mathcal{P}(\Sigma)$ gives meaning to abstract states. We say that a concrete state σ *models* an abstract state φ (equivalently, that φ *abstracts* σ), written $\sigma \models \varphi$, when $\sigma \in \gamma(\varphi)$.

Definition 3.1 (Local Abstract Interpreter Soundness). An abstract interpreter $\langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$ is locally sound if for all σ, φ, s , if $\sigma \models \varphi$ and $\llbracket s \rrbracket \sigma \neq \perp$ then $\llbracket s \rrbracket \sigma \models \llbracket s \rrbracket^\# \varphi$.

Local soundness can be extended to a global soundness property: if the abstract semantics are locally sound, then the abstract interpreter computes a sound over-approximation of the possible concrete states at each location.

Proposition 3.2 (Global Abstract Interpreter Soundness). If $\langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$ is locally sound and $\sigma_0 \models \varphi_0$ then it induces an abstract collecting semantics $\llbracket \cdot \rrbracket^\#_{\langle L, E, \ell_0 \rangle} : Loc \rightarrow \Sigma^\#$ such that for all $\sigma \in \llbracket \ell \rrbracket^\#_{\langle L, E, \ell_0 \rangle}, \sigma \models \llbracket \ell \rrbracket^\#_{\langle L, E, \ell_0 \rangle}$.

We elide the abstract collecting semantics $\llbracket \cdot \rrbracket^\#_{\langle L, E, \ell_0 \rangle}$; it is similarly a transitive closure of the abstract statement semantics over a flow graph. It is a well-known result that global abstract interpreter soundness is implied by local soundness [11] and that such a global fixed-point is computable using the chaotic iteration method with widening [8].

4 Demanded AI Graphs

Recall from Section 2 that a demanded abstract interpretation graph (DAIG) is a directed acyclic hypergraph, whose vertices are reference cells containing program syntax or intermediate analysis results, and whose edges reflect analysis dataflow relationships among those cells. In Fig. 6, we show a syntax for DAIGs. A DAIG $\mathcal{D} = \langle R, C \rangle$ is composed of a set $R \subseteq Ref$ of named reference cells connected by computation edges $C \subseteq Comp$. A computation $c : n \leftarrow f(n_1, \dots, n_k)$ is an edge connecting sources $\{n_1, \dots, n_k\}$ to a singleton destination $\{n\}$, labeled by some analysis function f .

Names $\underline{\ell}$, $\underline{f}$, $\underline{v}$ and $\underline{i}$ correspond respectively to locations ℓ , functions f , values v and integers i , supporting memoization of those syntactic constructs. Name products $n_1 \cdot n_2$ support the construction of more complicated names, and i -primed

functions	f	$::= \llbracket \cdot \rrbracket^\# \mid \sqcup \mid \nabla \mid \text{fix}$
values	v	$::= s \mid \varphi$
names	$n \in Nm$	$::= \underline{\ell} \mid \underline{f} \mid \underline{i} \mid \underline{v} \mid n_1 \cdot n_2 \mid n^{(i)}$
types	τ	$\in \{Stmt, \Sigma^\#\}$
reference cells	$r \in Ref$	$::= n[v : \tau] \mid n[\varepsilon : \tau]$
computations	$c \in Comp$	$::= n \leftarrow f(n_1, \dots, n_k)$
DAIGs	$\mathcal{D}$	$: \mathcal{P}(Ref) \times \mathcal{P}(Comp)$

Figure 6. Demanded Abstract Interpretation Graphs, edge-labelled by analysis functions and connecting named reference cells storing statements and abstract states.

names $n^{(i)}$ allow variants of a single name to be distinguished as loops are unrolled: $n^{(i)}$ is the i th unrolled copy of the name n in a loop. All name equalities are decided structurally.

Values include statements s and abstract states $\varphi \in \Sigma^\#$. Reference cells bind names to values or the absence thereof (denoted ε), while computations specify analysis data-flow dependencies between reference cells.

We denote by $\mathcal{D}[n \mapsto v]$ the DAIG identical to $\mathcal{D}$ except that the reference cell named n now holds value v . We also denote DAIG reachability by $n \rightsquigarrow_{\mathcal{D}} n'$ (eliding the subscript when it is clear from context), and define helper functions *name*, *srcs*, and *dest* to project out, respectively, the name n of a reference cell $n[v_\varepsilon : \tau]$ and the source names $\{n_1, \dots, n_k\}$ or destination name n of a computation $n \leftarrow f(n_1, \dots, n_k)$. Finally, the typing judgment $R \vdash n \leftarrow f(n_1, \dots, n_k)$ holds when n_1 through n_k name references in R with the same types as f 's inputs and n names a reference in R with the same type as f 's output.

Definition 4.1 (DAIG Well-formedness). A DAIG $\langle R, C \rangle$ is subject to the following well-formedness constraints.

(1) References are named uniquely:

$$\forall r, r' \in R. \text{name}(r) = \text{name}(r') \Leftrightarrow r = r'$$

(2) Computations have unique destinations:

$$\forall c, c' \in C. \text{dest}(c) = \text{dest}(c') \Leftrightarrow c = c'$$

(3) Dependencies are acyclic: $\nexists r \in R. \text{name}(r) \rightsquigarrow \text{name}(r)$

(4) Computations are well-typed with respect to references:

$$\forall c \in C. R \vdash c$$

(5) Empty references have dependencies:

$$\forall n[\varepsilon : \tau] \in R. \exists c \in C. n = \text{dest}(c)$$

Beyond these basic well-formedness conditions, a DAIG's structure must also properly encode an abstract interpretation computation over an underlying program. Given a program's CFG and an abstract interpreter interface (as defined in Section 3), there are three general cases shown in Fig. 7 to consider when examining a corresponding DAIG. The key property is that demand-driven query evaluation and incremental edits will evolve the DAIG but preserve the following consistency conditions:

- (1) A forward CFG edge $\ell' \rightarrow [s] \rightarrow \ell$ to a *non-join* location is encoded by a transfer function-labelled DAIG edge, connecting reference cells for its abstract pre-state (named $n_{\ell'}$) and statement label (named $\ell' \cdot \ell$) to a reference cell for its abstract post-state (named n_{ℓ}).⁵
- (2) Forward CFG edges to a *join* location ℓ are a bit more complex, introducing intermediate cells to encode the join $\sqcup$ into the DAIG. For each incoming edge to ℓ with statement s_i , we introduce a three-cell transfer function construct similar to case (1), with an output cell $i \cdot n_{\ell}$ named uniquely for that edge. Then, a single join edge connects each pre-join abstract state $i \cdot n_{\ell}$ to n_{ℓ} , the abstract post-state at ℓ .
- (3) As described informally in Section 2.3, our framework analyzes CFG back edges by unrolling the abstract fixed-point computation to evolve DAIGs on demand, so this diagram is parameterized by a number k of such unrollings. Given the CFG back edge $\ell' \rightarrow [s] \rightarrow \ell$, a transfer function DAIG edge connects the abstract state after one abstract iteration (named $\underline{\ell}^{(0)}$) and s to a pre-widen abstract state at the loop head ℓ (named $\underline{\ell}^{(0)} \cdot \underline{\ell}^{(1)}$), which is connected with the previous abstract state at the loop head ($\underline{\ell}^{(0)}$) to the next ($\underline{\ell}^{(1)}$) via a widen edge⁴. This acyclic structure is repeated k times in the DAIG (with $k = 1$ in the initial

⁵We write n_{ℓ} for the name of the abstract state at ℓ throughout this section: $\underline{\ell}^{(0)}$ if ℓ belongs to any natural loop and $\underline{\ell}$ otherwise. Loop heads ℓ are a special case: n_{ℓ} is $\underline{\ell}^{(0)}$ (the abstract state at loop entry) when the destination of a DAIG edge and $\underline{\ell}$ (the fixed point at ℓ) otherwise.

construction and further unrollings generated on demand as described in Section 5), thereby encoding the unbounded fixed-point computation. Lastly, the fix edge — indicating a dependency on the eventual fixed point — connects the two greatest abstract iterates to the reference cell ($\underline{\ell}$) for the fixed-point abstract state at the loop head.

Definition 4.2 (DAIG–CFG Consistency). A DAIG $\mathcal{D} = \langle R, C \rangle$ is consistent with a program CFG $\langle L, E, \ell_0 \rangle$, written $\mathcal{D} \cong \langle L, E, \ell_0 \rangle$, when it is well-formed and its structure is consistent with that of the program.

A formal statement of the $\mathcal{D} \cong \langle L, E, \ell_0 \rangle$ relation is given in the appendix [46], closely following the structure of the above description and Fig. 7.

The above establishes when the structure of a DAIG is consistent with the program’s CFG. A DAIG is consistent with an abstract interpretation of the program when the partial analysis results stored in the DAIG are consistent with the partial abstract interpretation, or formally as follows:

Definition 4.3 (DAIG–AI Consistency). A DAIG $\mathcal{D} = \langle R, C \rangle$ is consistent with an abstract interpreter $\langle \Sigma^{\#}, \varphi_0, [\![\cdot]\!], \sqsubseteq, \sqcup, \nabla \rangle$, written $\mathcal{D} \cong \langle \Sigma^{\#}, \varphi_0, [\![\cdot]\!], \sqsubseteq, \sqcup, \nabla \rangle$, when all partial analysis results stored in R are consistent with the computations encoded by C , such that $\underline{\ell}_0[\varphi_0 : \Sigma^{\#}] \in R$ and

$$\forall n[v : \Sigma^{\#}] \in R, n \leftarrow f(n_1, \dots, n_k) \in C.$$

$$\{n_i[v_i : \tau_i] \mid 1 \leq i \leq k\} \subseteq R \wedge \begin{cases} v = v_1 = v_2 & \text{if } f = \text{fix} \\ v = f(v_1, \dots, v_k) & \text{otherwise} \end{cases}$$

Figure 7. DAIG–CFG Consistency (Definition 4.2) in diagram form, illustrating how different CFG structures are encoded into DAIG structures. In subfigure (3), we apply some ad-hoc shorthands for the DAIG encoding of the loop body: $\mathcal{D}_{\text{stmt}}$ contains all of its statement reference cells, while $\mathcal{D}_{\Sigma^{\#}}^{(i)}$ contains all of its abstract state reference cells, with iteration counts set to i . Each dotted line from $\mathcal{D}_{\text{stmt}}$ thus represents one or more DAIG edges, from each statement to corresponding abstract states.

Given a program's CFG and a generic abstract interpretation interface, we can construct an initial DAIG that is consistent with a classical (batch) abstract interpretation:

Lemma 4.1 (Initial DAIG Construction, Well-Formedness, CFG-Consistency, and AI-Consistency). *There exists a constructive procedure $\mathcal{D}_{\text{init}}$ such that for all well-formed programs $\langle L, E, \ell_0 \rangle$, the initial DAIG*

$$\mathcal{D} = \mathcal{D}_{\text{init}}(\langle L, E, \ell_0 \rangle, \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle)$$

is well formed and consistent with both the target program (i.e. $\mathcal{D} \cong \langle L, E, \ell_0 \rangle$) and underlying abstract interpreter (i.e. $\mathcal{D} \cong \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$).

Proof sketch. We define such a $\mathcal{D}_{\text{init}}$ in the appendix [46] and show that it produces well-formed results consistent with both the target program and the underlying abstract interpreter. Its definition tracks closely with the informal and diagrammatic descriptions above, constructing DAIG structures that correspond to the input CFG. $\square$

5 Demanded AI by Evaluating DAIGs

In this section, we give an operational semantics for demand-driven and incremental evaluation of DAIGs. A state in the operational semantics consists of a DAIG $\mathcal{D}$ and also an auxiliary memoization table M , which can be used to reuse previously-computed analysis results independent of program location. Memoization tables are finite maps from names n to abstract states $\varphi \in \Sigma^\#$, and we write $M(n)$ for the abstract state mapped to by n in M , $\text{dom}(M)$ for the set of names in the domain of M , and $M[n \mapsto \varphi]$ for the extension of M with a new mapping from n to φ .

The DAIG operational semantics are split into two judgments, corresponding to *queries* and *edits* over both analysis results and program syntax. Both interaction modes are given in a small-step style, describing the effects of each operation on the DAIG and auxiliary memo table.

5.1 Query Evaluation Semantics

A query for the value of the reference cell with name n , given some initial DAIG $\mathcal{D}$ and auxiliary memo table M , yields a value v and (possibly unchanged) DAIG and memo-table structures $\mathcal{D}'$ and M' . This operation is defined inductively by the $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$ judgment form, whose inference rules are given in Fig. 8.

There are two potential ways to reuse previously-computed analysis results: either in DAIG $\mathcal{D}$ or the auxiliary memo table M . The Q-REUSE rule handles the case where the DAIG cell named by n already holds a value, returning that value and leaving the DAIG and memo table unchanged.

The Q-MATCH and Q-MISS rules handle the case where n is empty in $\mathcal{D}$. In both cases, queries are issued for the input cells n_1 through n_k to the computation f that outputs to n . In Q-MATCH, the auxiliary memo table is matched: f has already been computed for the relevant inputs, so the

$$\begin{array}{c}
 \boxed{\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'} \\
 \text{Q-REUSE} \quad \frac{n[v : \tau] \in R}{\langle R, C \rangle, M \vdash n \Rightarrow v; \langle R, C \rangle, M} \\
 \text{Q-MATCH} \quad \frac{\begin{array}{l} \mathcal{D}_0 = \langle R, C \rangle \quad n[\varepsilon : \tau] \in R \quad n \leftarrow f(n_1, \dots, n_k) \in C \\ \mathcal{D}_{i-1}, M_{i-1} \vdash n_i \Rightarrow v_i; \mathcal{D}_i, M_i \quad (\text{for } i \in [1, k]) \\ \underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k) \in \text{dom}(M_k) \quad v = M_k(\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k)) \end{array}}{\mathcal{D}_0, M_0 \vdash n \Rightarrow M_k(\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k)); \mathcal{D}_k[n \mapsto M_k(\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k))], M_k} \\
 \text{Q-MISS} \quad \frac{\begin{array}{l} \mathcal{D}_0 = \langle R, C \rangle \quad n[\varepsilon : \tau] \in R \quad n \leftarrow f(n_1, \dots, n_k) \in C \\ \mathcal{D}_{i-1}, M_{i-1} \vdash n_i \Rightarrow v_i; \mathcal{D}_i, M_i \quad (\text{for } i \in [1, k]) \\ \underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k) \notin \text{dom}(M_k) \quad v = f(v_1, \dots, v_k) \quad f \neq \text{fix} \end{array}}{\mathcal{D}_0, M_0 \vdash n \Rightarrow v; \mathcal{D}_k[n \mapsto v], M_k[\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k) \mapsto v]} \\
 \text{Q-LOOP-CONVERGE} \quad \frac{\begin{array}{l} n[\varepsilon : \tau] \in R \quad n \leftarrow \text{fix}(n_1, n_2) \in C \\ \langle R, C \rangle, M \vdash n_1 \Rightarrow v; \mathcal{D}', M' \quad \mathcal{D}', M' \vdash n_2 \Rightarrow v; \mathcal{D}'', M'' \end{array}}{\langle R, C \rangle, M \vdash n \Rightarrow v; \mathcal{D}''[n \mapsto v], M''} \\
 \text{Q-LOOP-UNROLL} \quad \frac{\begin{array}{l} n[\varepsilon : \tau] \in R \quad c = n \leftarrow \text{fix}(\underline{\ell}^{(k-1)}, \underline{\ell}^{(k)}) \in C \\ \langle R, C \rangle, M \vdash \underline{\ell}^{(k-1)} \Rightarrow v'; \mathcal{D}', M' \\ \mathcal{D}', M' \vdash \underline{\ell}^{(k)} \Rightarrow v''; \mathcal{D}'', M'' \\ v' \neq v'' \quad \text{unroll}(\mathcal{D}'', c), M'' \vdash n \Rightarrow v; \mathcal{D}''', M''' \end{array}}{\langle R, C \rangle, M \vdash n \Rightarrow v; \mathcal{D}''', M'''}
 \end{array}$$

Figure 8. Operational semantics rules governing *queries* for the contents of a DAIG. The judgment form $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$ is read as “Requesting n from DAIG $\mathcal{D}$ with auxiliary memo table M yields value v , updated DAIG $\mathcal{D}'$, and updated memo table M' .”

result is retrieved from M_k , the memo table after querying the input cells, and stored in n (i.e., via $\mathcal{D}_k[n \mapsto M_k(\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k))]$). Note that this notation is a low-level mutation of the reference cell named by n , not an external edit that would trigger invalidation (which we will describe below in Section 5.3).

Q-MISS handles memo table misses by computing and memoizing $f(v_1, \dots, v_k)$ before storing the result in both the DAIG $\mathcal{D}$ and the auxiliary memo table M (i.e., at names n and $\underline{f} \cdot (\underline{v}_1 \dots \underline{v}_k)$, respectively).

5.2 Demanded Fixed Points

Demanded unrolling is the process by which we compute abstract interpretation fixed-points over cyclic control flow graphs without introducing cyclic dependencies into DAIGs, as described informally in Section 2 and represented graphically in Fig. 4c. The semantics are formalized by the Q-LOOP-CONVERGE and Q-LOOP-UNROLL rules in Fig. 8.

Recall that fix is a special function symbol indicating an analysis fixed-point computation. The destination of a fix edge is a loop-head cell for storing a fixed-point invariant, and its sources are the two greatest abstract iterates of said loop head yet computed. When those abstract iterates have the same value v , the analysis has reached a fixed point⁴, and a query for the fixed point may return v . However, when they are unequal, a query triggers an unrolling in the DAIG: the loop body's abstract state reference cells are unrolled one more iteration, the fix edge's sources are shifted forward one iteration, and then the fixed-point query is reissued.

This procedure differs from concrete/syntactic loop unrolling – e.g. as applied by an optimizing compiler or bounded model checker – in that it applies to the DAIG's reified abstract interpretation computation, including joins and widens, *not* to the concrete syntax of the program under analysis. That is, the k -th demanded unrolling corresponds to the k -th application of the loop body's abstract semantics (i.e. the k -th abstract iteration) rather than the k -th concrete execution of the loop. As a result, it is sound with respect to the concrete semantics of the program under analysis and is guaranteed to converge.

As the name suggests, Q-LOOP-CONVERGE applies when the abstract interpretation has reached a fixed point. Since the dependencies of the fix edge, n_1 and n_2 , are consecutive abstract iterates at the head of the corresponding loop, their evaluation to the same value v indicates that loop analysis has converged, so v may be stored in the DAIG and returned.

On the other hand, Q-LOOP-UNROLL applies when the abstract interpretation has not yet reached a fixed point, since the two most recent abstract iterates are unequal. In this case, the loop is unrolled once by the unroll helper function and the query for the fixed point is reissued. The unroll helper function used in Q-LOOP-UNROLL takes a DAIG $\mathcal{D}$ and a fix edge and unrolls the loop corresponding to the fix edge by one iteration in $\mathcal{D}$. It is defined as follows:

$$\text{unroll} \left(\langle R, C \rangle, c = \underline{\ell} \leftarrow \text{fix}(\underline{\ell}^{(k-1)}, \underline{\ell}^{(k)}) \right) \triangleq \langle R', C' \rangle, \text{ where}$$

$$R' = R \cup \{ \text{incr}(n)[\varepsilon : \Sigma^\#] \mid \underline{\ell}^{(k-1)} \rightsquigarrow n \rightsquigarrow \underline{\ell}^{(k)} \} \text{ and}$$

$$C' = C / \{ c \} \cup \left\{ \underline{\ell} \leftarrow \text{fix} \left(\underline{\ell}^{(k)}, \underline{\ell}^{(k+1)} \right) \right\} \cup \left\{ \text{incr-c}(c) \mid \underline{\ell}^{(k-1)} \rightsquigarrow \text{dest}(c) \rightsquigarrow \underline{\ell}^{(k)} \right\}$$

where incr and incr-c increment the iteration counts of names. Intuitively, unroll takes the region of the DAIG forwards-reachable from the $(k-1)^{\text{th}}$ abstract iterate $\underline{\ell}^{(k-1)}$ and backwards-reachable from the k^{th} abstract iterate $\underline{\ell}^{(k)}$ and duplicates it while incrementing all name's iterations counts from $k-1$ to k , then shifts the fix edge forward one iteration. Crucially, this operation preserves the DAIG acyclicity invariant.

5.3 Incremental Edit Semantics

An *edit* to a DAIG $\mathcal{D}$ occurs when a value v is written to some reference cell named n in $\mathcal{D}$ by an external mutator. This edit must both update n and also clear the value of (or

$$\boxed{\mathcal{D} \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}'}$$

$$\begin{array}{c} \text{E-COMMIT} \\ \forall c \in C . n \in \text{srcs}(c) \implies \text{dest}(c)[\varepsilon : \tau] \in R \\ \underline{\ell} \leftarrow \text{fix}(\underline{\ell}^{(i)}, \underline{\ell}^{(i+1)}) \in C \implies \underline{\ell}^{(1)}[\varepsilon : \tau] \in R \\ v_\varepsilon = \varepsilon \implies \exists c \in C . n = \text{dest}(c) \\ v_\varepsilon \neq \varepsilon \implies \exists n[_ : \tau] \in R . v_\varepsilon : \tau \\ \hline \langle R, C \rangle \vdash n \Leftarrow v_\varepsilon ; \langle R, C \rangle[n \mapsto v_\varepsilon] \end{array}$$

$$\begin{array}{c} \text{E-PROPAGATE} \\ n \rightsquigarrow_{\mathcal{D}} n' \quad \mathcal{D} \vdash n' \Leftarrow \varepsilon ; \mathcal{D}' \quad \mathcal{D}' \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}'' \\ \hline \mathcal{D} \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}'' \end{array}$$

$$\begin{array}{c} \text{E-LOOP} \\ \underline{\ell} \leftarrow \text{fix}(\underline{\ell}^{(k-1)}, \underline{\ell}^{(k)}) \in C \quad \langle R, C' \rangle \vdash \underline{\ell}^{(1)} \Leftarrow \varepsilon ; \mathcal{D}' \\ C' = C / \left\{ \underline{\ell} \leftarrow \text{fix}(\underline{\ell}^{(k-1)}, \underline{\ell}^{(k)}) \right\} \cup \left\{ \underline{\ell} \leftarrow \text{fix}(\underline{\ell}^{(0)}, \underline{\ell}^{(1)}) \right\} \\ \hline \langle R, C \rangle \vdash \underline{\ell}^{(k)} \Leftarrow \varepsilon ; \mathcal{D}' \end{array}$$

Figure 9. Operational semantics rules governing *edits* to the contents of a DAIG. The judgment $\mathcal{D} \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}'$ is read as “Editing reference cell n of DAIG $\mathcal{D}$ with value v_ε yields updated DAIG $\mathcal{D}'$,” where v_ε ranges over values v and the “empty” symbol ε .

“dirty”) any reference cell that (transitively) depends on n . Here we give a full definition of the edit operation, using the $\mathcal{D} \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}'$ judgment given in Fig. 9.

As described informally in Section 2, invalidation proceeds by dirtying forwards in the acyclic DAIG, except that the implicit cyclic dependency from fix edges must be accounted for, by rolling back to a non-dirty source cell.

The E-COMMIT rule is a base case: if the edited cell's downstream dependencies are all empty, then the edit may be performed directly. Its second premise accounts for the implicit dependency of abstract iterate cells $\underline{\ell}^{(i)}$ for $i > 0$ on the fixed point cell $\underline{\ell}$ (corresponding to the loop back edge in the control-flow graph); it suffices to check that the 1st abstract iterate cell has been emptied, as all abstract iterates $\underline{\ell}^{(i)}$ for $i > 1$ are reachable from $\underline{\ell}^{(1)}$. The third and fourth premises ensure that DAIG well-formedness is preserved, by preventing emptying of source nodes and ill-typed edits respectively.

The E-PROPAGATE rule recursively empties reference cells that depend on the edited cell, eventually bottoming out when no such cells are non-empty and E-COMMIT can be used to derive the $\mathcal{D}' \vdash n \Leftarrow v_\varepsilon ; \mathcal{D}''$ premise. Note that there is no recomputation here in E-PROPAGATE, only emptying.

The E-LOOP rule applies when the final abstract iterate of a loop is dirtied. The sources of its fix edge are reset to its 0th and 1st abstract iterates and dirtying continues from the 1st abstract iterate. This handling is slightly more conservative than necessary in the case that an intermediate abstract iterate (i.e., with $k > 1$) is edited since results from some previous iterations (up to that k) may not need to be

discarded, but it simplifies the presentation and handles all program edits with maximal reuse.

6 Soundness, Termination, and From-Scratch Consistency

In this section, we state two key properties of demanded abstract interpretation graphs: *from-scratch consistency*, which guarantees that DAIG query results are identical to the analysis results computed by the underlying abstract interpreter at a global fixed-point, and *query termination*, which guarantees termination of the DAIG query semantics even in the presence of unbounded abstract loop unrolling. The proofs are deferred to the appendix [46] due to space constraints.

Both theorems rely on the preservation of DAIG well-formedness (Definition 4.1), DAIG-CFG consistency (Definition 4.2), and DAIG-AI consistency (Definition 4.3) under queries and program edits.

Lemma 6.1 (DAIG Well-Formedness Preservation).

If $\mathcal{D}$ is well-formed and either $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$ or $\mathcal{D} \vdash n \Leftarrow v_\varepsilon; \mathcal{D}'$, then $\mathcal{D}'$ is well-formed.

Lemma 6.2 (DAIG-CFG Consistency Preservation).

If $\mathcal{D} \approx \langle L, E, \ell_0 \rangle$ then:

- if $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$ then $\mathcal{D}' \approx \langle L, E, \ell_0 \rangle$;
- if $\mathcal{D} \vdash n \Leftarrow s; \mathcal{D}'$ then $\mathcal{D}' \approx \langle L, E', \ell_0 \rangle$, where E' is E with the edit applied (see appendix [46] for details).

Lemma 6.3 (DAIG-AI Consistency Preservation).

If $\mathcal{D} \approx \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$ and either $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$ or $\mathcal{D} \vdash n \Leftarrow s; \mathcal{D}'$, then $\mathcal{D}' \approx \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$.

With these preservation results, we can now prove that DAIG query results for abstract states at program locations are from-scratch consistent with the global fixed-point invariant map of the DAIG's underlying abstract interpreter.

Theorem 6.1 (DAIG From-Scratch Consistency). For all sound M and well-formed $\mathcal{D}$ such that $\mathcal{D} \approx \langle L, E, \ell_0 \rangle$ and $\mathcal{D} \approx \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$, if $\mathcal{D}, M \vdash \ell \Rightarrow v; \mathcal{D}', M'$ then $v = \llbracket \ell \rrbracket_{\langle L, E, \ell_0 \rangle}^{\#*}$.

Corollary 6.2. Query results are sound.

Since the global invariant map $\llbracket \cdot \rrbracket^{\#*}$ of the underlying abstract interpreter $\langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$ is sound (by Global Abstract Interpreter Soundness (Proposition 3.2)) and a DAIG query for the abstract state at a location ℓ returns $\llbracket \ell \rrbracket^{\#*}$, DAIG query results themselves are sound.

Theorem 6.3 (DAIG Query Termination). For all M and well-formed $\mathcal{D}$ such that $\mathcal{D} \approx \langle L, E, \ell_0 \rangle$ and $\mathcal{D} \approx \langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$, if n is in the namespace of $\mathcal{D}$ then there exist $v, \mathcal{D}', M'$ such that $\mathcal{D}, M \vdash n \Rightarrow v; \mathcal{D}', M'$.

7 Implementation and Evaluation

Here, we describe a prototype implementation and evaluation of demanded abstract interpretation via our DAIG framework. Our evaluation studied two research questions:

- **Expressivity:** Does the DAIG framework allow for clean and straightforward implementations of rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?
- **Scalability:** For these rich analysis domains, what degree of performance improvement can be obtained by performing incremental and/or demand-driven analysis, as compared to batch analysis?

7.1 Implementation

Our DAIG framework is implemented in approximately 2,500 lines of OCaml code [44, 45]. Incremental and demand-driven analysis logic, including demanded unrolling, operates over an explicit graph representation of DAIGs, but per-function memoization (i.e., the auxiliary memo table M in the Fig. 8 semantics) is provided via `adapton.ocaml`, an open-source implementation of the technique of Hammer et al. [21].

Our implementation is parametric in an abstract domain, and the effort required to instantiate the framework to a new abstract domain is comparable to the effort required to do so in a classical abstract interpreter framework. The required module signature is essentially the abstract interpreter signature $\langle \Sigma^\#, \varphi_0, \llbracket \cdot \rrbracket^\#, \sqsubseteq, \sqcup, \nabla \rangle$, extended with some standard utilities which can often be automatically derived.

Interprocedurality. Although the formalism is defined over control-flow graphs for clarity and brevity, our implementation supports context-sensitive analysis of non-recursive programs with static calling semantics (i.e., no virtual dispatch or higher-order functions).

In order to analyze such programs, we initially construct a DAIG only for the “main” procedure in the initial context. Then, when a query is issued for the abstract state after a call, we construct a DAIG for its callee in the proper context. When a query is issued at a location/context for which no DAIG has yet been constructed, we construct the DAIG for its containing function and analyze dataflow to its entry.

These operations are parametric in a context-sensitivity policy for choosing a context in which to analyze a callee at a call site. Our implementation includes functors that implement context-insensitivity and also 1- and 2-call-site-sensitivity [39].

7.2 Expressivity

To demonstrate the expressivity of our DAIG framework, we have instantiated it with three existing well-known abstract interpretation techniques — interval, octagon and shape analysis — all of which are inexpressible using existing incremental and/or demand-driven analysis frameworks. Here, we describe our experience applying the DAIG-based interval

and shape analyses to a small set of programs. In Section 7.3, we use the octagon domain to investigate the scalability of demanded analysis on synthetic benchmarks.

Together, these analysis implementations provide evidence for our approach’s agnosticity to the underlying abstract domain, including domains with black-box external dependencies and/or complicated non-monotone abstract operations.

Interval Analysis. The interval abstract domain is a textbook example of an infinite-height lattice, requiring widening to guarantee analysis convergence. An interval $[l, u]$ abstracts the set of numbers between lower bound l and upper bound u . Join and widen operations and abstract states are defined in the standard way [11]. Abstract interpretation in this domain is known as interval analysis and is commonly used, e.g., to verify the safety of array accesses. Interval analysis has been applied at industrial scale, for example by Cousot et al. [12].

In practice, it is common to use an optimized off-the-shelf interval abstract domain such as that of APRON [25] or Elina [40]. We have implemented an APRON-backed interval analysis for JavaScript programs in the DAIG framework. As an indication of the flexibility of our framework, we were able to use the APRON library *without modification*.

In order to validate our implementation, we analyzed 23 array-manipulating programs — with functions such as `contains`, `equals`, `swap`, and `indexOf` — from the test suite of Buckets.JS, a JavaScript data structure library [38].

Using the 2-call-string-sensitive context policy, our analysis verified the safety of all 85 array accesses in the programs; with 1-call-string-sensitivity, it verified 71/74 (96%), and with context-insensitive analysis it verified 4/18 (22%). These figures show that standard numerical analyses on DAIGs behave as they would in a batch analysis engine.

Shape Analysis. Precise analysis of recursive data structures such as linked lists is essential in many domains. Such analysis relies on complex abstract domains that cannot be expressed in existing frameworks for incremental and demand-driven analysis. We have implemented a DAIG-based demanded shape analysis for singly-linked lists. An abstract state in this shape domain is a triple consisting of

- A separation logic formula over points-to ($\alpha.f \mapsto \alpha'$) and list-segment ($\text{lseg}(\alpha, \alpha')$) atomic propositions, stating respectively that the f field of the object at symbolic address α points to α' and that there exists a sequence of next pointer dereferences from α to α' [34],
- A collection of pure constraints: equalities and disequalities over memory addresses, and
- An environment mapping variables to memory addresses.

Join, widen, and implication all rely on a collection of rewrite rules over such states from Chang et al. [10] (specialized to a fixed inductive definition for list segments). All told, the

implementation of this shape domain requires approximately 500 lines of OCaml code.

We have applied this DAIG-based shape analysis to successfully verify the correctness and memory-safety of the list append procedure of Fig. 2, along with several linked list utilities from the aforementioned Buckets.js library including `foreach` and `indexOf` [38]. Analysis of the ℓ_3 -to- ℓ_4 -to- ℓ_3 loop of the list append procedure converges in one demanded unrolling with a precise result.

7.3 Scalability

To study scalability, we conducted an initial investigation of what performance improvements are possible with demanded analysis variants in our framework. We compared the performance of analysis with and without incrementality and demand on interleaved sequences of program edits and queries. Throughout this section, our framework is instantiated with a context-insensitive APRON-backed octagon domain: a relational numerical domain representing invariants of the form $\pm x \pm y \leq y$, widely used in practice due to its balance of expressivity and efficiency [29].

To exercise the analyses, we created synthetic workloads consisting of 3,000 random edits to an initially-empty program. Programs are generated in a JavaScript subset with assignment, arrays, conditional branching, while loops, and (non-recursive) function calls of the form $x = f(y)$. An “edit” is an insertion of a randomly generated statement, if-then-else conditional, or while loop at a randomly-sampled program location, with 85%, 10%, and 5% probability respectively, and statements and expressions are generated probabilistically from their respective grammars.

We evaluate four analysis configurations on this workload:

- (1) *Batch* analysis: Classical whole-program abstract interpretation, fully re-analyzing the entire program from scratch in response to each edit.
- (2) *Incremental* analysis: An incremental-only configuration which applies the edit semantics to dirty as few previously-computed analysis results as possible, but eagerly recomputes all dirtied cells.
- (3) *Demand-driven* analysis: A demand-driven-only configuration which dirties the full DAIG after each edit, but applies the query semantics to avoid computing analysis results that aren’t demanded.
- (4) *Incremental & demand-driven* analysis: The full demanded abstract interpretation technique, which applies both the edit and query semantics to maximize reuse and minimize redundant computation.

In the demand-driven configurations, queries are issued at five randomly-sampled program locations between each edit. Note that since the first three configurations were implemented atop our DAIG framework, which is designed to support both incremental and demand-driven analysis, they may not be as tuned as specialized implementations.

Figure 10. Performance of octagon analysis on the synthetic workload of interleaved program edits and analysis queries described in Section 7.3. The four scatter plots show the scaling of each configuration as the program size is increased by edits, and their color-coding serves as a legend to the fifth figure: a cumulative distribution plot showing the fraction of analysis runs (y axis) completed by each configuration within some time interval (x axis). Lastly, the table shows summary statistics for each configuration, including the mean, median, 90th, 95th, and 99th percentile analysis latency.

Each plot includes data points from 9 separate trials, with fixed random seeds such that the same edits (and, in the two demand-driven configurations, queries) are issued to each configuration. In total, this data set includes 27,000 analysis executions in each exhaustive configuration and 135,000 queries in each demand-driven configuration.

The results, as shown in Fig. 10, indicate that while incremental and demand-driven analysis each significantly improve analysis latencies with respect to the batch analysis baseline, combining the two provides an additional large reduction in latency. This effect is most apparent in the tail of the distribution, since edits that dirty large regions of the program are costly for incremental analysis, and queries that depend on large regions of the graph are costly for demand-driven analysis. By combining incremental dirtying with demand-driven evaluation, demanded abstract interpretation mitigates these worst-case scenarios and consistently keeps analysis costs low even as the program grows.

In particular, at the 95th percentile, the 1.2s latency of incremental demand-driven analysis is more than five times lower than the next best configuration, and potentially low enough to support interactive use. Fig. 10 gives a cumulative distribution of analysis latencies, again showing the large advantage of the incremental demand-driven analysis over other configurations.

8 Related Work

Incremental Computation. Techniques for the efficient caching and reuse of computation results, particularly those based on memoization of pure functions [1, 20, 30] and dependency graphs [13, 32], have been the subject of a great deal of research and seen widespread practical application.

More recently, dependency graph-based approaches to incremental computation have improved on generic memoization and graph-based techniques, allowing for fine-grained automatic caching and reuse even in the presence of changes to inputs or an underlying data store [2, 3]. Building on these graph-based techniques for self-adjusting computation, some recent work has focused on support for interactive and demand-driven computations [21, 22]. Although this approach yields a general and powerful system for incremental computation, its low-level primitives make it difficult to express the complex fixed-point computation over cyclic control-flow graphs in arbitrary abstract interpretations. We take inspiration from demanded computation graphs but instead specialize the language of demanded computations to demanded abstract interpretations, both with syntactic structures and with a query/edit semantics which dynamically modifies the dependency graph to model such computations.

Incremental Analysis. The application of incremental computation to program analysis is similarly well-studied, going back at least to the development of incremental dataflow analyses to support responsive continuous compilation [35, 49]. Recent work has contributed incremental versions of several classes of program analysis, including IFD-S/IDE dataflow analyses [5, 16] and analyses based on extensions to Datalog [47, 48]. These specialized approaches offer effective solutions for certain classes of program analysis, but place restrictions on abstract domains that rule out arbitrary abstract interpretations in infinite-height domains.

Compositional program analysis, in which summaries are computed for individual files or compilation units rather than a whole program, naturally supports incrementality in the sense that results need only be recomputed for changed files. This has shown to be very effective for scaling program analyses to massive codebases in CI/CD systems [9, 14, 18], but it operates at a much coarser granularity than both the aforementioned approaches and our own, since it is designed to scale up to massive programs rather than to minimize analysis latencies at development-time.

Leino and Wüstholtz [26] propose a fine-grained incremental verification technique for the Boogie language, which verifies user-provided specifications of imperative procedures. These specifications include loop invariants, allowing their algorithm to ignore cyclic dependencies altogether.

Demand-Driven Analysis. Demand-driven techniques for dataflow analysis are also well-studied. The intra-procedural problem was studied by Babich and Jazayeri [6]. Several extensions to inter-procedural analysis have been presented, for example, by Reps [33], Duesterwald et al. [17], and Sagiv et al. [37]. In nearly all cases previous work has been focused on finite domains. The work of Sagiv et al. [37] allows for infinite domains of finite height, but does not consider infinite-height domains like intervals.

Any static analysis expressible as a context-free-language reachability (CFL-reachability) problem can be computed in a demand-driven fashion as a “single-source” problem [31]. As such, a number of papers have presented demand-driven algorithms for flow-insensitive pointer analysis [23, 42, 43].

Reference attribute grammars (RAGs) are declarative specifications of properties over ASTs (including potentially-cyclic flow analyses) which can be evaluated incrementally and on-demand [28, 41]. Termination of RAG evaluation requires that all cyclic computations converge to a fixed-point in finitely-many iterations [19, 28]; this convergence property holds for finite domains with monotone operators but may also be achieved through other means (e.g. widening).

Improving on prior work, our framework comes with proofs of termination and from-scratch consistency, and specifies the exact conditions required to ensure termination in infinite-height domains with non-monotone widening operators.

9 Conclusion

We have presented a novel framework for demanded abstract interpretation, in which an arbitrary abstract interpretation can be made both incremental and demand-driven. Unlike previous frameworks, ours supports arbitrary lattices and widening operators. The framework is based on a novel demanded abstract interpretation graph (DAIG) representation of the analysis problem, where careful handling of loops ensures the DAIG remains acyclic. We have proved various key properties of the framework, including soundness, termination, and from-scratch consistency. Our implementation shows that complex analyses can be easily implemented with our framework, with the potential for significant performance wins in incremental and demand-driven scenarios.

Acknowledgments

We thank Matthew A. Hammer and Jared Wright for their valuable contributions in the early stages of this research. We also thank the anonymous reviewers and members of the CUPLV lab for their helpful reviews and suggestions. This research was supported in part by the National Science Foundation under grants CCF-1619282, CCF-2008369, and CCF-2007024, and also by a gift from Oracle Labs.

References

- [1] Martín Abadi, Butler W. Lampson, and Jean-Jacques Lévy. 1996. Analysis and Caching of Dependencies. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/232627.232638>
- [2] Umut A. Acar, Amal Ahmed, and Matthias Blume. 2008. Imperative self-adjusting computation. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1328438.1328476>
- [3] Umut A. Acar, Guy E. Blelloch, and Robert Harper. 2002. Adaptive functional programming. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1186634>
- [4] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. 2006. *Compilers: Principles, Techniques, and Tools (2nd Edition)*.
- [5] Steven Arzt and Eric Bodden. 2014. Reviser: efficiently updating IDE-/IFDS-based data-flow analyses in response to incremental program changes. In *International Conference on Software Engineering (ICSE)*. <https://doi.org/10.1145/2568225.2568243>
- [6] Wayne A. Babich and Mehdi Jazayeri. 1978. The Method of Attributes for Data Flow Analysis: Part II. Demand analysis. *Acta Informatica* 3 (1978). <https://doi.org/10.1007/BF00264320>
- [7] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2005. Small-foot: Modular Automatic Assertion Checking with Separation Logic. In *Formal Methods for Components and Objects (FMCO)*. https://doi.org/10.1007/11804192_6
- [8] François Bourdoncle. 1993. Efficient chaotic iteration strategies with widenings. In *Formal Methods in Programming and Their Applications*. <https://doi.org/10.1007/BFb0039704>
- [9] Cristiano Calcagno and Dino Distefano. 2011. Infer: An Automatic Program Verifier for Memory Safety of C Programs. In *NASA Formal Methods (NFM)*. https://doi.org/10.1007/978-3-642-20398-5_33
- [10] Bor-Yuh Evan Chang, Xavier Rival, and George C. Necula. 2007. Shape Analysis with Structural Invariant Checkers. In *Static Analysis (SAS)*. https://doi.org/10.1007/978-3-540-74061-2_24
- [11] Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction

- or Approximation of Fixpoints. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/512950.512973>
- [12] Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, and Xavier Rival. 2005. The ASTREÉ Analyzer. In *European Symposium on Programming (ESOP)*. https://doi.org/10.1007/978-3-540-31987-0_3
- [13] Alan J. Demers, Thomas W. Reps, and Tim Teitelbaum. 1981. Incremental Evaluation for Attribute Grammars with Application to Syntax-Directed Editors. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/567532.567544>
- [14] Dino Distefano, Manuel Fähndrich, Francesco Logozzo, and Peter W. O'Hearn. 2019. Scaling static analyses at Facebook. *Commun. ACM* 8 (2019). <https://doi.org/10.1145/3338112>
- [15] Dino Distefano, Peter W. O'Hearn, and Hongseok Yang. 2006. A Local Shape Analysis Based on Separation Logic. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. https://doi.org/10.1007/11691372_19
- [16] Lisa Nguyen Quang Do, Karim Ali, Benjamin Livshits, Eric Bodden, Justin Smith, and Emerson R. Murphy-Hill. 2017. Just-in-time Static Analysis. In *Software Testing and Analysis (ISSTA)*. <https://doi.org/10.1145/3092703.3092705>
- [17] Evelyn Duesterwald, Rajiv Gupta, and Mary Lou Soffa. 1995. Demand-Driven Computation of Interprocedural Data Flow. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/199448.199461>
- [18] Manuel Fähndrich and Francesco Logozzo. 2010. Static Contract Checking with Abstract Interpretation. In *Formal Verification of Object-Oriented Software (FoVeOOS)*. https://doi.org/10.1007/978-3-642-18070-5_2
- [19] Rodney Farrow. 1986. Automatic generation of fixed-point-finding evaluators for circular, but well-defined, attribute grammars. In *Compiler Construction (CC)*. <https://doi.org/10.1145/12276.13320>
- [20] John Field and Tim Teitelbaum. 1990. Incremental Reduction in the lambda Calculus. In *LISP and Functional Programming*. <https://doi.org/10.1145/91556.91679>
- [21] Matthew A. Hammer, Jana Dunfield, Kyle Headley, Nicholas Labich, Jeffrey S. Foster, Michael W. Hicks, and David Van Horn. 2015. Incremental computation with names. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/2814270.2814305>
- [22] Matthew A. Hammer, Yit Phang Khoo, Michael Hicks, and Jeffrey S. Foster. 2014. Adaption: composable, demand-driven incremental computation. In *Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/2594291.2594324>
- [23] Nevin Heintze and Olivier Tardieu. 2001. Demand-Driven Pointer Analysis. In *Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/378795.378802>
- [24] Susan Horwitz, Thomas W. Reps, and Shmuel Sagiv. 1995. Demand Interprocedural Dataflow Analysis. In *Foundations of Software Engineering (FSE)*. <https://doi.org/10.1145/222124.222146>
- [25] Bertrand Jeannet and Antoine Miné. 2009. Apron: A Library of Numerical Abstract Domains for Static Analysis. In *Computer-Aided Verification (CAV)*. https://doi.org/10.1007/978-3-642-02658-4_52
- [26] K. Rustan M. Leino and Valentin Wüstholtz. 2015. Fine-Grained Caching of Verification Results. In *Computer-Aided Verification (CAV)*. https://doi.org/10.1007/978-3-319-21690-4_22
- [27] Stephen Magill, Aleksandar Nanevski, Edmund Clarke, and Peter Lee. 2006. Inferring invariants in separation logic for imperative list-processing programs. In *Semantics, Program Analysis, and Computing Environments for Memory Management (SPACE)*.
- [28] Eva Magnusson and Görel Hedin. 2007. Circular reference attributed grammars - their evaluation and applications. *Sci. Comput. Program.* 1 (2007). <https://doi.org/10.1016/j.scico.2005.06.005>
- [29] Antoine Miné. 2006. The octagon abstract domain. *High. Order Symb. Comput.* 1 (2006). <https://doi.org/10.1007/s10990-006-8609-1>
- [30] William Pugh and Tim Teitelbaum. 1989. Incremental Computation via Function Caching. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/75277.75305>
- [31] Thomas Reps. 1998. Program analysis via graph reachability. *Information and Software Technology* 11-12 (1998). [https://doi.org/10.1016/S0950-5849\(98\)00093-7](https://doi.org/10.1016/S0950-5849(98)00093-7)
- [32] Thomas W. Reps. 1982. Optimal-Time Incremental Semantic Analysis for Syntax-Directed Editors. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/582153.582172>
- [33] Thomas W. Reps. 1994. Solving Demand Versions of Interprocedural Analysis Problems. In *Compiler Construction (CC)*. https://doi.org/10.1007/3-540-57877-3_26
- [34] John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *Logic in Computer Science (LICS)*. <https://doi.org/10.1109/LICS.2002.1029817>
- [35] Barbara G. Ryder. 1983. Incremental Data Flow Analysis. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/567067.567084>
- [36] Caitlin Sadowski, Edward Aftandilian, Alex Eagle, Liam Miller-Cushon, and Ciera Jaspan. 2018. Lessons from building static analysis tools at Google. *Commun. ACM* 4 (2018). <https://doi.org/10.1145/3188720>
- [37] Shmuel Sagiv, Thomas W. Reps, and Susan Horwitz. 1996. Precise Interprocedural Dataflow Analysis with Applications to Constant Propagation. *Theor. Comput. Sci.* 1&2 (1996). [https://doi.org/10.1016/0304-3975\(96\)00072-2](https://doi.org/10.1016/0304-3975(96)00072-2)
- [38] Mauricio Santos. 2016. Buckets-JS: A JavaScript Data Structure Library. <https://github.com/mauriciosantos/Buckets-JS>.
- [39] Micha Sharir and Amir Pnueli. 1981. Two Approaches to Interprocedural Data Flow Analysis. In *Program Flow Analysis: Theory and Applications*.
- [40] Gagandeep Singh, Markus Püschel, and Martin T. Vechev. 2017. Fast polyhedra abstract domain. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3093333.3009885>
- [41] Emma Söderberg and Görel Hedin. 2012. Incremental Evaluation of Reference Attribute Grammars using Dynamic Dependency Tracking. *LU-CS-TR:2012-249* (2012).
- [42] Johannes Späth, Lisa Nguyen Quang Do, Karim Ali, and Eric Bodden. 2016. Boomerang: Demand-Driven Flow- and Context-Sensitive Pointer Analysis for Java. In *Object-Oriented Programming (ECOOP)*. <https://doi.org/10.4230/DARTS.2.1.12>
- [43] Manu Sridharan, Denis Gopan, Lexin Shan, and Rastislav Bodik. 2005. Demand-Driven Points-To Analysis for Java. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/1094811.1094817>
- [44] Benno Stein, Bor-Yuh Evan Chang, and Manu Sridharan. 2021. DAI: Demanded Abstract Interpretation. <https://github.com/cuplv/dai>.
- [45] Benno Stein, Bor-Yuh Evan Chang, and Manu Sridharan. 2021. Demanded Abstract Interpretation (artifact). <https://doi.org/10.5281/zenodo.4663292>.
- [46] Benno Stein, Bor-Yuh Evan Chang, and Manu Sridharan. 2021. Demanded Abstract Interpretation (Extended Version). <https://doi.org/10.1145/3453483.3454044>
- [47] Tamás Szabó, Gábor Bergmann, Sebastian Erdweg, and Markus Voelter. 2018. Incrementalizing lattice-based program analyses in Datalog. In *Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. <https://doi.org/10.1145/3276509>
- [48] Tamás Szabó, Sebastian Erdweg, and Markus Voelter. 2016. IncA: a DSL for the definition of incremental program analyses. In *Automated Software Engineering (ASE)*. <https://doi.org/10.1145/2970276.2970298>
- [49] F. Kenneth Zadeck. 1984. Incremental data flow analysis in a structured program editor. In *Compiler Construction (CC)*. <https://doi.org/10.1145/502874.502888>