Demanded Abstract Interpretation



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



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



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













* i.e. safe with respect to some program analyzer



* i.e. safe with respect to some program analyzer



* i.e. safe with respect to some program analyzer







































Batch Analysis





















Abstract Interpretation



Abstract Interpretation



Abstract Interpretation


Abstract Interpretation



Abstract Interpretation



- Expressivity of abstract domains
- Robust *metatheory*: soundness, termination, etc.
- Modularity w.r.t abstractions and I solvers













A DAIG reifies the *dependency* structure of an abstract interpretation.



A DAIG reifies the *dependency* structure of an abstract interpretation.

Reference cell vertices...

- ... contain intermediate analysis results & program syntax
- ... are uniquely named and potentially empty



A DAIG reifies the *dependency* structure of an abstract interpretation.

Reference cell vertices...

- ... contain intermediate analysis results & program syntax
- ... are uniquely named and potentially empty

Computation edges...

- ... <u>acyclically</u> connect reference cells
- ... denote analysis computations (e.g. $\llbracket \cdot \rrbracket^{\sharp}, \sqcup, \nabla$)



































Demand-driven query: "What is the abstract state at *l'*?" i.e. "What is the value of cell *l'*?"





Demand-driven query: "What is the abstract state at *l'*?" i.e. "What is the value of cell *l'*?"

$$\underline{l'} = \llbracket \underline{l \to l'} \rrbracket^{\sharp} \underline{l}$$







Demand-driven query: "What is the abstract state at l'?" i.e. "What is the value of cell l'?" $\underline{l' = [[S]]^{\ddagger} \underline{l}$











Demand-driven query: "What is the abstract state at l'?" $\underline{l'} = \llbracket S \rrbracket^{\ddagger} \varphi_l$ i.e. "What is the value of cell l'?"





Demand-driven query: "What is the abstract state at *l'*?" i.e. "What is the value of cell <u>*l'*</u>?"





Program edit: Modify statement s to s'
















Subsequent queries only recompute those analysis results potentially affected by the edit!



















Demand-driven query: "What is the *fixed-point* abstract state at *l* ?"



















Demand-driven query: "What is the *fixed-point* abstract state at *l* ?"

Otherwise, unroll the DAIG's loop representation, re-query, and continue.



Statement cells not unrolled!

Otherwise, unroll the DAIG's loop representation, re-query, and continue.



Statement cells not unrolled!

Otherwise, unroll the DAIG's loop representation, re-query, and continue.



A DAIG is an *explicit* representation of a partially-evaluated abstract interpretation



Query evaluation == (small-step) operational semantics



Query evaluation == (small-step) operational semantics



Given an initial DAIG*...

* this elides some details of the actual semantics

Query evaluation == (small-step) operational semantics



$\vdash n \Rightarrow v$

Given an initial DAIG*...

... a query for the value named n yields v...

Query evaluation == (small-step) operational semantics







Given an initial DAIG*...

... a query for the value named *n* yields *v*...

... and an updated DAIG*.

Query evaluation == (small-step) operational semantics



 $\vdash n \Rightarrow v$



Given an initial DAIG*...

... a query for the value named *n* yields *v*... ... and an updated DAIG*.

Edit handling/change propagation is *also* a small-step operational semantics

Query evaluation == (small-step) operational semantics



 $\vdash n \Rightarrow v$



Given an initial DAIG*...

... a query for the value named n yields v...

... and an updated DAIG*.

Edit handling/change propagation is *also* a small-step operational semantics



Given an initial DAIG...

Query evaluation == (small-step) operational semantics







Given an initial DAIG*...

... a query for the value named *n* yields *v*... ... and an updated DAIG*.

Edit handling/change propagation is *also* a small-step operational semantics



 $\vdash n \leftarrow v$

Given an initial DAIG...

... an edit that writes value v to cell n ...

Query evaluation == (small-step) operational semantics







Given an initial DAIG*...

... a query for the value named *n* yields *v*... ... and an updated DAIG*.

Edit handling/change propagation is *also* a small-step operational semantics



Given an initial DAIG... ... an edit that writes value v to cell n ...

 $n \leftarrow v$



... yields an updated DAIG.

Termination



Termination



If this initial DAIG is a valid abstract interpretation state for some program...




If this initial DAIG is a valid abstract interpretation state for some program...

 \dots then a query for any n therein \dots





If this initial DAIG is a valid abstract interpretation state for some program...

 \dots then a query for any n therein \dots

 \dots will terminate with some value v and an updated DAIG.





If the current DAIG is a valid abstract interpretation state for some program...



If the current DAIG is a valid abstract interpretation state for some program...

... and n is the name of the abstract state at a program location l ...



If the current DAIG is a valid abstract interpretation state for some program...

... and n is the name of the abstract state at a program location l ...

... then v is *precisely* the same value that a batch abstract interpreter would compute at l.



If the current DAIG is a valid abstract interpretation state for some program...

... and n is the name of the abstract state at a program location l ...

... then v is *precisely* the same value that a batch abstract interpreter would compute at l.

No loss of precision due to incrementality/demand!



If the current DAIG is a valid abstract interpretation state for some program...

... and n is the name of the abstract state at a program location l ...

... then v is *precisely* the same value that a batch abstract interpreter would compute at l.

No loss of precision due to incrementality/demand!

Corollary: DAIG query results are **sound**.

Does a DAIG-based analysis framework support rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?

Does a DAIG-based analysis framework support rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?

Prototype Implementation: <u>github.com/cuplv/dai</u>

- \sim 2500 lines of OCaml code
- Parametric in a statement language and abstract domain

Does a DAIG-based analysis framework support rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?

Prototype Implementation: <u>github.com/cuplv/dai</u>

- \sim 2500 lines of OCaml code
- Parametric in a statement language and abstract domain

module type Domain = sig
type t
val init : t
val interpret : stmt -> t -> t
val implies : t -> t -> bool
val join : t -> t -> t
val widen : t $->$ t $->$ t
(* elided: equal, hash, etc. *)
end

Does a DAIG-based analysis framework support rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?

Prototype Implementation: github.com/cuplv/dai

- \sim 2500 lines of OCaml code
- Parametric in a statement language and abstract domain



Does a DAIG-based analysis framework support rich analysis domains that cannot be handled by existing incremental and/or demand-driven frameworks?

Prototype Implementation: <u>github.com/cuplv/dai</u>

- \sim 2500 lines of OCaml code
- Parametric in a statement language and abstract domain



Do DAIGs support rich abstract domains that cannot be handled by existing incremental and/or demanddriven frameworks?

Do DAIGs support rich abstract domains that cannot be handled by existing incremental and/or demanddriven frameworks?

Interval Analysis

- Abstract values [x, y] model integers $\{i \mid x \le i \le y\}$
- Used to verify array accesses in-bounds in a JS data structure library.

Do DAIGs support rich abstract domains that cannot be handled by existing incremental and/or demanddriven frameworks?

Interval Analysis

- Abstract values [x, y] model integers $\{i \mid x \le i \le y\}$
- Used to verify array accesses in-bounds in a JS data structure library.

Shape Analysis

- Separation logic formulae over linked-list-segment primitive $lseg(\hat{x}, \hat{y})$
- Used to verify memory safety of linked-list append, reverse, etc.

Do DAIGs support rich abstract domains that cannot be handled by existing incremental and/or demanddriven frameworks?

Interval Analysis

- Abstract values [x, y] model integers $\{i \mid x \le i \le y\}$
- Used to verify array accesses in-bounds in a JS data structure library.

Shape Analysis

- Separation logic formulae over linked-list-segment primitive $lseg(\hat{x}, \hat{y})$
- Used to verify memory safety of linked-list append, reverse, etc.

Octagon Analysis

- Invariants of the form $\pm x \pm y \leq c$
- Used in scalability experiments

Do DAIGs support rich abstract domains that cannot be handled by existing incremental and/or demanddriven frameworks?

Interval Analysis

- Abstract values [x, y] model integers $\{i \mid x \le i \le y\}$
- Used to verify array accesses in-bounds in a JS data structure library.

Shape Analysis

- Separation logic formulae over linked-list-segment primitive $lseg(\hat{x}, \hat{y})$
- Used to verify memory safety of linked-list append, reverse, etc.

Octagon Analysis

- Invariants of the form $\pm x \pm y \leq c$
- Used in scalability experiments

Intervals & Octagons built with APRON — optimized open-source numerical domains in C















The *combination* of incrementality and demand consistently obtains lower latencies than either incrementality or demand alone.



The *combination* of incrementality and demand consistently obtains lower latencies than either incrementality or demand alone.



Thanks for watching!

Check out our paper or come chat at the Q&A for more details.

Conclusion:

- Batch whole-program analysis is too costly to support real-time developer interaction, but existing incremental and demanddriven analyses are often limited in expressivity or granularity
- By leveraging graph-based incremental computation techniques, we define an engine for incremental and demand-driven evaluation of *arbitrary* abstract interpreters (and prove it sound & from-scratch consistent)





Benno Stein University of Colorado Boulder <u>benno.stein@colorado.edu</u>

PLDI '21