

Representing Data Collections in an SSA Form

Tommy M^cMichen CGO 2024

Northwestern University









Motivations How Data Collections block compilers.



Motivations How Data Collections block compilers.

Proposal Introduce MEMOIR: the first general-purpose SSA IR for Data Collections.



Motivations How Data Collections block compilers.

Proposal Introduce MEMOIR: the first general-purpose SSA IR for Data Collections.

Evaluation Demonstrate optimizations that are now possible with MEMOIR.

Data Collection: A logical organization of data.

Motivations Data Collection: **A logical organization of data**.



Motivations

Data Collection: A logical organization of data.



Motivations

Data Collection: A logical organization of data.



Motivations **Most Accesses to Heap Memory are for Collections**



Figure: Breakdown of bytes read and written for each memory class in SPECINT 2017.

Motivations Most Accesses to Heap Memory are for Structured Data



Figure: Breakdown of bytes read and written for each memory class in SPECINT 2017.

The Problem

Data collection implementations introduce complex memory behavior.

Linked data structures



Linked data structures



13

Linked data structures



Expanding data structures











```
std::unordered_map<int, int> table = ...;
table[0] = 10;
table[1] = 20;
print(table[0]);
```

std::unordered_map<int, int> table = ...;

table[0] = 10; table[1] = 20; print(table[0]);

std::unordered_map<int, int> table = ...;

table[0] = 10; table[1] = 20; print(table[0]);

No production compiler can propagate **10** to the print statement



Simple operations \rightarrow complex memory behavior.

std::unordered_map<int, int> table = ...;

Complex memory behavior blocks optimizations!

Insights **Stems from premature lowering to fixed implementations manually or via libraries**



Insights Stems from premature lowering to fixed implementations manually or via libraries



Proposal Progressively lower to MEMOIR before library implementation



Proposal Implement Memory Optimizations within the Compiler for Easy, Automatic Reuse















MEMOIR

Representing data collections for analysis and transformation

MEMOIR Goal: Provide a compiler intermediate representation with...

Unambiguous memory operations

MEMOIR Goal: Provide a compiler intermediate representation with...

Unambiguous memory operations

Element-level analyzability

MEMOIR Goal: Provide a compiler intermediate representation with...

Unambiguous memory operations

Element-level analyzability

Ability to **transform memory layout** of collections and single objects

MEMOIR Our Approach

Unambiguous memory operations

Element-level analyzability

Ability to **transform memory layout** of collections and single objects

Achieved by **decoupling** memory used to **store** data from memory used to **logically organize** data.

MEMOIR Decoupling Data from its Logical Organization

Example: Linked List



MEMOIR Decoupling Data from its Logical Organization

Example: Linked List


MEMOIR Decoupling Data from its Logical Organization

Example: Linked List



MEMOIR Decoupling Data from its Logical Organization

Example: Linked List



Abstract away the memory used to *logically organize* the collection.

MEMOIR Capture two common cases:

Associative Uniqueness in index space

MEMOIR Capture two common cases:

Associative Uniqueness in index space

Sequential Contiguous in index space

MEMOIR Element-Level Analysis

With this decoupling, we can analyze collections at the *granularity of elements*.

MEMOIR Where scalar analysis and transformation fails.

LLVM IR, etc.

Constant Scalar Propagation

... table = ...;

table[0] = 10; table[1] = 20; print(table[0]);

Element-Level analysis and transformation can prevail.

LLVM IR, etc.

















LLVM IF			Number of Collections		tions	IEMOIR
	(Benchmark	Source	SSA		it
		mcf	5	13		••• • 9
		deepsjeng	2	14		0, 10);
		LLVM opt	8	37		1, 20);

SSA Construction introduces new collections



But SSA Destruction introduces no spurious copies!

MEMOIR Performing a sparse data flow analysis on collections.

LLVM IR, etc.

Constant Scalar Propagation



Constant *Element* Propagation

Propagate *element-level* constants to optimize the program.

LLVM IR, etc.

Constant Scalar Propagation



Constant *Element* Propagation

Propagate element-level constants to optimize the program.

LLVM IR, etc.

Constant Scalar Propagation



Constant *Element* Propagation

Generalizing scalar optimizations to operate on *collections* and single *objects* with MEMOIR

LLVM IR, etc.



Generalizing scalar optimizations to operate on *collections* and single *objects* with MEMOIR

LLVM IR, etc.



Generalizing scalar optimizations to operate on *collections* and single *objects* with MEMOIR

LLVM IR, etc.



Production compilers provide negligible performance improvements on mcf_s.



Figure: Execution time of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.

Evaluation MEMOIR provides significant performance improvements with several optimizations.



Figure: Execution time of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.

Evaluation Example Application: mcf_s from SPEC2017

Quick sort accounts for ~40% of exec. time

```
Seq<T> sorted = qsort(in);
for (i = 0 to K)
    r = READ(sorted, i);
    if (r > threshold)
        use(r);
```

Evaluation Live Range Analysis propagates liveness information

```
Seq<T> sorted = qsort(in);
for (i = 0 to K)
    r = READ(sorted, i);
    if (r > threshold)
        use(r);
```

Evaluation Live Range Analysis propagates liveness information

```
Seq<T> sorted = qsort(in);
for (i = 0 to K)
    r = READ(sorted, i);
    if (r > threshold)
        use(r);
```



Dead Element Elimination converts sort \rightarrow partial sort!



Dead Element Elimination converts sort → partial sort!



With no primitive knowledge of sort!

Dead Element Elimination converts sort → partial sort!



With no primitive knowledge of sort!

See the paper for more details.



Figure: Execution time of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.



Figure: Execution time of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.

Conclusion
How can I use MEMOIR today?



Write a pass with our open source compiler
LLVM NOELLE

github.com/arcana-lab/memoir

Conclusion **How can I use MEMOIR today?**





github.com/arcana-lab/memoir



Conclusion
How can I use MEMOIR today?





github.com/arcana-lab/memoir







github.com/arcana-lab/memoir

Representing Data githe Collections in an SSA Form

Tommy M^cMichen, Nathan Greiner, Peter Zhong, Federico Sossai, Atmn Patel, Simone Campanoni





www.mcmichen.cc

Northwestern University




Motivations **Most Heap Memory is for Structured Data**



Figure: Breakdown of bytes read and written for each memory class in SPECINT 2017.

Motivations Most Reads from Heap Memory are for Structured Data



Figure: Breakdown of bytes read and written for each memory class in SPECINT 2017.

Motivations Most Writes to Heap Memory are for Structured Data



Figure: Breakdown of bytes read and written for each memory class in SPECINT 2017.

Evaluation MEMOIR requires reasonable compilation time

	Compile Time (ms)						
	MEM	IOIR	LLVM				
Benchmark	-00	-03	-00	-03			
mcf	70.6	776.4	20.9	663.2			
deepsjeng	246.0	1867.6	34.8	852.8			
LLVM opt	225.9	668.4	52.0	414.7			

Motivations Constant Folding Rarely Succeeds with Memory Operations



Figure 4: Breakdown of attempts to perform constant folding.

mcf_s Execution Time with Pass Breakdown

Lower is better.



Figure 3: Execution time of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.

mcf_s Max RSS with Pass Breakdown

Lower is better.



Figure: Maximum resident set size usage of mcf_s with refspeed input. 10 trials. Normalized to LLVM9.

MEMOIR Additional Optimizations



- - Genverts fields of objects into local associative arrays when profitable



- - Section Converts associative arrays into sequences when possible



Motivations Linked Data Structures

↓ Linked data structures are the root of undecidability for the memory aliasing problem¹



₲ Commonly used for fast insertion and deletion to lists



But this hides the **underlying index space**, blocking compiler optimizations



¹G. Ramalingam. "*The undecidability of aliasing*." ACM TOPLAS'94, Volume 16, Issue 5.

Motivations **Premature Memory Optimizations**

 Performance engineers employ manual memory optimizations at the source-level to improve memory performance and utilization



Short-lived objects are aggregated into the same, long-lived object



 Alias analysis can only glean information about accesses to *memory locations* rather than *memory objects*

MEMOIR Example C to MEMOIR Translation

```
def interp(size n, i8 in[n]) \rightarrow i32 {
  S0 = new Sequence<i32>(0)
  for (size i = 0; i < n; i++) {</pre>
    S1 = \phi(S0, S5)
    i8 tok = in[i]
    if (tok = '+') {
      i32 lhs = S0[end]
      S1 = USE(S0)
      S2 = S1[0..end-1] // pop
      i32 rhs = S2[end]
      S3 = USE(S2)
      S4 = S3[0..end-1] // pop
      i32 sum = lhs + rhs
      S5 = CONCAT(S4, [sum]) // push
    } else if (tok ≥ '0' &
               tok ≤ '9') {
      i32 val = tok - '0'
      S6 = CONCAT(S1, [val]) // push
    }
    S7 = \phi(S5, S6)
  i32 top = S7[end]
  return stack[sp];
```

MEMOIR Intermediate Representation

Data collections as *first-class citizens*

- ₲ Collections are *immutable* and *read-once*
 - Their size and elements are fixed upon creation
- New collections can be created by *adding or removing elements* from existing collections and *modifying the elements* of existing collections
- Gollections and their elements have static, strong types

MEMOIR General-Purpose Data Collections



MEMOIR General-Purpose Data Collections

Associative



MEMOIR Capture two common cases:

Linked List

Vector

Array

Hashtable Association List Bitmap

MEMOIR Capture two common cases:



MEMOIR Read-once Form

 A read-once form is constructed by creating a new collection variable after each read

 $\dots = S_0[\%i]$ $S_1 = USE\phi(S_0[\%i], S_0)$ $\dots = A_0[\%k]$ $A_1 = USE\phi(A_0[\%k], A_0)$

MEMOIR Data Flow of Collections

- ₲ Collections may only be accessed via DEF-USE data flow
 - *Cannot escape* into memory locations

MEMOIR Object Types

↓ User-defined product types are supported

struct ⊺ {	\Rightarrow	type	T = {
int x;		<i>x:</i>	i32,
float y;	<i>y</i> :	f32,	
int * <i>z</i> ;		Ζ:	ptr,
T *n;		n:	ЪТ
};		}	

- Support for primitive types
 - Raw pointers are handled by the ptr type
- Support for reference types
 - ▶ Nullable reference to a MEMOIR object of the given type
- Additional constraints for field ordering and alignment can be specified

MEMOIR Objects

◦ Objects are an instantiation of a given object type

a0 = **new** T

- ₲ Fields are laid out according to the type
 - If packing or padding are needed, it must be explicit in the type definition

MEMOIR Field Arrays

➡ Fields of an object are accessed via *field arrays*: an associative array from an object reference to the field's value

F^{T.x} = new Assoc<&T,i32>F^{T.y} = new Assoc<&T,f32>F^{T.z} = new Assoc<&T,ptr>F^{T.n} = new Assoc<&T,&T>

• Decouples the memory layout of fields within an object from their access

MEMOIR Live-Variable Analysis

• Determine if a *variable* is alive at a given program point

a	=	1			\triangleright	{	a	}	
b	=	а	+	1	$\[\] \]$	{	a,	b	}
С	=	а	*	2	\triangleright	{	b,	С	}
d	=	b	+	С	\triangleright	{	d }		

MEMOIR Live-Element Analysis

Getermine if an *element* is alive at a given program point

 $a[0] = 1 \qquad \qquad \triangleright \ \{ \ a[0] \ \}$ $a[1] = a[0] + 1 \qquad \triangleright \ \{ \ a[0], \ a[1] \ \}$ $a[2] = a[0] + 2 \qquad \triangleright \ \{ \ a[1], \ a[2] \ \}$ $a[3] = a[1] + a[2] \qquad \triangleright \ \{ \ a[3] \ \}$

MEMOIR Live-Element Analysis

• Unlike scalars, the *element index* is not necessarily known at compile time:

MEMOIR Live-Range Analysis

▶ Value range analysis informs us of elements that **may be used**:



MEMOIR Live-Range Analysis

So we can determine *live ranges* of collection index spaces:



Evaluation Example Application: mcf_s from SPEC2017

Quick sort accounts for 39.9% of exec. time



Evaluation

Live-Range Analysis reports only a portion of the result is needed.



Evaluation Dead Element Elimination converts Sort → Partial Sort



Evaluation Dead Element Elimination converts Sort → Partial Sort



mcf_s parallel speedup with DEE optimization



Cores