

Loop-Oriented Array- and Field-Sensitive Pointer Analysis for Automatic SIMD Vectorization

Yulei Sui, Xiaokang Fan, Hao Zhou and Jingling Xue

School of Computer Science and Engineering
The University of New South Wales
2052 Sydney Australia

Jun 13, 2016

Contributions

- A new **loop-oriented** array- and field-sensitive **inter-procedural** pointer analysis using access-based location sets built in terms of a **lazy memory modeling**.
- The technique improves the effectiveness of both SLP and Loop-Level Vectorization by **vectorizing more basic blocks** and **reducing runtime checks**
- Improves the performance of LLVM's SLP (best speedup of **2.95%**) and Loop vectorizer (best speedup of **7.18%**)

Outline

- Background and Motivation
- Our approach: LPA
- Evaluation

Pointer Alias Analysis and SIMD Vectorization

Pointer Analysis

- Statically approximate runtime values of a pointer.
- Serves as the foundation for compiler optimisations and software bug detection.
- Generally answers the questions, such as does two pointer expressions (e.g., *a and *b) may access the same memory.

Pointer Alias Analysis and SIMD Vectorization

Pointer Analysis

- Statically approximate runtime values of a pointer.
- Serves as the foundation for compiler optimisations and software bug detection.
- Generally answers the questions, such as does two pointer expressions (e.g., *a and *b) may access the same memory.

Pointer Alias Analysis and SIMD Vectorization

Pointer Analysis

- Statically approximate runtime values of a pointer.
- Serves as the foundation for compiler optimisations and software bug detection.
- Generally answers the questions, such as does two pointer expressions (e.g., *a and *b) may access the same memory.

Automatic SIMD Vectorization

- *Superword-Level Parallelism (SLP)* vectorization packs isomorphic scalar instructions in the same basic block into a vector instruction
- *Loop-Level Vectorization (LLV)* combines multiple consecutive iterations of a loop into a single iteration of a vector instruction.

Pointer Alias Analysis and SIMD Vectorization

Pointer Analysis

- Statically approximate runtime values of a pointer.
- Serves as the foundation for compiler optimisations and software bug detection.
- Generally answers the questions, such as does two pointer expressions (e.g., *a and *b) may access the same memory.

Automatic SIMD Vectorization

- *Superword-Level Parallelism (SLP)* vectorization packs isomorphic scalar instructions in the same basic block into a vector instruction
- *Loop-Level Vectorization (LLV)* combines multiple consecutive iterations of a loop into a single iteration of a vector instruction.

Aim of this work:

- Study and develop interprocedural pointer analysis to **generate more vectorized code (SLP) and reduce dynamic dependence checks (LLV)**.

Pointer Alias Analysis and SLP Vectorization

```
void foo(float* A, float* B){  
    A[0] = B[0];  
    A[1] = B[1];  
    A[2] = B[2];  
    A[3] = B[3];  
}
```

vectorization



```
void foo(float* A, float* B){  
    A[0:3] = B[0:3];  
}
```

SLP vectorization: pack isomorphic non-alias memory accesses

Pointer Alias Analysis and SLP Vectorization

```
void foo(float* A, float* B){  
    A[0] = B[0];  
    A[1] = B[1];  
    A[2] = B[2];  
    A[3] = B[3];  
}
```

vectorization
→

```
void foo(float* A, float* B){  
    A[0:3] = B[0:3];  
}
```

SLP vectorization: pack isomorphic non-alias memory accesses

Imprecise alias information (e.g., $A[i]$ and $B[i]$ are aliases)
miss the vectorization opportunity!

Pointer Alias Analysis and LLV Vectorization

```
void foo(float* A, float* B){  
  for(int i = 0; i < N; i++)  
    A[i] = B[i] + K;  
}  
  
vectorization →  
void foo(float* A, float* B){  
  if( (&A[N-1] >= &B[0]) && (&B[N-1] >= &A[0]))  
    for(int i = 0; i < N; i++)  
      A[i] = B[i] + K;  
  else  
    for(int i = 0; i < N; i+=4)  
      A[i:i+3] = B[i:i+3] + K;  
}
```

Loop vectorization: Dynamic checks due to imprecise aliases

Pointer Alias Analysis and LLV Vectorization

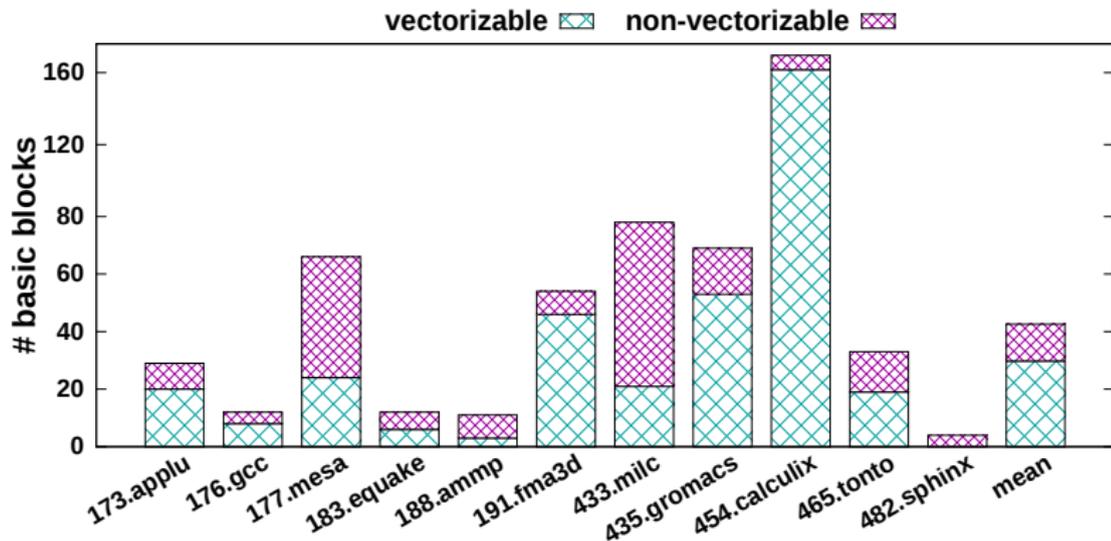
```
void foo(float* A, float* B){  
  for(int i = 0; i < N; i++)  
    A[i] = B[i] + K;  
}  
  
vectorization → void foo(float* A, float* B){  
  if( (&A[N-1] >= &B[0]) && (&B[N-1]) >= &A[0])  
    for(int i = 0; i < N; i++)  
      A[i] = B[i] + K;  
  else  
    for(int i = 0; i < N; i+=4)  
      A[i:i+3] = B[i:i+3] + K;  
}
```

Loop vectorization: Dynamic checks due to imprecise aliases

**Imprecise alias information (e.g., A[i] alias B[i])
increases the runtime overhead!**

Motivation

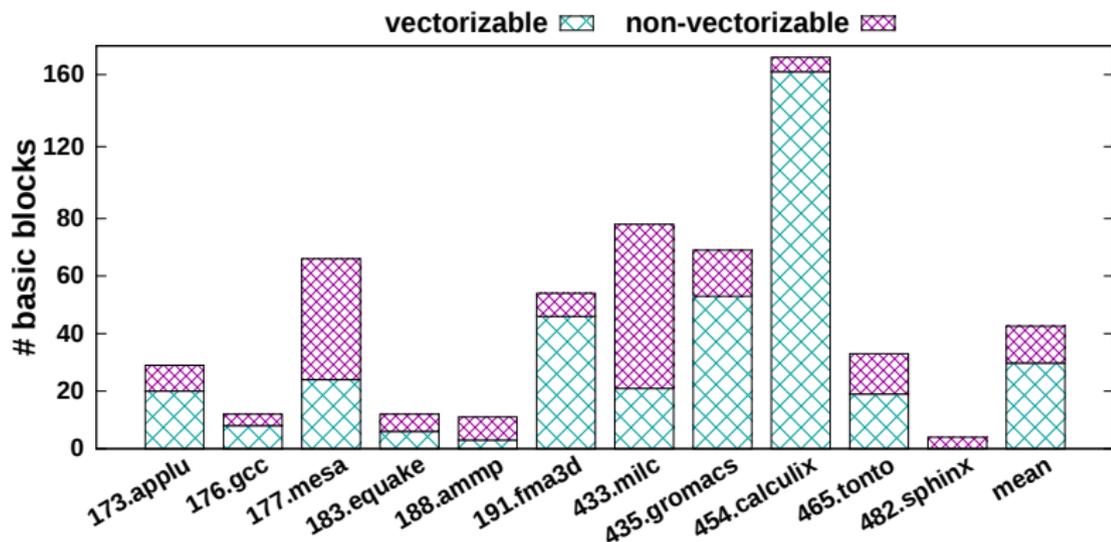
Impact of LLVM's Basic Alias Analysis on the effectiveness of SLP in LLVM



SLP: number of vectorizable and non-vectorizable basic blocks

Motivation

Impact of LLVM's Basic Alias Analysis on the effectiveness of SLP in LLVM

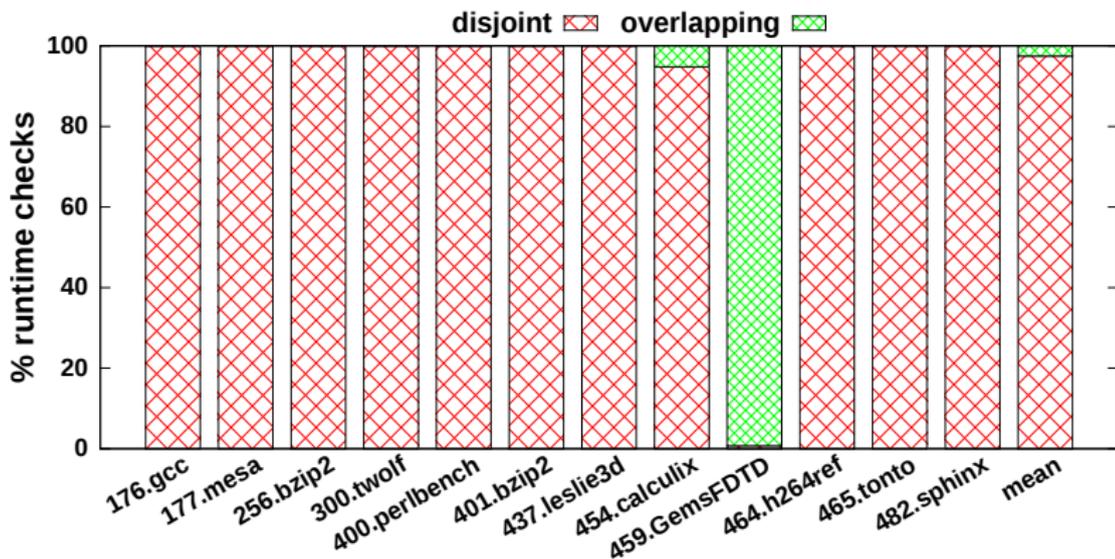


SLP: number of vectorizable and non-vectorizable basic blocks

On average, up to 30.04% of basic blocks are vectorizable
if more precise alias analysis is used in the above benchmarks!

Motivation

Impact of LLVM's Basic Alias Analysis on the effectiveness of LLV in LLVM



LLV: percentage of runtime checks for disjoint and overlapping memory

On average, up to 96.35% of dynamic alias checks

which return disjoint regions can be removed in the above benchmarks!

Precision of Pointer Alias Analysis

- Analysis dimensions (**Most previous works**):
 - flow-sensitivity
 - context-sensitivity
 - path-sensitivity
- Abstract memory modeling (**This work**)
 - Partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

Abstract Memory modeling for Pointer Analysis

Abstract memory modeling is to partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

```
struct ST{  
  int f1;  
  int f2;  
  int f3;  
}
```

```
struct ST st;  
int* p = &st.f1;  
int* q = &st.f2;
```

Field-Insensitive Modeling:



Abstract Memory modeling for Pointer Analysis

Abstract memory modeling is to partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

```
struct ST{  
  int f1;  
  int f2;  
  int f3;  
}
```

```
struct ST st;  
int* p = &st.f1;  
int* q = &st.f2;
```

Field-Insensitive Modeling:



Field-Sensitive Modeling:



Abstract Memory modeling for Pointer Analysis

Abstract memory modeling is to partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

```
struct ST{  
  int f1;  
  int f2;  
  int f3;  
}
```

```
struct ST st;  
int* p = &st.f1;  
int* q = &st.f2;
```

Field-Insensitive Modeling:



$\text{Alias}(*p, *q) = \text{true}$

Field-Sensitive Modeling:



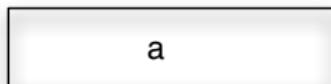
$\text{Alias}(*p, *q) = \text{false}$

Abstract Memory modeling for Pointer Analysis

Abstract memory modeling is to partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

Array-Insensitive Modeling:

```
int a[3];
```

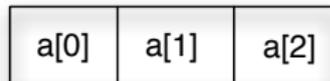


$\text{Alias}(*p, *q) = \text{true}$

Array-Sensitive Modeling:

```
int* p = &a[0];
```

```
int* q = &a[1];
```



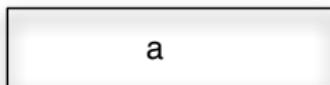
$\text{Alias}(*p, *q) = \text{false}$

Abstract Memory modeling for Pointer Analysis

Abstract memory modeling is to partition the infinite-size concrete addresses (stack/global/heap) into a finite number of abstract objects.

Array-Insensitive Modeling:

```
int a[3];
```

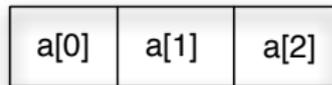


$\text{Alias}(*p, *q) = \text{true}$

Array-Sensitive Modeling:

```
int* p = &a[0];
```

```
int* q = &a[1];
```



$\text{Alias}(*p, *q) = \text{false}$

Insensitive modeling: coarse-grained (commonly used in pointer analysis)

Sensitive modeling: costly and overkill for precision

Challenges

- How to find the right balance between efficiency and precision to model abstract objects?
- How to model an array access when its index is variant including nested aggregates (e.g., array of struct, struct of array)?
- How to integrate byte-precise abstract modeling into an inter-procedural pointer analysis to improve vectorization?

LPA: Loop-oriented Pointer Analysis

- **Loop-oriented** array- and field-sensitive inter-procedural pointer analysis using **access-based location sets** built in terms of lazy memory modeling.
 - *Statically evaluate the symbolic range of pointers according to loop information.*
 - *Generate location sets lazily during points-to resolution.*

LPA: Loop-oriented Pointer Analysis

- **Loop-oriented** array- and field-sensitive inter-procedural pointer analysis using **access-based location sets** built in terms of lazy memory modeling.
 - *Statically evaluate the symbolic range of pointers according to loop information.*
 - *Generate location sets lazily during points-to resolution.*
- Separates **memory modeling as an independent concern** from the rest of the pointer analysis.
 - *Facilitating the development of pointer analyses with desired efficiency and precision tradeoffs by reusing existing pointer resolution algorithms.*

LPA: Loop-oriented Pointer Analysis

- **Loop-oriented** array- and field-sensitive inter-procedural pointer analysis using **access-based location sets** built in terms of lazy memory modeling.
 - *Statically evaluate the symbolic range of pointers according to loop information.*
 - *Generate location sets lazily during points-to resolution.*
- Separates **memory modeling as an independent concern** from the rest of the pointer analysis.
 - *Facilitating the development of pointer analyses with desired efficiency and precision tradeoffs by reusing existing pointer resolution algorithms.*
- Generate efficient vector code and **improves the performance** of both SLP and loop vectorizer (best speedup over 7%).

Access-based location set

A location set σ represents memory locations in terms of numeric offsets from the beginning of an abstract memory block¹.

Our **array-sensitive modeling**, e.g., $arr[i]$ inside a loop:

- **Interval range** $i \in [lb, ub]$
- **Access step** $X \in \mathbb{N}^+$ (e.g., $X = 1$ if arr is accessed consecutively inside the loop)

¹R. P. Wilson and M. S. Lam. Efficient context-sensitive pointer analysis for C programs. In PLDI '95
(*Field-Sensitive array-insensitive modeling based on location set*)

Access-based location set

A location set σ represents memory locations in terms of numeric offsets from the beginning of an abstract memory block¹.

Our **array-sensitive modeling**, e.g., $arr[i]$ inside a loop:

- **Interval range** $i \in [lb, ub]$
- **Access step** $X \in \mathbb{N}^+$ (e.g., $X = 1$ if arr is accessed consecutively inside the loop)
- **Access trip** is a pair (t, s) consists of
 - A trip count $t = \frac{(ub - lb)}{(X - 1)}$
 - A stride $s = es * X$ where es is the size of an array element.

¹R. P. Wilson and M. S. Lam. Efficient context-sensitive pointer analysis for C programs. In PLDI '95
(*Field-Sensitive array-insensitive modeling based on location set*)

Access-based location set

An access-based location set derived from an object a is:

$$\sigma = \langle \text{off}, \llbracket (t_1, s_1), \dots, (t_m, s_m) \rrbracket \rangle_a$$

where $\text{off} \in \mathbb{N}$ is an offset from the beginning of object a , and $T = \llbracket (t_1, s_1), \dots, (t_m, s_m) \rrbracket$ is an access-trip stack containing a sequence of (trip count, stride) pairs for handling a nested struct of arrays.

Access-based Location Set (Examples)

```
1 float a[16]; float *p = &a[0];
2 for(i=0;i<8;i++)                i∈[0, 7], X=1, es=4
3   p[i] = p[i+8];                i+8∈[8, 15], X=1, es=4
```

0 4 8 12 16 20 24 28 32 36 40 44 48 52 56 60 64

■ $p[i]: \langle 0, [(8, 4)] \rangle_a$ ■ $p[i+8]: \langle 32, [(8, 4)] \rangle_a$

(a) An array with **consecutive** accesses

Access-based Location Set (Examples)

```
1 float a[16]; float *p = &a[0];
2 for(i=0;i<16;i+=4){
3   p[i]   = i ;           i∈[0, 12], X=4, es=4
4   p[i+1] = i + 1;       i+1∈[1, 13], X=4, es=4
5   p[i+2] = i + 2;       i+2∈[2, 14], X=4, es=4
6   p[i+3] = i + 3;       i+3∈[3, 15], X=4, es=4
7 }
```

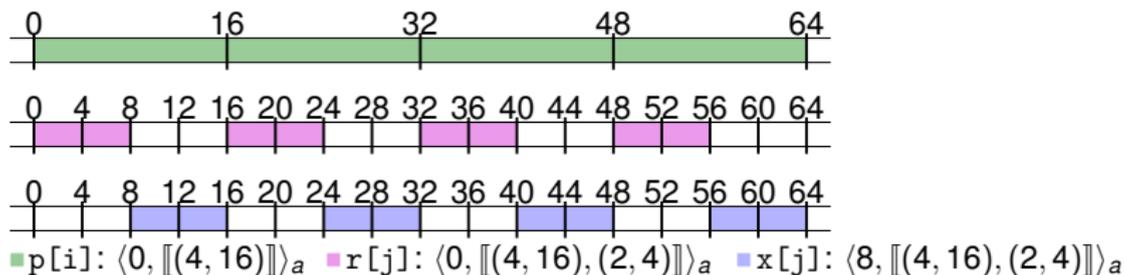
0 4 8 12 16 20 24 28 32 36 40 44 48 52 56 60 64

■ p[i]: $\langle 0, \llbracket (4, 16) \rrbracket \rangle_a$ ■ p[i+1]: $\langle 4, \llbracket (4, 16) \rrbracket \rangle_a$
■ p[i+2]: $\langle 8, \llbracket (4, 16) \rrbracket \rangle_a$ ■ p[i+3]: $\langle 12, \llbracket (4, 16) \rrbracket \rangle_a$

(b) An array with **non-consecutive** accesses

Access-based Location Set (Examples)

```
1 struct {float f1[2]; float f2[2];} a[4], *p, *q;
2 float *r;
3 p = &a[0];
4 for(i=0;i<4;i++){
5     q = &p[i];                                 $i \in [0, 3], X=1, es=16$ 
6     r = q->f1;
7     x = q->f2;
8     for(j=0;j<2;j++){
9         r[j] = x[j];                             $j \in [0, 1], X=1, es=4$ 
10 }
```



(c) **Nested arrays and structs** with consecutive accesses

Andersen's Pointer Analysis based on Access-based Location Set

$$[S\text{-ALLOC}] \frac{p = \&a \quad \sigma = \langle 0, \square \rangle_a}{\{\sigma\} \subseteq \text{pt}(p)}$$

$$[S\text{-LOAD}] \frac{q = e_p \quad \sigma \in \text{pt}(q) \quad \sigma' = \text{GetLS}(\sigma, e_q)}{\text{pt}(\sigma') \subseteq \text{pt}(q)}$$

$$[S\text{-COPY}] \frac{p = q}{\text{pt}(q) \subseteq \text{pt}(p)}$$

$$[S\text{-STORE}] \frac{e_p = q \quad \sigma \in \text{pt}(p) \quad \sigma' = \text{GetLS}(\sigma, e_p)}{\text{pt}(q) \subseteq \text{pt}(\sigma')}$$

$$\text{GetLS}(\langle \text{off}, T \rangle_a, e_p) = \begin{cases} \langle \text{off}, T \rangle_a & \text{if } e_p \text{ is } *p \\ \langle \text{off} + \text{off}_f, T \rangle_a & \text{else if } e_p \text{ is } p \rightarrow f, \text{ where } \text{off}_f \text{ is the offset of} \\ & \text{field } f \text{ in array object } a \\ \langle \text{off} + C * \text{es}, T \rangle_a & \text{else if } e_p \text{ is } p[i], \text{ where } i \text{ is constant } C \\ \langle \text{off} + lb * \text{es}, T.\text{push}(\frac{ub' - lb'}{X} + 1, X * \text{es}) \rangle_a & \text{else if } e_p \text{ is } p[i], \text{ where } i \in [lb, ub] \text{ with step } X, \\ & [lb', ub'] = [lb, ub] \cap [0, m - 1] \\ & \text{and } m \text{ is size of array object } a \end{cases}$$

Points-to target in the points-to set of a pointer is not an abstract object but rather a location set derived from it.

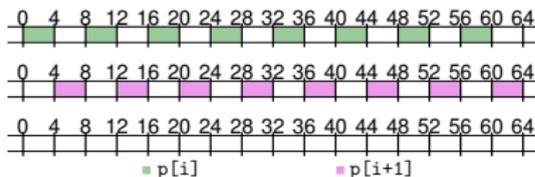
Disambiguation of location sets

$$\text{alias}(e_p, e_q) = \begin{cases} \text{true} & \text{if } \exists \sigma'_p \in \text{pt}(p) \wedge \sigma'_q \in \text{pt}(q) : (\sigma_p, \text{sz}_p) \bowtie (\sigma_q, \text{sz}_q), \\ & \text{where } \sigma_p = \text{GetLS}(\sigma'_p) \wedge \sigma_q = \text{GetLS}(\sigma'_q) \\ \text{false} & \text{otherwise} \end{cases} \quad (1)$$

$$(\sigma_p, \text{sz}_p) \bowtie (\sigma_q, \text{sz}_q) = \begin{cases} \text{true} & \text{if } \text{obj}(\sigma_p) = \text{obj}(\sigma_q) \text{ and} \\ & \exists l_p \in \text{LS}(\sigma_p) \wedge l_q \in \text{LS}(\sigma_q) : (l_p < l_q + \text{sz}_q) \wedge (l_q < l_p + \text{sz}_p) \\ \text{false} & \text{otherwise} \end{cases} \quad (2)$$

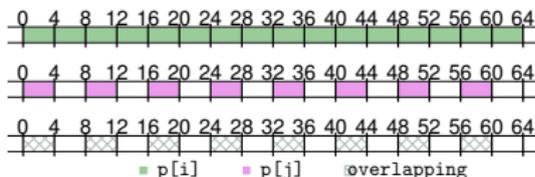
Disambiguation of location sets

```
1 float a[16]; float *p = &a[0];
2 for(i=0;i<16;i+=2){
3   p[i] = ...;
4   p[i+1] = ...;          ((0, [(8, 8)])a, 4) ∩ ((4, [(8, 8)])a, 4)
5 }
```



(a) Disjoint location sets

```
1 float a[16]; float *p = &a[0];
2 for(i=0;i<16;i++)
3   p[i] = ...;
4 for(j=0;j<16;j+=2)          ((0, [(16, 4)])a, 4) ∩ ((0, [(8, 8)])a, 4)
5   p[j] = ...;
```

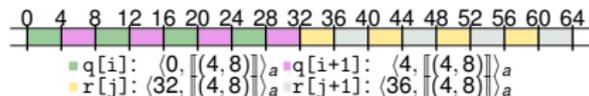


(b) Overlapping location sets

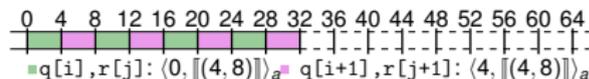
Field Unification

```
1 struct {float f1[8]; float f2[8];} a, *p;
2 p = &a;
3 float *q = p->f1, *r = p->f2;
4 for(int i=0;i<8;i=i+2){
5     q[i] = ...;
6     q[i+1] = ...;
7 }
8 for(int j=0;j<8;j=j+2){
9     r[j] = ...;
10    r[j+1] = ...;
11 }
```

(a) Code



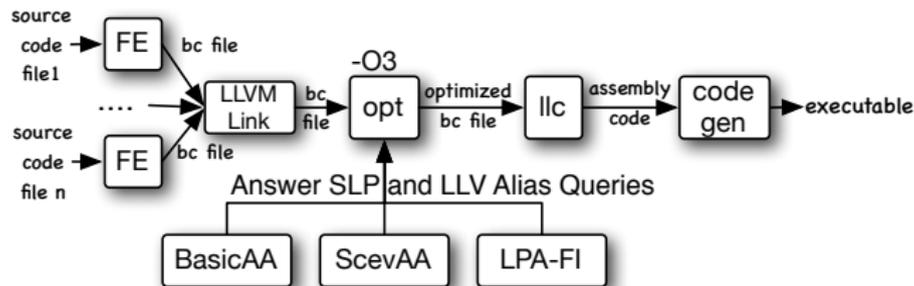
(b) Default location sets



(c) Location sets with **max offset limit: F = 32**

Experiments

Compilation Process



Our experiments are conducted on

- An Intel Core i7-4770 CPU (3.40GHz) with an AVX2 SIMD extension, which supports 256 bit floating point and integer SIMD operations.
- 64-bit Ubuntu (14.0.4) with 32 GB memory.

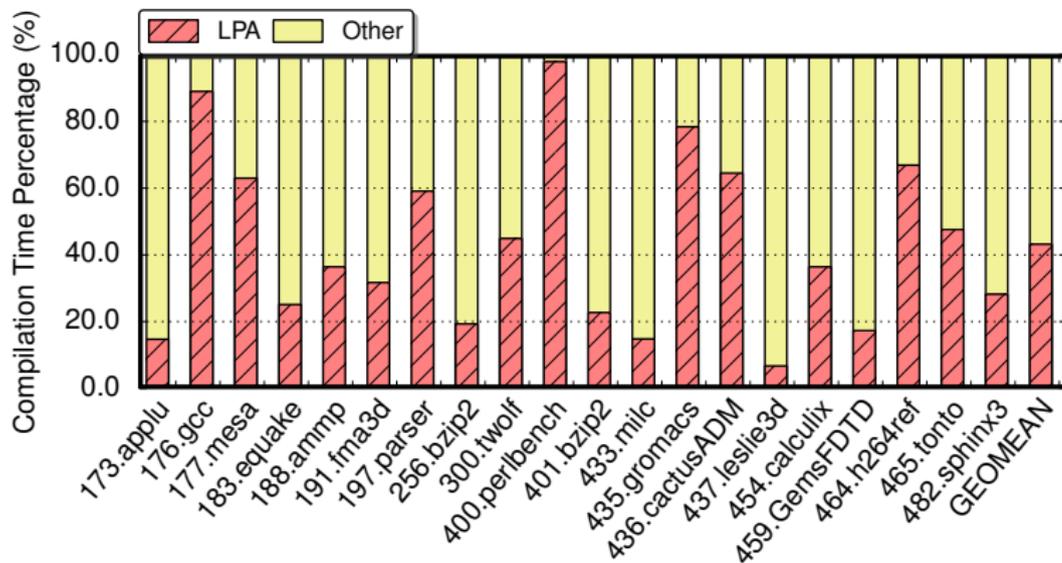
Experiments

Program Characteristics

Program	KLOC	#Stmt	#Ptrs	#Objs	#CallSite
173.applu	3.9	3361	20951	159	346
176.gcc	226.5	215312	545962	16860	22595
177.mesa	61.3	99154	242317	9831	3641
183.quake	1.5	2082	6688	236	235
188.amp	13.4	14665	56992	2216	1225
191.fma3d	60.1	119914	276301	6497	18713
197.parser	11.3	13668	36864	1177	1776
256.bzip2	4.6	1556	10650	436	380
300.twolf	20.4	23354	75507	1845	2059
400.perlbench	168.1	130640	296288	3398	15399
401.bzip2	8.2	7493	28965	669	439
433.milc	15	11219	30373	1871	1661
435.gromacs	108.5	84966	224967	12302	8690
436.cactusADM	103.8	62106	188284	2980	8006
437.leslie3d	3.8	12228	38850	513	2003
454.calculix	166.7	135182	532836	18814	23520
459.GemsFDTD	11.5	25681	107656	3136	6566
464.h264ref	51.5	55548	184660	3747	3553
465.tonto	143.1	418494	932795	28704	58756
482.sphinx3	25	20918	60347	1917	2775
Total	1208.2	1457541	3898253	117308	182338

Analysis Time

Percentage of analysis time over total compilation time



LPA's analysis times ranging from 94.4 secs to 240.8 secs

On average, LPA's analysis time occupies 42% over the total compilation time.

Experiments

SLP Vectorization Static Statistics

Benchmark	BASICAA	SCEVAA	LPA
173.applu	20	4	26
176.gcc	4	3	6
177.mesa	24	23	64
183.quake	2	1	4
188.amp	1	2	4
191.fma3d	46	23	53
433.milc	21	13	69
435.gromacs	53	35	57
454.calculix	161	92	166
465.tonto	19	21	32
482.sphinx	0	0	1
Total	351	217	482

Number of basic blocks vectorized by SLP under the three alias analyses (larger is better).

Experiments

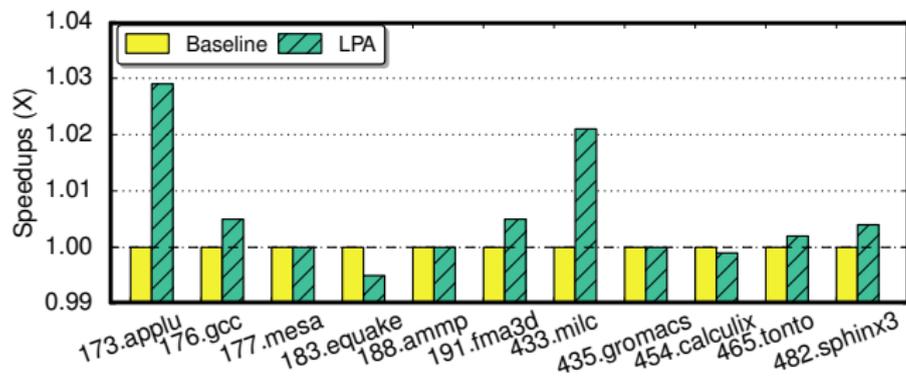
Loop Vectorization Static Statistics

Benchmark	BASICAA	SCEVAA	LPA
176.gcc	4	8	2
177.mesa	121	137	88
197.parser	1	1	0
256.bzip2	1	6	0
300.twolf	11	13	10
400.perlbench	23	21	13
401.bzip2	6	9	5
436.cactusADM	71	112	2
437.leslie3d	21	21	4
454.calculix	83	90	57
459.GemsFDTD	65	79	16
464.h264ref	30	32	2
465.tonto	110	118	38
482.sphinx3	4	5	1
Total	551	652	238

Number of static alias checks inserted by LLV under the three alias analyses (smaller is better).

Experiments

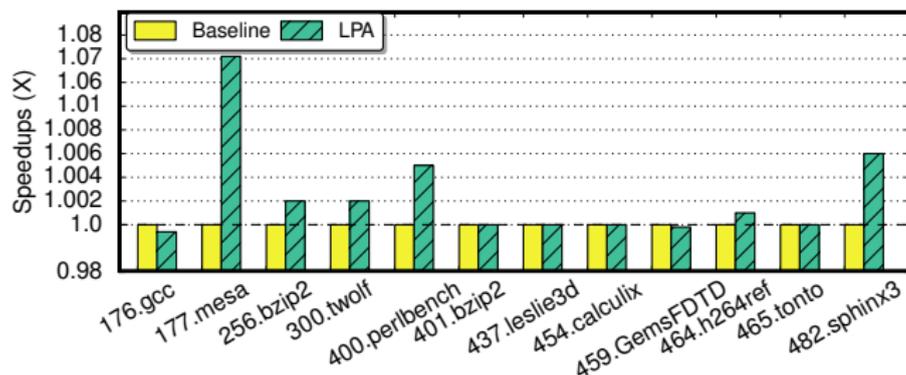
SLP: whole-program performance speedups



The whole-program speedups achieved by SLP under LPA normalized with respect to LLVM's alias analyses

Experiments

Loop Vectorization: whole-program performance speedups



The whole-program speedups achieved by LLV under LPA normalized with respect to LLVM's alias analyses

Conclusion

- A new **loop-oriented** array- and field-sensitive **inter-procedural** pointer analysis using access-based location sets built in terms of **lazy memory model**.
- The technique improves the effectiveness of both SLP and Loop-Level Vectorization by **vectorizing more basic blocks** and **reducing dynamic checks**
- Improves the performance of LLVM's SLP (best speedup of **2.95%**) and Loop vectorizer (best speedup over **7%**)

Thanks!

Q & A

Limitations

- Range analysis
 - SCEV in LLVM
 - Indirect array access e.g., $a[*p]$
 - Irregular loops, e.g., iterating arrays inside a loop with variant bounds
- More precise analysis methods, e.g, context-,heap-sensitivity

Andersen's Analysis based on Field-Insensitive Modeling

$$\frac{p = \&a}{\{a\} \subseteq \text{pt}(p)}$$

$$\frac{p = q}{\text{pt}(q) \subseteq \text{pt}(p)}$$

$$\frac{*p = q \quad o \in \text{pt}(p)}{\text{pt}(q) \subseteq \text{pt}(o)}$$

$$\frac{p = *q \quad o \in \text{pt}(q)}{\text{pt}(o) \subseteq \text{pt}(p)}$$

Every allocation site is treated as a single memory object. Array and field accesses like $p[j] = ..$ and $p \rightarrow f = ...$ are treated as copies.