

Inferring Method Effect Summaries for Nested Heap Regions

Mohsen Vakilian, Danny Dig, Robert Bocchino, Jeffrey Overbey, Vikram Adve, Ralph Johnson
University of Illinois at Urbana-Champaign
{mvakili2,dig,bocchino,overbey2,vadve,rjohnson}@illinois.edu

Abstract

Effect systems are important for reasoning about the side effects of a program. Although effect systems have been around for decades, they have not been widely adopted in practice because of the large number of annotations that they require. A tool that infers effects automatically can make effect systems practical. We present an effect inference algorithm and an Eclipse plugin, DPJIZER, that alleviate the burden of writing effect annotations for a language called Deterministic Parallel Java (DPJ). The key novel feature of the algorithm is the ability to infer effects on nested heap regions. Besides DPJ, we also illustrate how the algorithm can be used for a different effect system based on object ownership. Our experience shows that DPJIZER is both useful and effective: inferring effects annotations automatically saves significant programming burden; and inferred effects are comparable to those in manually-annotated programs, while in many cases they are more accurate.

1. Introduction

Programs written in mainstream, imperative languages have side effects on program’s memory. Programmers have embraced this paradigm because it avoids passing program’s state between different functions. However, this paradigm also makes it harder for programmers or tools to understand or analyze programs in a modular fashion.

Knowing what parts of the program’s state are mutated by a function can help programmers modify large programs without introducing subtle mutation errors and can serve as explicit, machine-checkable documentation. It can enable safety tools to detect inconsistencies between intended usage of API methods and their actual usage, it is a building block for several other compiler analyses (e.g., MODREF analysis), and it can enable compilers to check the safety of parallel programs [1]–[3].

Although effect systems have been around for decades, they have not been used much in the software engineering

practice. The reason is that manually writing such effects is tedious and error-prone. In this paper we present an algorithm that automatically infers the effects of each program statement and summarizes them at the level of method declarations as *method effect summaries*. The key novel capability in the algorithm is that it is able to handle effect inference for programs even on “nested heap structures,” including recursive as well as non-recursive data structures.

We have used the algorithm to develop an effect inference tool for a previously developed extension to Java, called *Deterministic Parallel Java (DPJ)* [4], [5], which aims to enable programmers to write safe parallel programs. The effect system in DPJ is based on “regions,” which are partitions of the heap, down to the granularity of individual fields. We also illustrate briefly how the algorithm could be used for a very different class of effect systems based on “object ownership” [6].

DPJ is an explicitly parallel language that uses a region and effect system to guarantee that any program that type-checks will have deterministic visible behavior, i.e., the program will produce identical externally visible results in all executions for a given input. Such deterministic semantics can greatly simplify parallel program design, debugging, testing and maintenance because it allows programs to reason about programs with a simple, sequential semantic model and debug and test the program using versions of familiar sequential tools. Using DPJ, we have safely parallelized several programs [5]; the parallel programs are deterministic and they exhibit good speedup. However, to get these benefits the programmer has to write region and effect annotations by hand. This job is non trivial, error-prone, and time consuming. For example, a Monte Carlo financial application contains 2877 LOC, 29 region annotations, and 24 effect annotations. A Barnes-Hut N-body application contains 682 LOC, 91 region annotations, and 46 effect annotations.

DPJIZER alleviates the programmer’s burden when writing effect annotations. Given a program with region annotations, DPJIZER infers the method effect summaries and annotates the program. We implemented DPJIZER as an

extension to Eclipse’s refactoring engine, thus it offers all the convenient features of a practical refactoring engine: previewing changes, selection of edits to be applied, undo/redo, etc.

At the heart of `DPJIZER` lies an algorithm that statically infers side effects on field regions. The input of the algorithm is a program where shared fields are annotated with region information. The algorithm infers for each method the *method effect summary* that covers the effects (reads/writes) on declared regions. When summarizing the effect information, `DPJIZER` eliminates redundant effects, which makes the effect annotations concise and easier to understand.

The inference algorithm is built on a classical constraint-based type-inference approach, but we use it to infer effects. The algorithm generates constraints from primitive operations (variable access, assignment, method calls, and method overriding declarations), using the appropriate parameter and type substitutions at method invocations. It then solves these constraints by processing them iteratively and propagating the constraints through the call graph until a fixed point is reached and no more constraints are discovered. The novelty in the algorithm lies in the constraint solving phase. This phase handles nested regions by taking advantage of the structure of region specifications in the target language (e.g., Region Path Lists [5] in DPJ or object “levels” in the object ownership system, JOE [6]). It handles *recursive* structures by summarizing these nested structures in each case.

Although `DPJIZER` is designed to help in porting a Java program to DPJ, its applicability goes well beyond DPJ. Given a concurrent program that uses shared memory, by inferring the method effects, `DPJIZER` helps a programmer discover the patterns of shared data. This information is crucial in helping the programmer find out the accesses to shared data that need to be protected. Moreover, the underlying algorithm is useful beyond concurrent programs. For example, as noted earlier, we show how the algorithm can be used to infer effects for a different effect specification system based on object ownership, which is a general mechanism to reason about and express the side effects of object-oriented programs.

This paper makes the following contributions:

1. Algorithm. To the best of our knowledge, this paper presents the first algorithm for inferring method effect summaries for a full Object-Oriented language (e.g., aliasing, recursion, polymorphism, generics, arrays, etc.) with a sophisticated effect system (e.g., parameterized regions, nested regions for recursive data-structures, etc.).

2. Tool. We implemented the effect inference algorithm in an *interactive* tool called `DPJIZER`. A programmer can use `DPJIZER` to infer method effects for a Java or a DPJ program. `DPJIZER` writes the inferred effects into the source code as DPJ annotations or as code comments.

`DPJIZER` is built as an Eclipse plugin that extends Eclipse’s refactoring engine.

3. Evaluation. We used `DPJIZER` to infer method effects in several real programs. We compare the effects inferred with `DPJIZER` against effects manually inferred by programmers. The comparison shows that `DPJIZER` can drastically reduce the burden of writing annotations manually, while the automatically inferred effects are more accurate.

2. Overview of DPJ

In this section we briefly introduce the Deterministic Parallel Java (DPJ) language [5]. We say that two tasks are *noninterfering* if for each pair of memory accesses, one from each task, either both accesses are reads, or the two accesses operate on disjoint sets of memory locations.¹ Noninterfering tasks can be run in parallel while still exhibiting the same behavior as if they were run sequentially.

DPJ provides a type system that *guarantees* noninterference of parallel tasks for a well-typed program. In DPJ, the programmer assigns every object field and array cell to a *region* of memory and annotates each method with a summary (called a *method effect summary*) of the method read and write effects. The programmer also marks which code sections to run in parallel, using several standard constructs, such as `cobegin` for parallel statement execution and `foreach` for parallel loops. The compiler uses the region annotations and method effect summaries to check that all pairs of parallel tasks are noninterfering.

2.1. Region Names

Figure 1 illustrates the use of region names to distinguish writes to different fields of an object. Line 2 declares `Mass` and `Force` as region names that are available within the scope of class `Node`. These are called *field region declarations*. Lines 3 and 4 declare fields `mass` and `force` and place them in regions `Mass` and `Force`, respectively. Field region declarations are static, so there is one for each class. For example, all `mass` fields of all `Node` instances are in the same region, `Mass`.

Each method must have a *method effect summary* recording the effects that it performs on the heap, in terms of reads and writes to regions. For example, method `setMass` (line 5) has the summary `writes Mass`, because the effect of line 6 is to write the field `mass`, located in region `Mass`; and similarly for `setForce` (line 9). It is permissible for a method effect summary to be overly conservative; for example, `setMass` could have

1. The full DPJ language [5], [7] also allows commutativity annotations that specify noninterference directly, without checking reads and writes. Here we focus on inferring read and write effects.

```

1 class Node {
2     region Mass, Force;
3     double mass in Mass;
4     double force in Force;
5     void setMass(double mass) writes Mass {
6         /* writes Mass */
7         this.mass = mass;
8     }
9     void setForce(double force) writes Force {
10        /* writes Force */
11        this.force = force;
12    }
13    void initialize(double mass, double force)
14        writes Mass, Force {
15        cobegin {
16            /* writes Mass */
17            this.setMass(mass);
18            /* writes Force */
19            this.setForce(force);
20        }
21    }
22 }

```

Figure 1. Using field region names to distinguish writes to different object fields. In Section 3, we will show how to infer the underlined method effect summaries.

said `writes Mass, Force`. However, this may inhibit parallelism. It is an error for a method effect summary to be not conservative enough, for example if `setMass` had said `pure`, meaning “no effect.”

Together, the DPJ annotations allow the compiler to efficiently analyze noninterference of parallel code sections, as illustrated in the `initialize` method. From the method effect summaries, the compiler can infer that the effect of line 17 is `writes Mass` and the effect of line 19 is `writes Force`. The compiler can then use the distinctness of the names `Mass` and `Force` to prove noninterference: although both statements in the `cobegin` perform writes, the writes are to disjoint regions of the heap.

2.2. Region Parameters

As shown in section 2.1, region names are useful for distinguishing parts of an object from each other. Often, however, we need to distinguish different *object instances* from each other. To do this, DPJ uses *region parameters*, which operate similarly to Java generic parameters [8] and allow us to instantiate different object instances of the same class with different regions.

Figure 2 illustrates the use of region parameters to distinguish writes to different object instances. In line 1, we declare class `Node` to have one region parameter `P` (we use the keyword `region` to distinguish DPJ region parameters from Java type parameters). As with Java generics, when we write a type using a class with region parameters, we provide an argument to the parameter, as shown in lines 4 and 5. The argument must be a valid region name in scope.

```

1 class Node<region P> {
2     region L, R;
3     double mass in P;
4     Node<L> left in L;
5     Node<R> right in R;
6     void setMass(double mass) writes P {
7         /* writes P */
8         this.mass = mass;
9     }
10    void setMassOfChildren(double mass) writes L, R {
11        cobegin {
12            /* writes L */
13            if (left != null) left.setMass(mass);
14            /* writes R */
15            if (right != null) right.setMass(mass);
16        }
17    }
18 }

```

Figure 2. Using region parameters to distinguish writes to different object instances.

We can use the region name `P` within the scope of the class. For example, line 3 declares field `mass` in region `P`. When `this.mass` is accessed, the effect is on region `P`, as shown in line 8. However, when we access `mass` through a selector other than `this`, we resolve the region `P` by substituting the actual argument given in the type of the selector. For example, the effect of `left.setMass` in line 13 is `writes L`. We get this by looking at the declaration `writes P` of `setMass` and substituting `L` for `P` from the type of `left`. (The read of field `left` also generates a read effect on region `L`; but in DPJ, write effects imply read effects, so the read is covered by `writes L`.) We can then use an analysis similar to the one discussed in Section 2.1 to prove that the statements in lines 13 and 15 are noninterfering, since their write effects are on the disjoint regions `L` and `R`.

2.3. Region Path Lists (RPLs)

In conjunction with array index regions and index-parameterized arrays (discussed further in Section 3.4.1), basic region names and region parameters can be used to express important parallel algorithms. However, it is often essential to be able to express a *nesting* relationship between regions. For example, to express tree-like recursive updates we need a nested hierarchy of regions. DPJ provides two ways to express nesting: *region path lists* and *owner regions*. Here we focus on region path lists; we defer the discussion of owner regions until after we have presented the effect inference algorithm.

A region path list (RPL) extends the idea of a simple region name introduced in Section 2.1. An RPL is a colon-separated list of names that expresses the nesting relationship syntactically: if `P` and `R` are names, then `P:R` is nested under `P`. Nested RPLs are particularly useful in conjunction with region parameters: if we append names to parameters, such as `P:L` and `P:R`, then by recursive

substitution we can generate arbitrarily long chains of names, such as $P:L:L:R$.

```

1 class Node<region P> {
2   region L,R;
3   double mass in P;
4   Node<P:L> left in P:L;
5   Node<P:R> right in P:R;
6   void setMassForTree(double mass) writes P:* {
7     /* writes P */
8     this.mass = mass;
9     cobegin {
10      /* writes P:L:* */
11      if (left != null) left.setMassForTree(mass);
12      /* writes P:R:* */
13      if (right !=null) right.setMassForTree(mass);
14    }
15  }
16 }

```

Figure 3. Using RPLs and region parameters to recursively update a tree in parallel.

Figure 3 illustrates the use of this technique to write a recursive tree update. The example is similar to the one shown in Figure 2, except that lines 4 and 5 use regions $P:L$ and $P:R$ instead of L and R , and the method invocations in lines 11 and 13 are recursive. To write the method effect summary in line 6, we need some new syntax: because the tree can be arbitrarily deep, and the RPLs arbitrarily long, we use a star (*) to stand in for any sequence of RPL elements. Then we can write the method effect summary `writes P:*`, as shown in line 6. Note that the rules discussed in Section 2.1 for method effect summaries are still followed: by substituting the RPL arguments in the types of `left` and `right` in for P , we get the inferred effects shown in lines 10 and 12; and those effects are covered by the summary. Further, because the RPLs form a tree, we can conclude that all regions under $P:L$ and all regions under $P:R$ are disjoint, so lines 11 and 13 are noninterfering.

3. Effect Inference Algorithm

We present our algorithm using *Core DPJ* [5], a small skeleton language that illustrates the ideas yet is tractable to formalize. To make the presentation easier to follow, we start with a simplified form of Core DPJ corresponding to the features introduced in Section 1, i.e., basic region names with no region parameters or nesting. Then we build up the language to add region parameters and region path lists. We also discuss how to handle owner regions, array regions, and inheritance. Finally, we discuss how to adapt the algorithm for use with other languages.

3.1. Basic region names

We start by showing how to infer effects for Core DPJ with basic region names, i.e., with no region parameters

or nested regions. Figure 4 shows the syntax of the initial language. The algorithm consists of two phases, constraint generation and constraint solving.

Meaning	Symbol	Definition
Programs	<i>program</i>	<i>region-decl</i> * <i>class</i> *
Region decls	<i>region-decl</i>	<i>region r</i>
Classes	<i>class</i>	<i>class C {field* μ*}</i>
Fields	<i>field</i>	<i>T f in r</i>
Types	<i>T</i>	<i>C</i>
Methods	μ	<i>T m(T x) { e }</i>
Expressions	<i>e</i>	<i>this.f this.f = e e.m(e) new T z</i>
Variables	<i>z</i>	<i>this x</i>

Figure 4. Syntax of Core DPJ with basic region names. C , f , m , x and r are identifiers.

3.1.1. Constraint generation. The goal of the constraint generation phase is to associate with each method μ a *constraint set* K_μ , where each element of K_μ is one of the constraints *reads r*, *writes r*, and *invokes μ'* . The first two constraints indicate the presence of a read or write effect in the method itself. The *invokes* constraint asserts that one method is invoking another, either directly or indirectly; these constraints will cause the solving phase (Section 3.1.2) to account for the read and write effects of callees.

The constraint generation phase visits each method body and walks the AST to generate a set of constraints. Figure 5 shows the constraint generation rules for the simplified language. The rules are similar to the ones for typing DPJ expressions [5], except that we do not check assignments or method formal parameter bindings for soundness (we assume that full DPJ type checking has been done as a separate pass).

$$\begin{array}{l}
 \text{(FIELD-ACCESS)} \quad \frac{(\text{this}, C) \in \Gamma \quad \text{field}(C, f) = T f \text{ in } r}{\Gamma \vdash \text{this}.f : T, \{\text{reads } r\}} \\
 \text{(ASSIGN)} \quad \frac{(\text{this}, C) \in \Gamma \quad \Gamma \vdash e : T, K \quad \text{field}(C, f) = T' f \text{ in } r}{\Gamma \vdash \text{this}.f = e : T, K \cup \{\text{writes } r\}} \\
 \text{(INVOKE)} \quad \frac{\Gamma \vdash e_1 : C, K_1 \quad \Gamma \vdash e_2 : T, K_2 \quad \text{method}(C, m) = \mu}{\Gamma \vdash e_1.m(e_2) : T_r, K_1 \cup K_2 \cup \{\text{invokes } \mu\}} \\
 \text{(NEW)} \quad \frac{}{\text{new } C : C, \emptyset} \quad \text{(VARIABLE)} \quad \frac{(z, T) \in \Gamma}{z : T, \emptyset}
 \end{array}$$

Figure 5. Rules for computing the constraints generated by an expression.

The judgment $\Gamma \vdash e : T, K$ means that expression e has type T and generates constraint set K in environment Γ . The environment Γ is a set of pairs (z, T) binding variable z to type T . The term *method*(C, m) means the method named m defined in class C , while *field*(C, f) means the field named f defined in class C . For each method $\mu = T_r m(T_x x) \{ e \}$, let C_μ be the class where μ is

defined. Then K_μ is just the set of constraints such that

$$\{(\text{this}, C_\mu), (x, T_x)\} \vdash e : T, K_\mu.$$

As an example, we show how to generate the constraints for the bodies of methods `setMass`, `setForce` and `initialize` in Figure 1. In line 7, rule `FIELD-ACCESS` generates the constraint $\text{reads}_{\text{Mass}}$ for reading the right-hand side of the assignment. Similarly, rule `ASSIGN` generates the constraint $\text{writes}_{\text{Mass}}$ for assignment to the field in region `Mass` and the constraint $\text{writes}_{\text{Force}}$ for method `setForce`. In method `initialize`, there are two method invocations (lines 17 and 19). Therefore, rule `INVOKE` generates two constraints $\text{invokes}_{\text{setMass}}$ and $\text{invokes}_{\text{setForce}}$.

3.1.2. Constraint solving. The goal of the constraint solving phase is to associate with each method μ an *effect set* E_μ , where each element of E_μ is one of the effects $\text{reads } r$ or $\text{writes } r$. This phase comprises the following steps:

- 1) For each method μ , for each constraint $\text{invokes } \mu'$ in $K_{\mu'}$, add the elements of $K_{\mu'}$ to K_μ .
- 2) Repeat step 1 until no more constraints are added to any K_μ .

In step 1, we prune the constraint sets K_μ by never adding redundant constraints. For example, since writes cover reads in our effect system, there is no need for any K_μ to contain both $\text{reads } r$ and $\text{writes } r$; the second constraint suffices.

The algorithm terminates, because the total number of regions is bounded, so the total number of constraints that can be added to the K_μ is bounded. At the end of this process, for each μ we let $E_\mu = \text{effects}(K_\mu)$, where the function effects extracts the read and write constraints (i.e., the effects) from K_μ . As an example, from Figure 1, the constraints $\text{invokes}_{\text{setMass}}$ and $\text{invokes}_{\text{setForce}}$ generate the effects $\text{writes}_{\text{Mass}}$ and $\text{writes}_{\text{Force}}$.

3.2. Region Parameters

In this section, we extend Core DPJ by adding region parameters. Figure 6 shows the new syntax.

Meaning	Symbol	Definition
Classes	class	$\text{class } C\langle P \rangle \{ \text{field}^* \mu^* \}$
Regions	R	$r \mid P$
Types	T	$C\langle R \rangle$

Figure 6. New syntax of Core DPJ with region parameters. P is an identifier.

3.2.1. Constraint generation. Figure 7 shows the rules for generating reads , writes , and invokes constraints in Core DPJ with region parameters. The new rule `INVOKE`

records the *region substitution* $\theta = \{P \mapsto R\}$ that the constraint solver will need when translating the effects of one method to another. The term $\text{param}(C)$ represents the region parameter P of class C .

$$\begin{array}{l}
 \text{(FIELD-ACCESS)} \quad \frac{(\text{this}, C\langle P \rangle) \in \Gamma \quad \text{field}(C, f) = T \ f \ \text{in } R}{\Gamma \vdash \text{this}.f : T, \{ \text{reads } R \}} \\
 \text{(ASSIGN)} \quad \frac{(\text{this}, C\langle P \rangle) \in \Gamma \quad \Gamma \vdash e : T, K \quad \text{field}(C, f) = T' \ f \ \text{in } R}{\Gamma \vdash \text{this}.f = e : T, K \cup \{ \text{writes } R \}} \\
 \text{(INVOKE)} \quad \frac{\Gamma \vdash e_1 : C\langle R \rangle, K_1 \quad \Gamma \vdash e_2 : T, K_2 \quad \text{method}(C, m) = \mu}{\Gamma \vdash e_1.m(e_2) : \theta(T_r), K_1 \cup K_2 \cup \{ \text{invokes } \mu \ \text{where } \theta \}} \\
 \text{(NEW)} \quad \frac{}{\text{new } C\langle R \rangle : C\langle R \rangle, \emptyset} \quad \text{(VARIABLE)} \quad \frac{(z, T) \in \Gamma}{z : T, \emptyset}
 \end{array}$$

Figure 7. Rules for generating constraints in Core DPJ with region parameters.

As an example, we show how to generate the constraints for the code in Figure 2. According to rule `ASSIGN`, line 8 generates the constraint $\text{writes } P$. In line 13, Rule `FIELD-ACCESS` generates the constraint $\text{reads } L$ for accessing the field `left`. Then, because the type of `this.left` is `Node<L>`, rule `INVOKE` generates the constraint set $\{ \text{reads } L, \text{invokes}_{\text{setMass}} \ \text{where } \{P \mapsto L\} \}$. Line 15 generates similar constraints, using region `R` instead of `L`.

3.2.2. Constraint solving. The constraint solving phase is identical to the one described in Section 3.1.2, except that the algorithm applies substitutions θ in resolving invokes constraints:

- 1) For each method μ , for each constraint ($\text{invokes } \mu'$ where θ) in $K_{\mu'}$, add the elements of $\theta(K_{\mu'})$ to K_μ .
- 2) Repeat step 1 until no more constraints are added to any K_μ .

Here we apply the substitution θ elementwise to sets K_μ , and we apply θ to constraints as follows:

$$\begin{aligned}
 \theta(\text{reads } R) &= \text{reads } \theta(R) \\
 \theta(\text{writes } R) &= \text{writes } \theta(R) \\
 \theta(\text{invokes } \mu \ \text{where } \theta') &= \text{invokes } \mu \ \text{where } \theta(\theta') \\
 \theta(\{P \mapsto R\}) &= \{P \mapsto \theta(R)\}
 \end{aligned}$$

The algorithm terminates for the same reason given in Section 3.1.2.

Figure 8 illustrates the constraints and effects inferred by each iteration of the algorithm on the method `setMassOfChildren` in Figure 2. For brevity, we show only the effects coming from `left.setMass()`. The effect of method `setMass` is summarized as $\text{writes } P$ in iteration 0 (just before execution of step 3). In iteration 1, the invokes effect leads the algorithm to infer the effect

writes L by applying the substitutions $P \mapsto L$ on the effect of `setMass`. The algorithm does not infer any new effects in iteration 2, so it terminates after the second iteration.

	Iteration 0	Iteration 1
Effects	reads L	writes L
Constraints	invokes <code>setMass</code> where $\{P \mapsto L\}$	

Figure 8. Effects and constraints inferred in each iteration of the algorithm for method `setMassOfChildren` in Figure 2

3.3. Region Path Lists (RPLs)

In this section, we add RPLs to Core DPJ. Only the syntax for regions, shown in Figure 9, is new. `Root` is a special name representing the root of the region tree.

Meaning	Symbol	Definition
Regions	R	<code>Root r P R : r R : *</code>

Figure 9. Syntax of RPLs.

Constraint generation is the same as explained in Section 3.2.1. However, we need to extend the solving phase to handle recursion that would not terminate if we naively applied the algorithm from Section 3.2.2. For example, that algorithm would not terminate on the code in Figure 3, because it would try to infer effects on infinite chains of RPL elements, such as $P : L : R : \dots$. In such a case, we want to cut off the recursion and summarize the infinite set of RPLs with a *partially specified* RPL ending in $*$, as described in Section 2.3.

3.3.1. Algorithm description. Figure 10 shows the modified constraint solving algorithm. The algorithm iterates until all the effect and constraint sets stabilize. As before, each iteration of the main loop iterates over the method set \mathcal{M} and adds effects implied by the *invokes* constraints of K_μ .

However, instead of just adding an *invokes* constraint, we first check whether the constraint is *recursive* and has an *expanding substitution*. A constraint (*invokes* μ where θ) $\in K_{\mu'}$ is recursive if $\mu = \mu'$, i.e., the method includes its own effects, through a chain of one or more invocations. A substitution $P \mapsto R$ is expanding if P is the first RPL element of R , and R has more than one element. For example, $P \mapsto P : R$ is expanding but $P \mapsto P$ and $P \mapsto P' : R$ are not. If the constraint is recursive and has an expanding substitution, then we summarize the effects of K_μ by replacing each appearance of P in $effects(K_\mu)$ with $P : R : *$, and we add all new effects generated this way back into K_μ . In line 5, the function *summarize* takes $P \mapsto R$ to $P \mapsto R : *$.

If the constraint is either not recursive or not expanding (line 6), then we add the effects of $E_{\mu'}$ to E_μ after applying the substitution θ . We also add the non-recursive or non-expanding *invokes* constraints of $K_{\mu'}$ to K_μ , again after applying θ (line 9), so we can continue down the call graph until we hit an expanding recursion.

As before, all unions are up to redundant constraints and effects. For instance, once *writes* $R : *$ appears in K_μ , we never again add *reads* R or *writes* R to K_μ in line 5. Similarly, once *invokes* μ with $P \mapsto P : R$ appears in K_μ , we never add *invokes* μ with $P \mapsto P : R : R$ in line 10. This pruning ensures that the algorithm terminates (Section 3.3.3).

```

input : Program  $\mathcal{P}$  with region annotations
       Set  $\mathcal{M}$  of methods
       Set  $K_\mu$  of constraints for each method  $\mu$ 
output: A set of effects,  $E_\mu$ , for each method  $\mu$ 
1 repeat
2   foreach  $\mu$  in  $\mathcal{M}$  do
3     foreach  $c = (\text{invokes } \mu' \text{ where } \theta) \in K_\mu$  do
4       if  $\mu' = \mu$  and isExpanding( $\theta$ ) then
5          $K_\mu \leftarrow K_\mu \cup \text{summarize}(\theta)(\text{effects}(K_{\mu'}))$ 
6       else
7          $K_\mu \leftarrow K_\mu \cup \theta(\text{effects}(K_{\mu'}))$ 
8         foreach  $c' = (\text{invokes } \mu'' \text{ where } \theta')$  in  $K_{\mu'}$  do
9           if  $\mu'' \neq \mu'$  or not isExpanding( $\theta'$ ) then
10             $K_\mu \leftarrow K_\mu \cup \theta(c')$ 
11 until no  $K_\mu$  changes
12 foreach  $\mu$  in  $\mathcal{M}$  do
13    $E_\mu = \text{effects}(K_\mu)$ 

```

Figure 10. The inference algorithm for RPLs.

	Iteration 0	Iteration 1
Effects	reads $P : L$, $P : R$ writes P	writes $P : L : *$, $P : R : *$
Constraints	invokes <code>setMassForTree()</code> where $\{P \mapsto P : L\}$, <code>setMassForTree()</code> where $\{P \mapsto P : R\}$	

Figure 11. Effects and constraints inferred in each iteration of Figure 10 for the method `setMassForTree` in Figure 3.

3.3.2. Example. Figure 11 illustrates the constraints and effects inferred by each iteration of the while loop in lines 1–10 of Figure 10 for the method `setMassForTree` in Figure 3. Iteration 0 lists the constraints and effects as of line 2 in Figure 10. In iteration 1, the algorithm detects the recursive *invokes* constraints with expanding substitutions, appends stars to these two substitutions, and applies the new substitutions on the effects discovered in iteration 0 to get the two new effects of iteration 1. The algorithm terminates after iteration 2 because it does not find any new effects in the second iteration.

Note that even though the effect summary *writes* $P : *$ shown in Figure 3 is correct (i.e., it type-checks), `DPJZER` infers a more accurate (i.e., a more refined) summary. For

this program, `DPJIZER` infers `writes P, P:L:*, P:R:*`. The effect `writes P` comes from the write access in line 8. `writes P:L:*` comes from the recursive function in line 11. `DPJIZER` recognizes the recursive traversal of the data structure, and partially specifies the affected regions as `P:L:*`. `writes P:R:*` comes from the recursive function in line 13.

3.3.3. Termination and algorithmic complexity. The algorithm terminates because the pruning of redundant constraints (Section 3.3.1) ensures that no repeating RPL can appear in an effect or constraint with a period longer than the longest acyclic path in the call graph times the longest RPL appearing on the right-hand side of a substitution θ in the original constraint sets K_μ . The length of such an RPL is bounded by the length of the longest RPL appearing in the program text. Therefore, again the total number of effects and constraints that can be added to the K_μ is finite.

The algorithm is context-sensitive, so in some cases it will enumerate several call paths between two functions. While this could cause exponential complexity in the worst case, this is not likely to be a problem for realistic programs because it is very unlikely that the analysis will generate distinct constraints on exponentially many call paths. For each of the programs discussed in Section 5, `DPJIZER` inferred the effects in less than one second.

3.4. Other DPJ Features

We now show how we extended the algorithm described in Section 3.3 to handle the key remaining features of DPJ: arrays, owner regions, and inheritance.

3.4.1. Arrays. DPJ provides two capabilities for computing with arrays: *array RPL elements* and *index-parameterized arrays*. An array RPL element is $[e]$, where e is an integer expression. Since array regions are just RPL elements (e.g., `Root: [e] : r`), the algorithm can handle them in exactly the same way as described for name RPL elements. We just need a constraint collection rule that says that if expression e uses a method-local variable that is out of scope at the point of the method prototype, then we replace the element $[e]$ in any RPL with $[?]$, representing an *unknown* array index element in the DPJ type system.

An index-parameterized array allows the programmer to use an array index expression in the type of an indexed element. For example, the programmer can specify that the type of array index expression $A[e]$ is $C<[e]>$. To handle index-parameterized arrays, we just add constraint generation rules for assignment and access through array index expressions that are nearly identical to the rules for field assignment and access shown in Figure 7. The rules

are also similar to the rules for array access typing shown in [5].

3.4.2. Owner regions. DPJ provides a mechanism called *owner regions* for recursively partitioning a flat data structure (such as an array) in a divide-and-conquer manner. Figure 12 illustrates how to use owner regions to write parallel quicksort. Here the class `DPJArray` wraps an ordinary Java array and can be used to partition the array into subranges, and class `DPJPartition` is used to split the `DPJArray` into left and right segments `segs.left` and `segs.right`, as shown in line 7.

In line 10, the type of `segs.left` is `segs:DPJPartition.Left`, where `DPJPartitionLeft` is a field region name (Section 2.1) and the final local variable `segs` functions as an RPL. When a variable appears as an RPL, the region it represents is associated with the object reference stored in the variable at runtime, as in ownership type systems [6], [9]. The region of the variable is nested under the region bound to the first parameter of the variable's type. Here, `segs` has type `DPJPartition<P>`, so `segs` is nested under P . This fact allows us to write the method effect summary `writes P:*` covering both the write to P in line 6 and the recursive invocations of `sort` in lines 10 and 12. Given this effect summary, the compiler can use the inferred effects shown in lines 10 and 13 to prove that the statements in the `cobegin` block are noninterfering.

```

1 class QSort<region P> {
2   final DPJArray<P> A in P;
3   QSort(DPJArray<R> A) pure { this.A = A; }
4   void sort() writes P:* {
5     /* Quicksort partition: writes P */
6     int p = qsPartition(A);
7     final DPJPartition<P> segs =
8       new DPJPartition<P>(A, p);
9
10    cobegin {
11      /* writes segs:DPJPartition.Left:* */
12      new QSort<segs:DPJPartition.Left>
13        (segs.left).sort();
14      /* writes segs:DPJPartition.Right:* */
15      new QSort<segs:DPJPartition.Right>
16        (segs.right).sort();
17    }
18  }

```

Figure 12. Using owner regions to write quicksort.

Figure 13 shows the syntax of Core DPJ extended to support owner regions. Note that we have changed the syntax of expressions in the following ways: (1) we add a `let` construct to simulate `final` local variables; and (2) we require the selector and actual argument of a method invocation expression to be variables to keep the typing rules simple.

Constraint generation works exactly as described in Section 3.2.1 except for rules `LET` and `INVOKE`, shown in Figure 14. In rule `LET`, we have to account for the fact that

Meaning	Symbol	Definition
Regions	R	Root r P z R : r R : *
Expressions	e	let z = e in e this.f this.f = e z.m(z) new C⟨R⟩ z

Figure 13. Core DPJ with owner regions.

RPLs generated inside the `let` expression may contain a local variable z that is not in scope outside the body of the expression. Therefore, we coarsen any such RPL $z : R$ to $R' : *$, where the type of z is $C\langle R' \rangle$. Rule `INVOKE` is nearly identical to the one shown in Figure 7, except that we record the substitutions `this` $\mapsto z_1$ and $x \mapsto z_2$ as well as the substitution $param(C) \mapsto R$. With these changes, the solving algorithm works exactly as described in Section 3.3.1.

$$\begin{array}{c}
\text{(LET)} \quad \frac{\Gamma \vdash e_1 : C\langle R \rangle, K_1 \quad \Gamma \cup \{(x, T_1)\} \vdash e_2 : T_2, K_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \theta(T_2), \theta(K_1 \cup K_2)} \\
\theta = \{x \mapsto R : *\} \\
\text{(INVOKE)} \quad \frac{\{(z_1, C\langle R \rangle), (z_2, T)\} \subseteq \Gamma \quad \text{method}(C, m) = \mu}{\mu = T_r, m(T_x, x) \{e'\}} \quad \theta = \{param(C) \mapsto R, \text{this} \mapsto z_1, x \mapsto z_2\}}{\Gamma \vdash z_1.m(z_2) : \theta(T_r), K \cup \{\text{invokes } \mu \text{ where } \theta\}}
\end{array}$$

Figure 14. Rules for generating constraints in Core DPJ with owner regions.

3.4.3. Inheritance. The real DPJ language supports inheritance, e.g., `class B<P> extends A<R>`. Inheritance raises two issues for the inference algorithm. First, we must translate inherited methods and fields from the superclass to the subclass. We do this by applying the *translating substitution* θ implied by the chain of `extends` clauses from the superclass to the subclass. For example, if class $C_1\langle P_1 \rangle$ extends $C_2\langle R_1 \rangle$, and $C_2\langle P_2 \rangle$ extends `Object`, then the translating substitution from C_2 to C_1 is $\{P_2 \mapsto R_1\}$.

Second, DPJ requires that the declared effects of a method include the effects of all overriding methods [5]. This gives rise to a constraint similar to the one we represented by an *invokes* constraint, except that it is simpler, because there is no recursion in the inheritance graph. To handle this constraint, we make two simple extensions to the algorithm. First, in the constraint collection phase, after collecting constraints from each method body, we add to each K_μ the constraint *isOverriddenBy* μ' where θ , for each method μ' such that μ is overridden by μ' . Here θ is the translating substitution defined above. Second, in the constraint solving phase, in each iteration of the `repeat` loop in Figure 10, between lines 10 and 11, we add another iteration over all methods $\mu \in \mathcal{M}$ to add $\theta(K_{\mu'})$ to K_μ for each constraint *isOverriddenBy* μ' where θ in K_μ .

3.5. Applicability Beyond DPJ

The relevance of our effect inference algorithm is not limited to DPJ: with suitable modifications, the algorithm can be adapted to infer effects for other object-oriented effect systems, such as ownership-based systems [6], [10], [11], that have features similar to DPJ's. Here we illustrate how the inference algorithm might be adapted to work on the ownership-based effect system by Clarke and Drossopoulou called Java with Ownership and Effects, or JOE [6].

JOE also employs method effect summaries and supports effects on regions similar to DPJ's owner regions, except that JOE has no RPLs. Instead, JOE uses *effect shapes* of the form $p.n$ and under $p.n$, where p is a *final* local variable or context parameter (similar to a DPJ region parameter), and $n \geq 0$ is a natural number. The shape $p.n$ refers to all descendants of p in the region tree that are n levels below p in the tree, with $p.0$ being equivalent to p . The shape under $p.n$ is similar to an RPL with $*$ at the end and refers to all $p.n'$ such that $n' \geq n$. The key rule of JOE, which defines the region tree, is that if variable z has type $C\langle o \rangle$, then the shape $z.n$ is covered by the shape $o.n+1$, where $o = owner(z)$ is the region bound to the owner parameter in the type of z .

To adapt our algorithm to JOE, we make the following modifications. First, instead of substitutions $P \mapsto R$, we use substitutions $p.n \mapsto p'.n+k$, for $k \geq 0$. Second, in rule `LET` (Figure 14), when a variable z goes out of scope, we generate an effect for the outer scope by applying the substitution $z.n \mapsto o.n+1$, where $o = owner(z)$. (We could also replace $z.n$ with `under o`, as our `LET` rule does for DPJ, but this would be less precise in the context of JOE.) Third, we define an expanding substitution to be $p.n \mapsto p.n+k$, for $k > 0$, and when we see an expanding substitution, we replace its right-hand side with `under p.n+k`. Otherwise, the algorithm works as described in the previous sections.

```

1 class List<o> {
2   int data;
3   List<this> next;
4   void update(int data) writes this, under this+1 {
5     /* writes this */
6     this.data = data;
7     /* invokes update where {this.n  $\mapsto$  this.n+1} */
8     let z = next in
9       /* invokes update where {this.n  $\mapsto$  z.n+0} */
10      if (z != null) z.update(data);
11  }
12 }

```

Figure 15. Example of inferring effects for JOE.

Figure 15 shows how the algorithm infers effects for a simple recursive JOE program. This code traverses and updates a list such that each node of the list owns the next node. The initial constraints gathered in the constraint

collecting phase are shown in the comments. Rule `INVOKE` generates the effect in line 9, which is adjusted to the effect in line 7 by the `LET` rule discussed above. At the end of initial constraint collection, the constraints are as shown in lines 5 and 7. In the first iteration of the solving algorithm, the expanding substitution shown in line 7 gets summarized as $\text{this}.n \mapsto \text{under } \text{this}.n + 1$. Applying this substitution to the effect `writes this` and putting the result back into the constraint set yields the inferred effects shown in line 4. The algorithm then terminates because there are no new effects to add.

4. The DPJIZER Tool

We have built an interactive tool, `DPJIZER`, as an Eclipse plug-in, which incorporates the algorithm discussed in Section 3. Given a partial DPJ program with legal region annotations, `DPJIZER` produces a legal DPJ program with region and effect annotations. In addition, the tool has some valuable interactive features. A programmer can select an effect in a method summary and `DPJIZER` highlights the statements or expressions that generated that effect (as seen in the screen fragment in Fig. 16). Alternately, the programmer can select a statement or expression, and `DPJIZER` highlights its corresponding effect in the effect summary. Thus, when the compiler reports interference warnings, the developer can use `DPJIZER` to localize the problem and refine the region annotations accordingly.

We now describe in more detail how `DPJIZER` helps programmers write DPJ programs. Typically, a DPJ programmer carries out the following steps to convert a given sequential Java program to DPJ. First, choose which sections of code are to be run in parallel, but don't yet insert the parallel constructs (`cobegin`, `foreach`, etc.). Second, devise a strategy for using region declarations, region parameters, and RPLs to express the noninterference of the parallel sections. Add these annotations and make sure they pass the type checker. At this point, the methods all have an empty summary, which in DPJ means the most conservative effect, i.e., `writes Root : *`. Such effect summaries will pass the type checker but will not allow parallelism to be safely expressed. Third, for methods transitively invoked by a parallel section, refine the method summaries as necessary to make the parallel tasks mutually noninterfering. Fourth, add the parallel constructs to the parallel sections. Fifth, if there are any interference warnings, then revisit steps two and three to revise the region annotations and/or method effect summaries to eliminate the interference.

`DPJIZER` helps this process in the following ways. First, step three is completely automated. This automation removes a lot of work from the development process, particularly if the user has to do two or more iterations of steps two and three. While the compiler provides

```
void setMass(double mass)
    writes Root : BinaryTree.M,
    Root : BinaryTree.L,
    Root : BinaryTree.R {
    this.mass = mass;
    initTree();
}
```

Figure 16. The programmer selects an effect in the effect summary and `DPJIZER` highlights the statement that generated that effect.

error information of the form “effect E is missing from the summary of method m ” that helps the user fix bad summaries, step three is still time consuming and difficult. In code with many methods and invocations, there are a lot of summaries to write. Further, it is difficult for a user to manually propagate effects backwards along the call graph and around cycles. `DPJIZER` automates this process.

Second, `DPJIZER` helps step five by identifying the statements and expressions within a method that are contributing “bad” effects. A key part of this step is understanding the statements and methods that contribute these effects; the tool simplifies that greatly by allowing users to map effects back to the expressions that produce them, and vice versa. With additional programming (not yet implemented for lack of time), the tool will also help understand better how effects are propagated from methods to call sites, along with the relevant substitutions and how they propagate around cycles in the call graph, in some cases leading to summarization with ‘*’.

5. Evaluation

Research Questions. To evaluate the effectiveness of `DPJIZER`, we answer the following two questions:

- **Q1:** Is `DPJIZER` useful? Does it alleviate the burden of writing effect annotations?
- **Q2:** Is the inference accurate? What is the granularity of the inferred effects?

We answer these questions in two ways: with a case study running the tool ourselves, and with a survey in which we asked other programmers who have written DPJ programs to run the tool and describe their experience using it. The case study (Section 5.1 provides *quantitative* answers, while the survey (Section 5.2) provides *qualitative* answers.

5.1. Case Study

5.1.1. Methodology. Table 1 lists the programs that we used as case studies. These programs were manually annotated with regions and effects by other programmers before the existence of `DPJIZER`. We took these programs, erased the effect annotations and left only the region annotations, and we used `DPJIZER` to infer the method effects.

Program	SLOC	Effects Inferred	
		by Programmer	by DPJIZER
Barnes-Hut	682	46	145
IDEA	228	6	6
K-means	501	3	42
ListRanking	105	4	30
MergeSort	295	17	28
MonteCarlo	2877	38	273
QuadTree	117	11	28
QuickSort	144	12	13
StringMatching	373	61	60
SumReduce	57	3	12
Total	5379	201	637

Table 1. Programs used as case studies. Program size is given in non-blank, non-comment lines of source code, counted by *sloccount*. The last two columns show the number of effects written by programmers or inferred by DPJIZER.

To answer the first question (usefulness), we report the number of effects that programmers wrote manually, as well as the number of effects that DPJIZER inferred automatically. To answer the second question (accuracy) we compared the effects written manually with the effects inferred by DPJIZER.

5.1.2. Quantitative Results.

Q1: Is DPJIZER useful? From Table 1 one can see that if the programmers had used DPJIZER to infer the method effects, they would have saved writing 201 effects. Further, the programmers would have saved the time it took to generate these effects by manually propagating constraints backwards through the call graph, around cycles, and up the inheritance graph.

Also, note that in some cases (e.g., Barnes-Hut, KMeans), DPJIZER generated many more effects than the programmers. This is because the programmers put in effects only for methods that were transitively invoked in a parallel construct. Because DPJIZER infers effects for *all* methods in a program, these effects can serve for future parallelization of new code fragments, thus encouraging an incremental parallelization approach. In addition, a programmer can look over the inferred effects to detect patterns of data accesses.

Q2: Is the inference accurate? We carefully analyzed the programs in Table 1 and compare the effects written by programmers with the effects inferred automatically by DPJIZER. Since programmers did not write effects for all methods, as explained above, we can only compare the effects for the methods that were annotated. Table 2 shows the number of inaccurate effects, i.e., effects that are too coarse-grained in comparison with the effects inferred by DPJIZER. Note that in all cases, DPJIZER infers effects that are the same as, or more precise (i.e., finer-

grained) than, those written by the programmer: although DPJIZER must be conservative, it is also quite accurate. Programmers can be at least as accurate if they choose, but sometimes choose to summarize effects, e.g., if they think the coarser effects will not inhibit parallelism.

Table 2 shows three sources of inaccuracy in the manually inferred effects. First, some of these effects are too coarse-grained in the choice of effect keyword, e.g., *writes_R* instead of *reads_R*. This is legal (i.e., it type-checks) but unnecessarily coarse and forbids the parallel execution of two methods (e.g., two `get()` methods) that only read region *R* and otherwise could have been executed in parallel.

Second, some manually inferred effects are too coarse-grained in the region specification. For example, the programmer specified *writes_{P:*}* when the appropriate effect inferred by DPJIZER was *writes_{P, P:L:*, P:R:*}*. The coarser region inferred by programmer forbids any other method that writes in a subregion of *P* to run in parallel. This is an unnecessary restriction because the method only writes in subregions created using the *L* or *R* prefixes, so that another method that writes into *P:M* should be allowed to run in parallel.

Third, some manually inferred effects are redundant. For example, the programmer may specify *reads_R writes_R*, but the read effect is subsumed by the write. Alternatively, the programmer inferred *writes_{P, P:*}*, where the first region is redundant since it is subsumed by the second region. Such redundancies do not hinder parallelism but make the method effect summary unnecessarily verbose, which can hinder program understanding.

We carefully analyzed the source code of the methods, and indeed DPJIZER inferred the most fine-grained effects that are possible to express with the current DPJ language, and the summaries do not contain redundant effects.

5.2. User Survey

We also conducted a preliminary survey of other programmers who have previously written DPJ programs. This study took the following steps:

- 1) We elided the effect annotations on the programs previously parallelized by these users, but retained the region annotations they had inserted.
- 2) The users then used DPJizer to infer the effect annotations for those programs.
- 3) The users finally filled out a brief questionnaire asking about the results, usability and overall experience of using DPJIZER.

This study is limited because it only has a small number of users and they all know the study authors. Nevertheless, it provides some preliminary feedback on the usefulness of the tool from experienced DPJ programmers not involved in designing or building DPJIZER (none of

Program	# of Effects Too Coarse By keyword	region	# of Redundant Effects
Barnes-Hut	1	0	3
IDEA	0	0	0
K-means	0	0	0
ListRanking	1	2	0
MergeSort	0	4	0
MonteCarlo	1	3	6
QuadTree	1	2	1
QuickSort	0	3	3
StringMatching	5	24	10
SumReduce	0	0	2
Total	9	38	25

Table 2. Number of inaccurate effects written by programmers. We report effects that are too coarse-grained by keyword (e.g., programmer wrote *writes_R* instead of *reads_R*), by region (e.g., programmer wrote *reads_{R:*}* instead of *reads_R*). Last column shows the number of redundant effects (e.g., programmer wrote both *reads_R* *writes_R*).

them had seen or even participated in discussions about the tool before the survey).

Usefulness. The users said the tool saved “a significant fraction” of porting effort. One user said the tool saved “a lot of time in the process of writing/adding annotations ... and then compile to find more methods to annotate.”

Accuracy. One user thought the tool inferred too many annotations: he would prefer to see fewer effect annotations, and could re-run the tool if more were needed. Conversely, he said the tool did help eliminate some redundant annotations (compared with his manual effect summaries).

Requested features. The most requested features included incremental addition of annotations; presenting choices of annotations to the user and letting him choose; and recommend better region structure to produce more fine-grain effects. (The latter is outside the scope of the current work but is a subject of future work, as described in Section 7.)

Summary. Overall, all users said that they would use DPJIZER to help write DPJ programs. One user said “I think it will also help me redesign region structures to be more precise and effective.”

6. Related Work

Method effect summaries. Many effect systems employ effect summaries to enable modular analysis and composability of program components. The original proposals for an object-oriented effect system [3], [12] use summaries, as do several systems combining object ownership with some form of effects [6], [11], [13]. Our

work presents an algorithm and tool that can be used to infer such summaries.

Effect inference. The seminal work on inferring effects is from Jouvelot and Gifford [14]. They use a technique called *algebraic reconstruction* to infer types and effects in a mostly functional language. Talpin and Jouvelot [15] build on this work to develop a constraint-based solving algorithm. These algorithms are tailored to a mostly-functional language with a much simpler effect system than DPJ’s: nested effects cannot be expressed, so no summaries such as $R : *$ have to be inferred.

Bierman and Parkinson [16] present an inference algorithm for Greenhouse and Boyland’s effect system [3]. The features they consider are similar to the smallest subset of Core DPJ we covered in section 3.1, plus support for unique reference annotations and a limited form of nesting. Again there is no unbounded nesting.

Side-effect analysis [17]–[20] uses interprocedural alias analysis and dataflow propagation algorithms to compute the side effects of functions. There are two major differences between these algorithms and DPJizer. First, DPJizer operates on programmer-specified region types, which identify and express effects more precisely than alias analysis. Second, DPJizer exploits the structure of RPLs to do a custom solution for recursive calls, which should significantly speed up convergence of the constraint solver.

Commutativity analysis [21] uses symbolic execution to collect the side effects of methods and reason about which pairs of methods commute with each other. The analysis is fully automatic, but less expressive than DPJ, because programs must be written in a certain restricted style in order for the analysis to work.

Region and type inference. There is extensive literature on *region inference* for region-based memory management [22]–[25]. Several researchers have studied the problem of inferring types or type qualifiers for imperative programs with references. Kiezun et al. [26] show how to infer Java generic parameters and arguments. Agarwal and Stoller show how to do type inference for parameterized race-free Java [27]. Quinonez et al. [28] present a tool called Javarifier for inferring *reference immutability* for variables (i.e., that the reference is never used to update the state of any object that it transitively points to). Terauchi and Aiken [29] present a type inference algorithm for deterministic parallelism using *linear types* supplemented with fractional permissions [30].

These algorithms are broadly similar to ours, in that they collect constraints across the whole program and solve them. However, the technical details are quite different because the problem domains differ from our problem of inferring effects for nested regions. The region and type inference techniques may be useful in extending DPJIZER to infer DPJ region annotations.

7. Conclusions

We have presented an effect inference algorithm and a tool, `DPJIZER`, that ease the burden of writing DPJ programs. The `DPJIZER` algorithm is also applicable to other effect systems that rely on method effect summaries. As future work, we plan to extend the capabilities of `DPJIZER` so that it can help with *region inference*, i.e., inferring region declarations, region parameters, and region arguments. Region inference in DPJ is challenging, but preliminary work indicates that it should be possible to infer regions for many common parallel patterns.

References

- [1] J. M. Lucassen *et al.*, “Polymorphic effect systems,” in *POPL*, 1988.
- [2] R. T. Hammel and D. K. Gifford, “FX-87 performance measurements: Dataflow implementation,” Tech. Rep. MIT/LCS/TR-421, 1988.
- [3] A. Greenhouse and J. Boyland, “An object-oriented effects system,” *ECOOP*, 1999.
- [4] “DPJ homepage,” <http://dpj.cs.uiuc.edu>.
- [5] R. L. Bocchino, V. S. Adve, D. Dig, S. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian, “A Type and Effect System for Deterministic Parallel Java,” to appear in *OOPSLA 2009*.
- [6] D. Clarke and S. Drossopoulou, “Ownership, encapsulation and the disjointness of type and effect,” in *OOPSLA*, 2002.
- [7] R. L. Bocchino, V. S. Adve, S. V. Adve, and M. Snir, “Parallel Programming Must Be Deterministic By Default,” in *First USENIX Workshop on Hot Topics in Parallelism (HotPar)*, 2009.
- [8] J. Gosling, B. Joy, G. Steele, and G. Bracha, *The Java Language Specification, Third Edition*. Addison-Wesley Longman, 2005.
- [9] D. G. Clarke *et al.*, “Ownership types for flexible alias protection,” *OOPSLA*, 1998.
- [10] C. Boyapati, B. Liskov, and L. Shriram, “Ownership types for object encapsulation,” in *POPL*, 2003.
- [11] C. Boyapati, R. Lee, and M. Rinard, “Ownership types for safe programming: preventing data races and deadlocks,” in *OOPSLA*, 2002.
- [12] K. R. M. Leino *et al.*, “Using data groups to specify and check side effects,” 2002.
- [13] N. R. Cameron, S. Drossopoulou, J. Noble, and M. J. Smith, “Multiple ownership,” *OOPSLA*, 2007.
- [14] P. Jouvelot and D. Gifford, “Algebraic reconstruction of types and effects,” in *POPL*, 1991.
- [15] J.-P. Talpin and P. Jouvelot, “Polymorphic type, region and effect inference,” *J. Funct. Prog.*, July 1992.
- [16] G. Bierman and M. Parkinson, “Effects and effect inference for a core java calculus,” *Workshop on Object Oriented Developments (WOOD)*, 2003.
- [17] J. P. Banning, “An efficient way to find the side effects of procedure calls and the aliases of variables,” in *POPL*, 1979.
- [18] B. G. Ryder, W. A. Landi, P. A. Stocks, S. Zhang, and R. Altucher, “A schema for interprocedural modification side-effect analysis with pointer aliasing,” *TOPLAS*, 2001.
- [19] A. Salcianu and M. C. Rinard, “Purity and side effect analysis for Java programs,” in *VMCAI*, 2005.
- [20] A. Rountev, “Precise identification of side-effect-free methods in java,” in *ICSM*, 2004.
- [21] P. C. Diniz, “Commutativity analysis: A new analysis technique for parallelizing compilers,” *TOPLAS*, 1997.
- [22] M. Tofte and L. Birkedal, “A region inference algorithm,” *TOPLAS*, 1998.
- [23] W.-N. Chin, F. Craciun, S. Qin, and M. Rinard, “Region inference for an object-oriented language,” in *PLDI*, 2004.
- [24] A. Banerjee, M. Barnett, and D. A. Naumann, “Boogie meets Regions: A verification experience report,” in *VSTTE*, 2008.
- [25] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney, “Region-based memory management in Cyclone,” *PLDI*, 2002.
- [26] A. Kiezun, M. D. Ernst, F. Tip, and R. M. Fuhrer, “Refactoring for parameterizing Java classes,” in *ICSE*, 2007.
- [27] R. Agarwal and S. D. Stoller, “Type inference for parameterized race-free Java,” in *VMCAI*, 2004.
- [28] J. Quinonez, M. S. Tschantz, and M. D. Ernst, “Inference of reference immutability,” in *ECOOP*, 2008.
- [29] T. Terauchi and A. Aiken, “A capability calculus for concurrency and determinism,” *TOPLAS*, 2008.
- [30] J. Boyland, “Checking interference with fractional permissions,” *SAS*, 2003.