

Evaluating the applyAllIter() Java Fork-Join Framework Programming Model

Douglas C. Schmidt

d.schmidt@vanderbilt.edu

www.dre.vanderbilt.edu/~schmidt

Professor of Computer Science

Institute for Software
Integrated Systems

Vanderbilt University
Nashville, Tennessee, USA



Learning Objectives in this Part of the Lesson

- Evaluate different fork-join framework programming models in practice
- Evaluate the applyAllIter() method
 - This method uses “work-stealing” to disperse tasks to worker threads

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
    return fjPool  
        .invoke(new  
            RecursiveTask  
                <List<T>>() {  
            protected List<T>  
                compute() {  
                    ...  
                }  
            } );  
    }
```

Implementing the applyAllIter() Method

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) {  
  
    return fjPool  
        .invoke(new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
  
                ...  
            }  
        }) ;  
}
```

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) {  
  
    return fjPool  
        .invoke(new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
                ...  
            }  
        }) ;  
}
```

These parameters are treated as "effectively final" variables in the anonymous inner class below

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) {  
  
    return fjPool  
        .invoke(new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
                ...  
            }  
        }) ;  
}
```

Create an anonymous RecursiveTask instance whose compute() method iterative creates many sub-tasks

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) {  
  
    return fjPool  
        .invoke(new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
                ...  
            }  
        }) ;  
}
```



Code is verbose due to lack of functional interface (& thus can't use lambdas)..

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) {  
  
    return fjPool  
        .invoke(new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
                ...  
            }  
        }) ;  
}
```

TWO WAY



*Invoke the task on the fork-join pool
& then wait for & return the results*

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> res = new LinkedList<>();  
  
    for (T t : list) Hook method implements the main fork-join task  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return op.apply(t); }  
        }.fork());  
  
    for (ForkJoinTask<T> task : forks) res.add(task.join());  
    return res;  
} ...
```

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> result = new LinkedList<>();  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return op.apply(t); }  
        }.fork());  
  
    for (ForkJoinTask<T> task : forks) result.add(task.join());  
    return result;  
} ...
```

Lists that hold the forked tasks & the results

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,  
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> result = new LinkedList<>();  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return op.apply(t); }  
        }.fork());  
  
    for (ForkJoinTask<T> task : forks) result.add(task.join());  
    return result;  
} ...
```

*Iterate through input list,
fork all the tasks, & add
them to the forks list*

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> result = new LinkedList<>();  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return  
                .fork(); }  
  
            for (ForkJoinTask<T> task : forks) result.add(task.join());  
            return result;  
        } ...  
}
```



This implementation relies on “work-stealing” to disperse tasks to worker threads

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> result = new LinkedList<>();  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return op.apply(t); }  
        }.fork());  
  
    for (ForkJoinTask<T> task : forks) result.add(task.join());  
    return result;  
} ...
```

Join all results of forked tasks & add to results list



"Collaborative Jiffy Lube" model of processing!

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...  
  
protected List<T> compute() {  
    List<ForkJoinTask<T>> forks = new LinkedList<>();  
    List<T> result = new LinkedList<>();  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute() { return op.apply(t); }  
        }.fork());  
  
    for (ForkJoinTask<T> task : forks) result.add(task.join());  
    return result;  
} ...
```

Return the results list

Implementing the applyAllIter() Method

- Apply an 'op' to all items in the list by calling fork-join methods iteratively

```
<T> List<T> applyAllIter(List<T> list, Function<T, T> op,
                           ForkJoinPool fjPool) { ...

protected List<T> compute() {
    List<ForkJoinTask<T>> forks = new LinkedList<>();
    List<T> result = new LinkedList<>();

    for (T t : list)
        forks.add(new RecursiveTask<T>() {
            protected T compute() { return op.apply(t); }
        }.fork());
}

for (ForkJoinTask<T> task : forks) result.add(task.join());
return result;
} ...
```

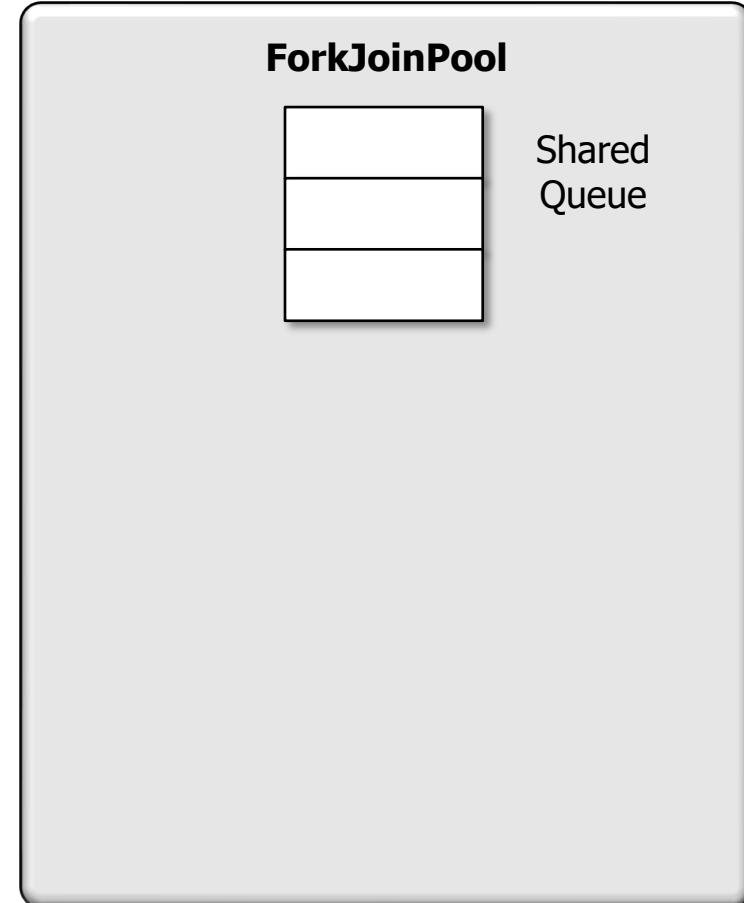
This implementation is very simple to program & understand since it's iterative

Visualizing the applyAllIter() Method

Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
  
    fjPool.invoke  
        (new RecursiveTask<List<T>>() {  
            protected List<T> compute() {  
                ...  
            }  
        }  
        ...  
    )  
}
```

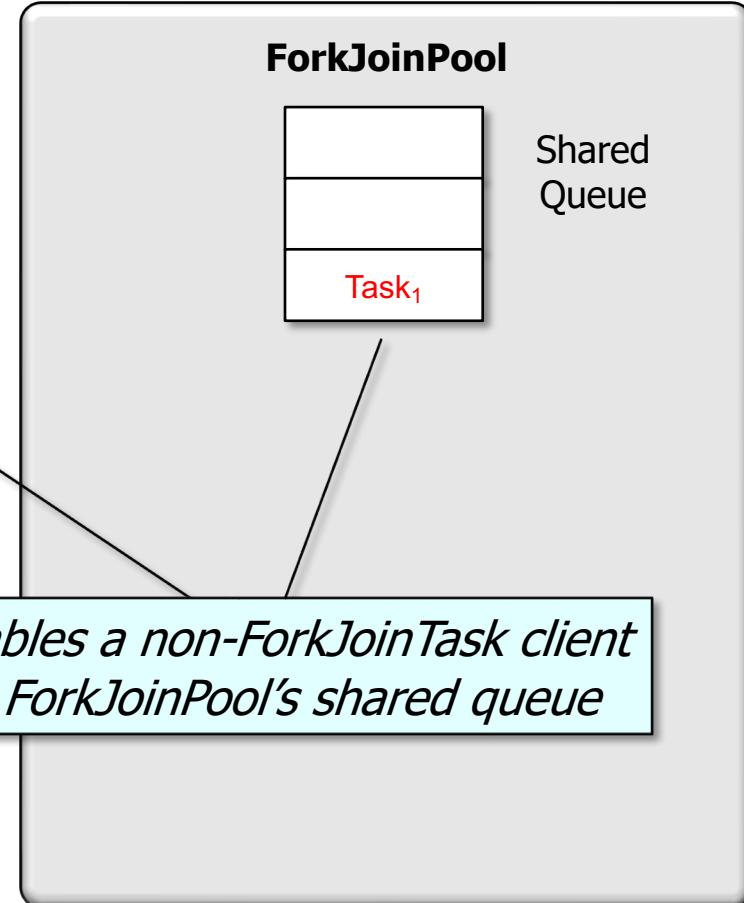


See [LiveLessons/blob/master/Java8/ex22/src/utils/ForkJoinUtils.java](#)

Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
fjPool.invoke  
    (new RecursiveTask<List<T>>() {  
        protected List<T> compute() {  
            ...  
        }  
    }  
    ...  
}
```



Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter(List<T> lis
```

Worker thread WT_1 gets Task₁ & creates n new sub-tasks that run 'op' on each list element

```
    Function<T, T> op,  
    ForkJoinPool fjPool) {
```

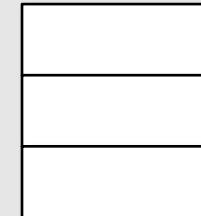
```
...
```

```
for (T t : list)  
    forks.add(new RecursiveTask<T>() {  
        protected T compute()  
        { return op.apply(t); }  
        .fork());
```

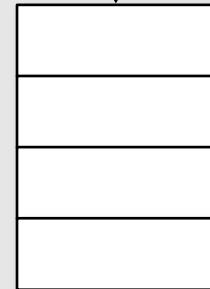
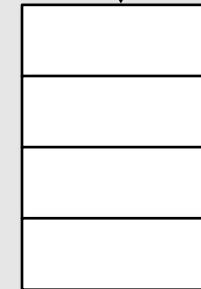
```
for (ForkJoinTask<T> task : forks)  
    result.add(task.join());
```

```
...
```

ForkJoinPool



Shared Queue

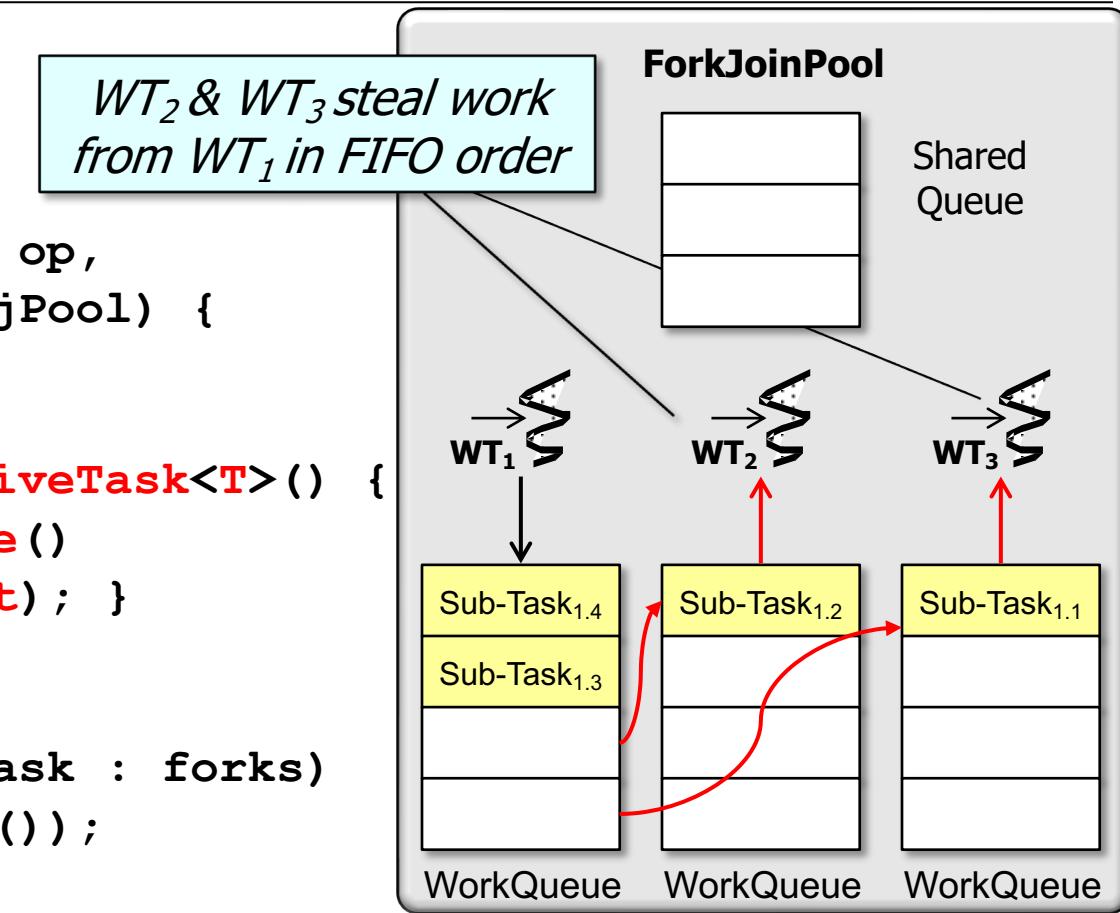


The highlighted code runs in the RecursiveTask's compute() method

Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
  
    ...  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute()  
            { return op.apply(t); }  
            .fork());  
  
    for (ForkJoinTask<T> task : forks)  
        result.add(task.join());  
    ...
```

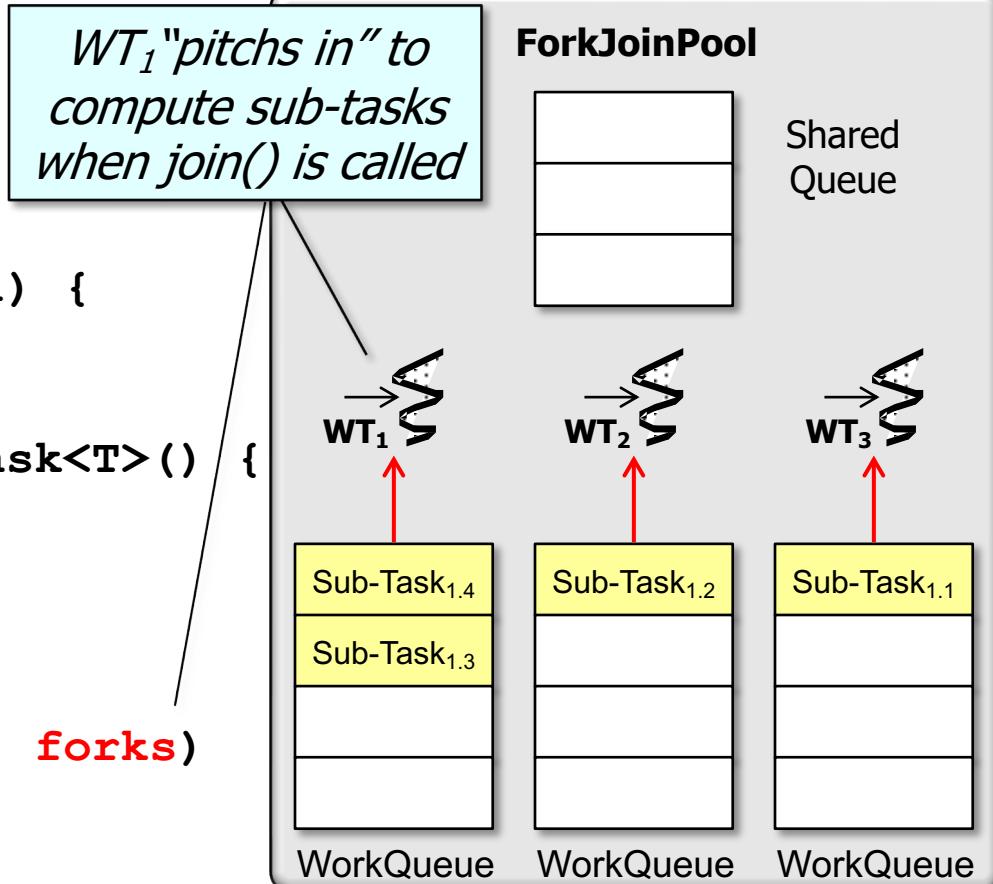


“Work-stealing” overhead is high, but copying & method call overhead is low

Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

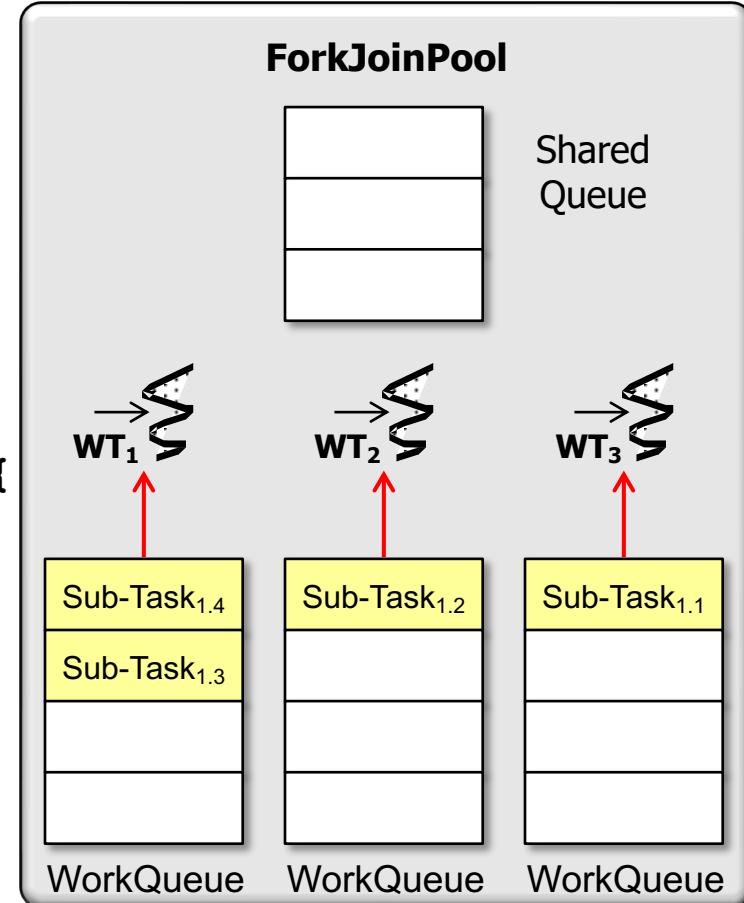
```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
  
    ...  
  
    for (T t : list)  
        forks.add(new RecursiveTask<T>()  
            {  
                protected T compute()  
                { return op.apply(t); }  
            }.fork());  
  
    for (ForkJoinTask<T> task : forks)  
        result.add(task.join());  
  
    ...
```



Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {  
  
    ...  
    for (T t : list)  
        forks.add(new RecursiveTask<T>() {  
            protected T compute()  
            { return op.apply(t); }  
            .fork());  
  
    for (ForkJoinTask<T> task : forks)  
        result.add(task.join());  
    ...
```



"Collaborative Jiffy Lube" model of processing!

Visualizing the applyAllIter() Method

- Visualizing applyAllIter()

```
<T> List<T> applyAllIter  
    (List<T> list,  
     Function<T, T> op,  
     ForkJoinPool fjPool) {
```

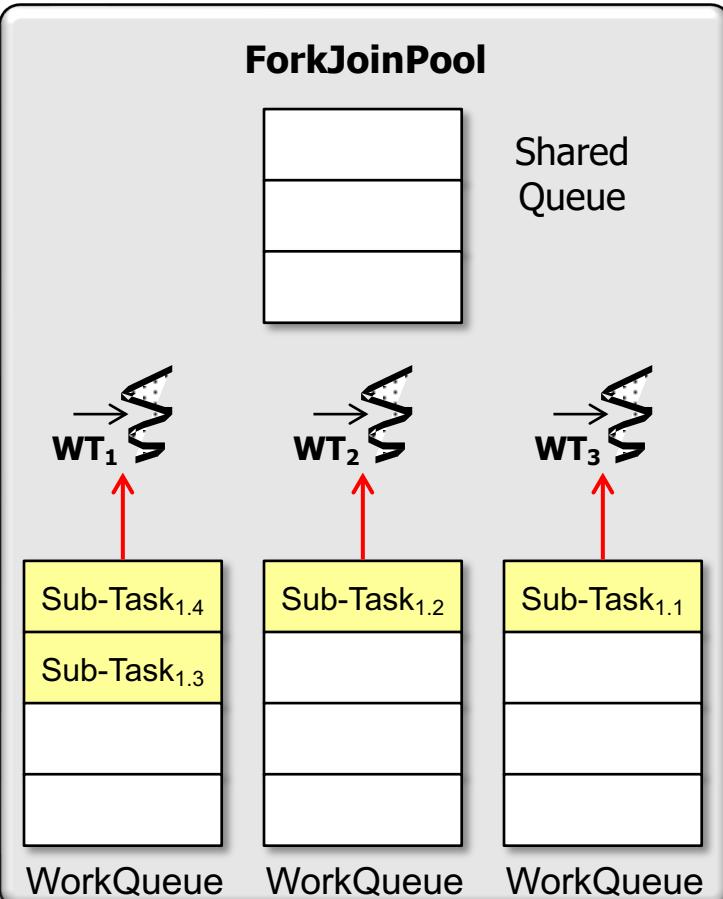
...

```
for (T t : list)  
    forks.add(new  
        protected T  
        { return op.apply(t); }  
    ).fork());
```

*This loop implements
"barrier synchronization"*

```
for (ForkJoinTask<T> task : forks)  
    result.add(task.join());
```

...



See [en.wikipedia.org/wiki/Barrier_\(computer_science\)](https://en.wikipedia.org/wiki/Barrier_(computer_science))

End of Evaluating the apply
AllIter() Java Fork-Join
Framework Programming
Model