

# When is Parallelism Fearless and Zero-Cost with Rust?

Javad Abdi, Gilead Posluns, Guozheng Zhang,  
Boxuan Wang, Mark C. Jeffrey

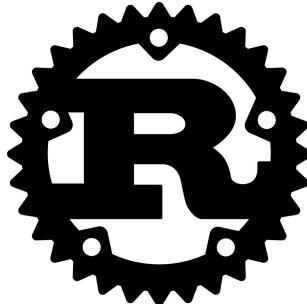


UNIVERSITY OF  
**TORONTO**

**SPAA 2024**

# Rust claims to solve the arduous task of parallel programming

---



**The Rust  
Programming  
Language**

: Rust provides fearless concurrency

Does Rust make expressing all parallel programs easy?

As it is,

- **It makes expressing easy parallelism fearless**
- **Its support is limited for hard cases of parallelism**

# Contributions

---

A **case study** of Rust's support for regular and irregular parallelism

- Classify parallel patterns found in PBBSv2 and MultiQueue-based algorithms
- Qualitatively evaluates the level of support for each pattern

Rust Parallel Benchmark Suite (**RPB**):

A benchmark suite with regular and irregular parallelism in Rust

- Performance evaluation of Rust vs. C++ benchmarks



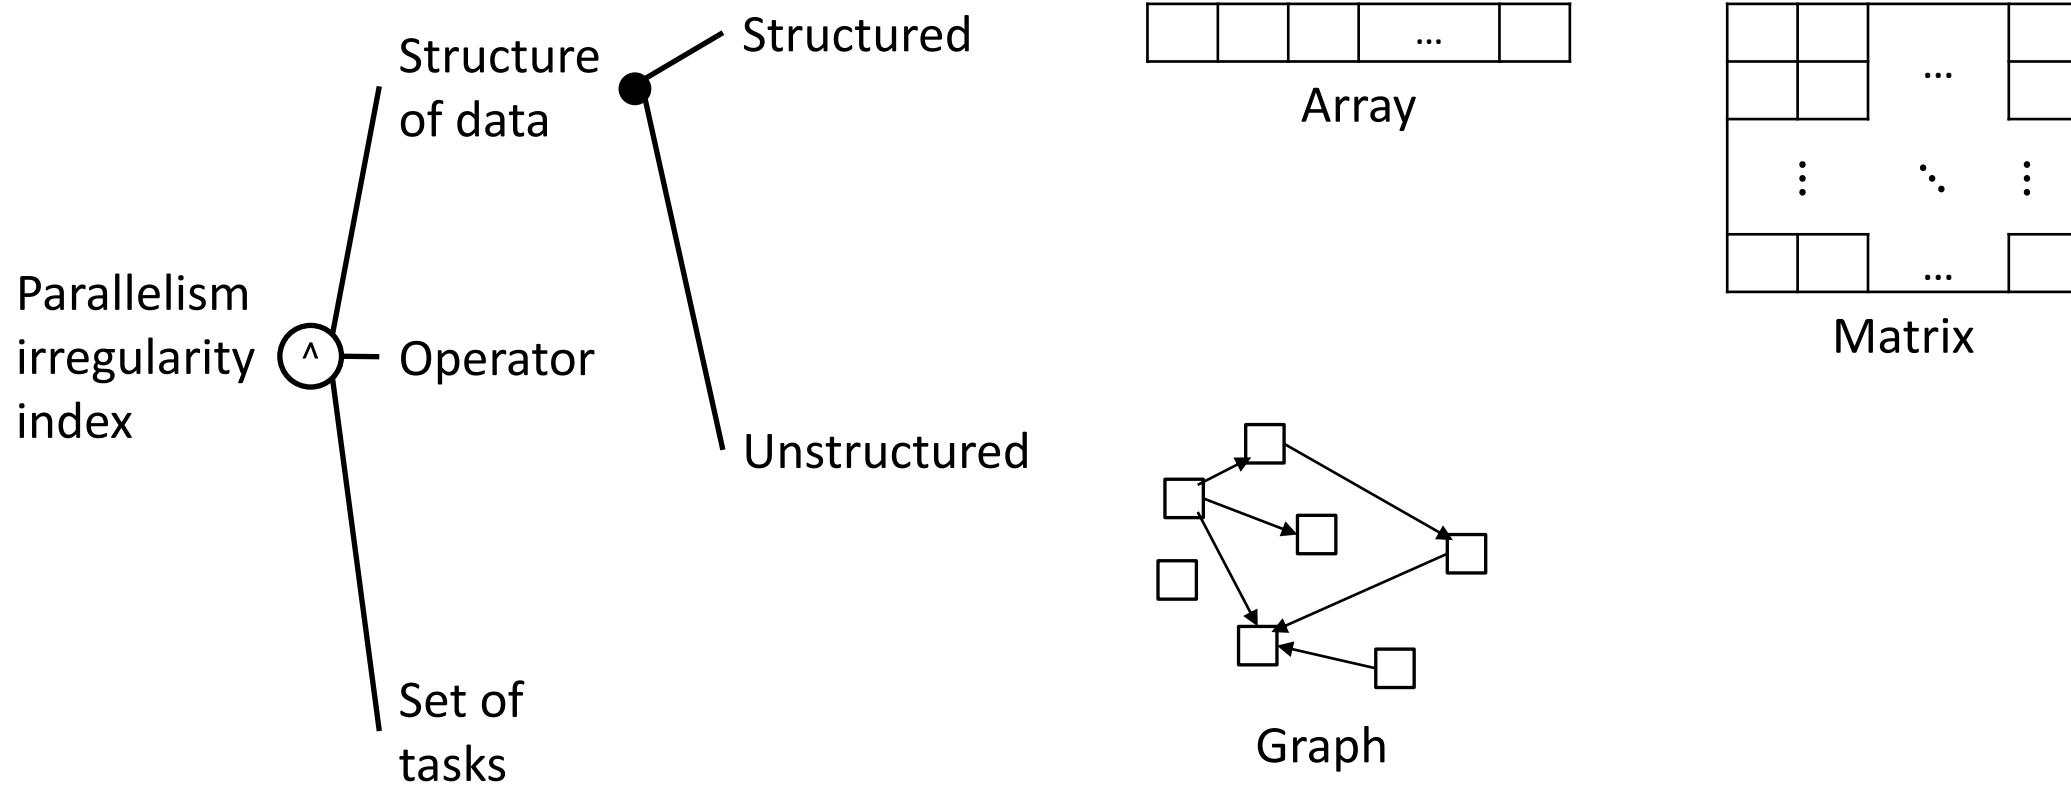
# Background

---

Types of parallelism

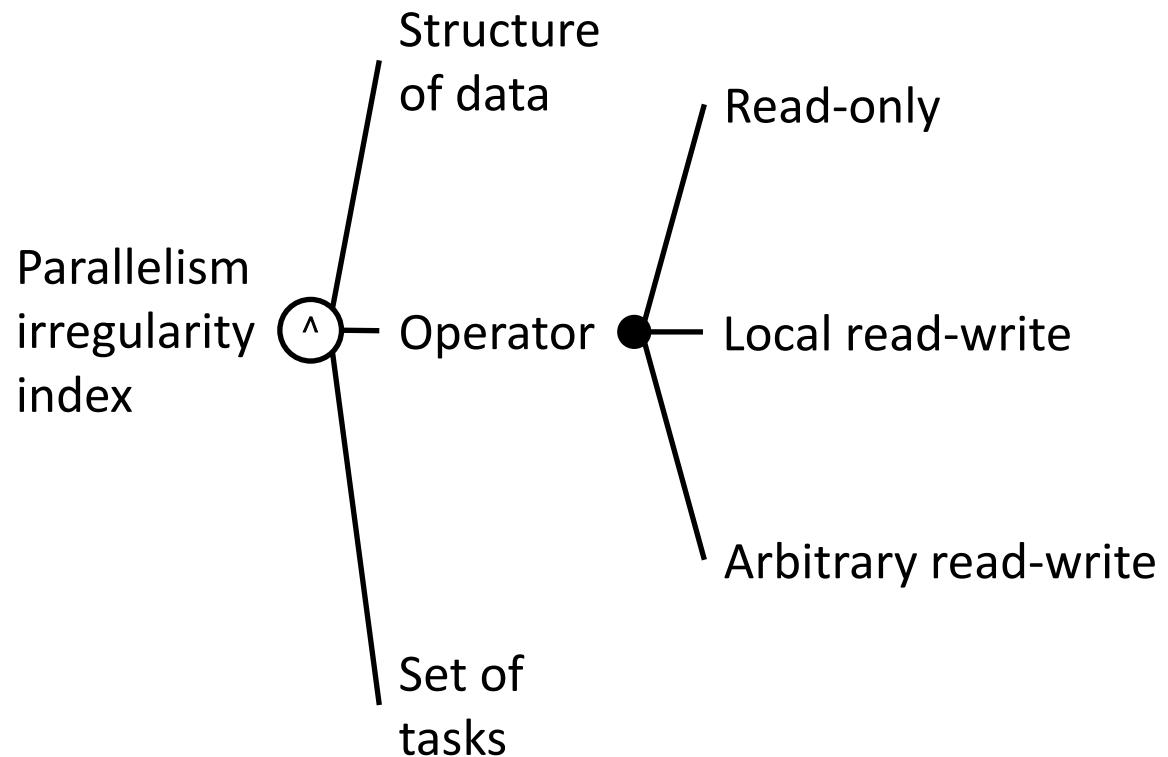
# The regularity of algorithms affects ease of parallelization

[Pingali et al., PLDI 2011]



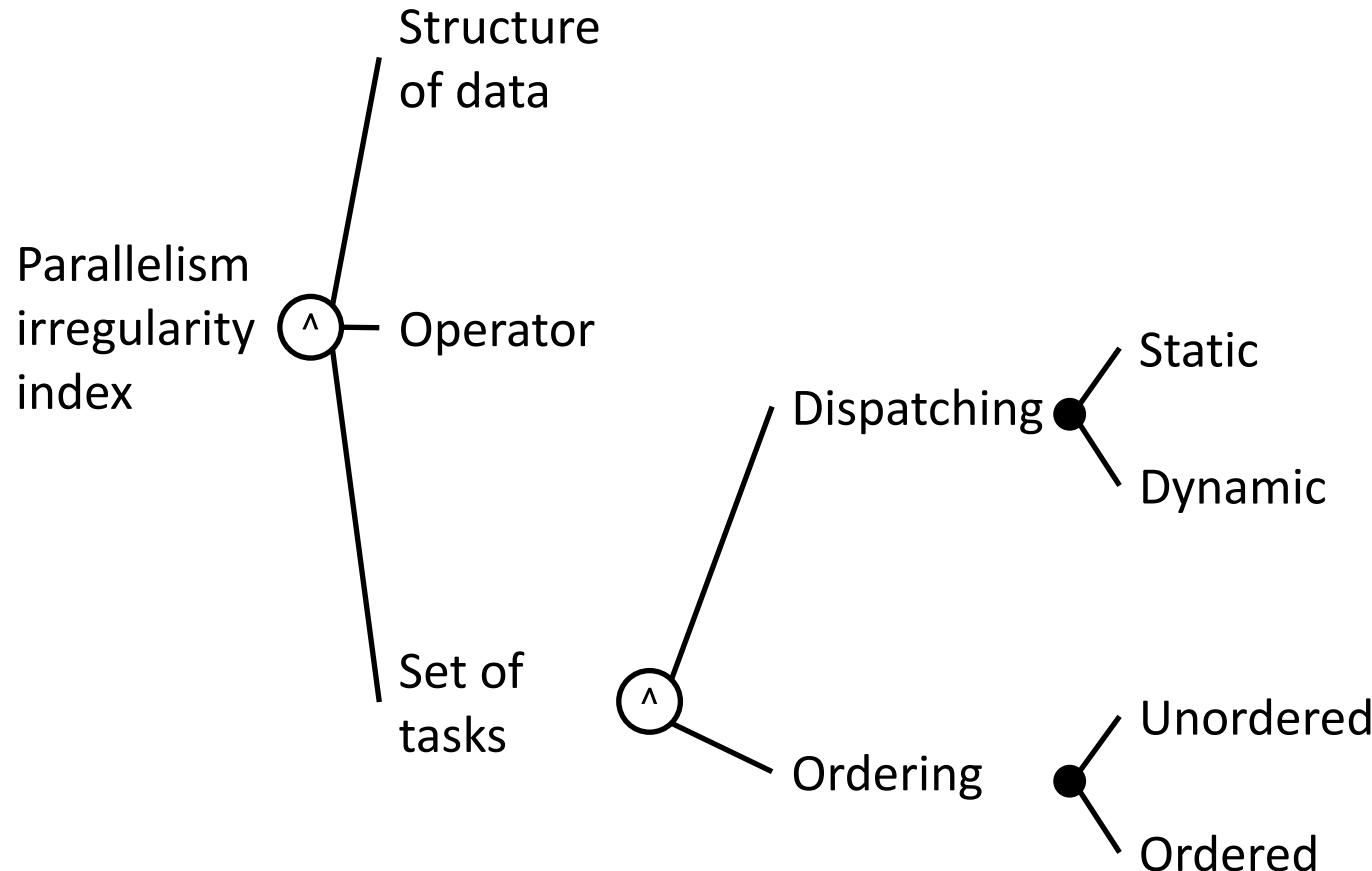
# The regularity of algorithms affects ease of parallelization

[Pingali et al., PLDI 2011]



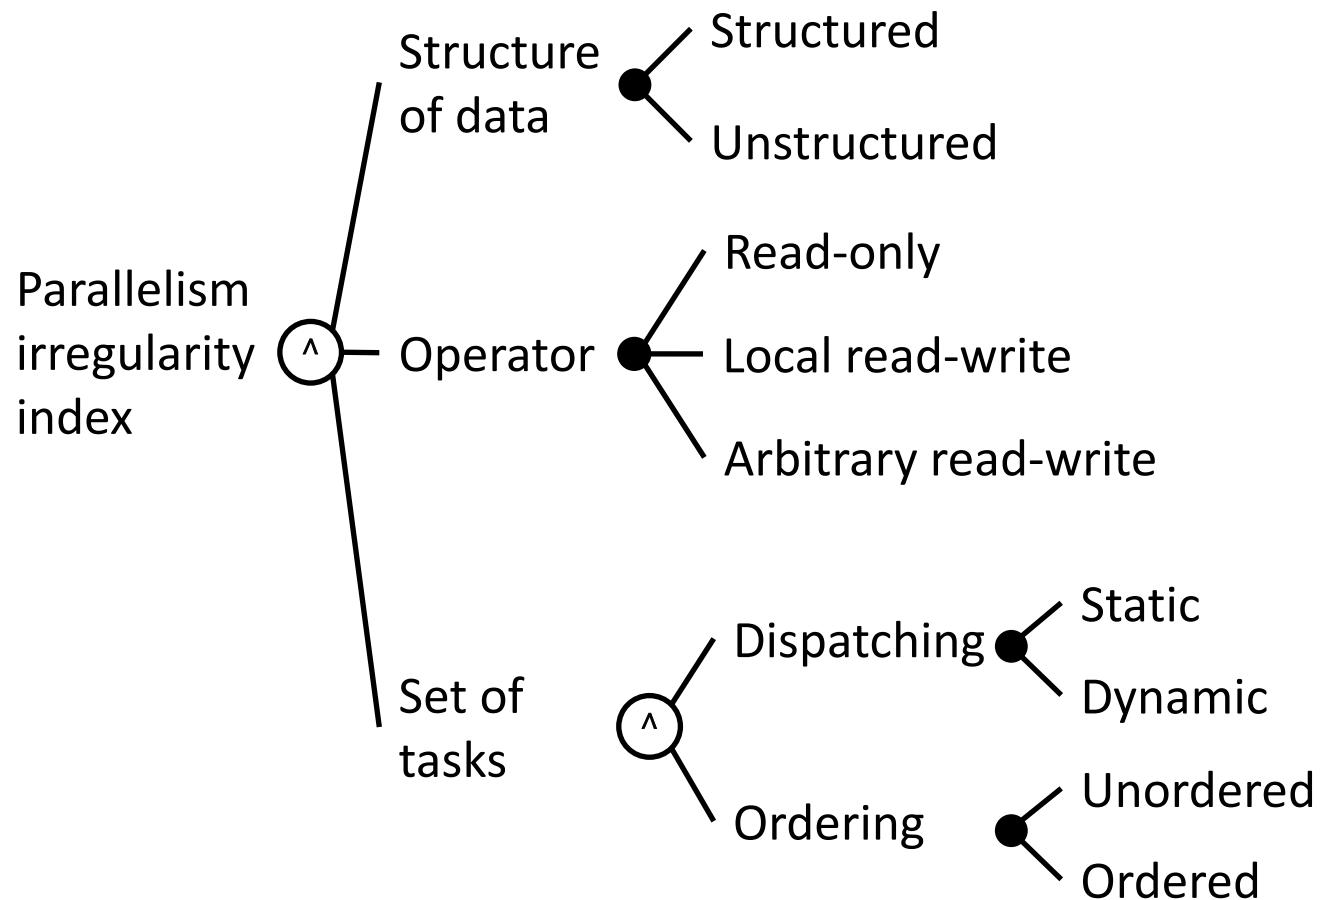
# The regularity of algorithms affects ease of parallelization

[Pingali et al., PLDI 2011]



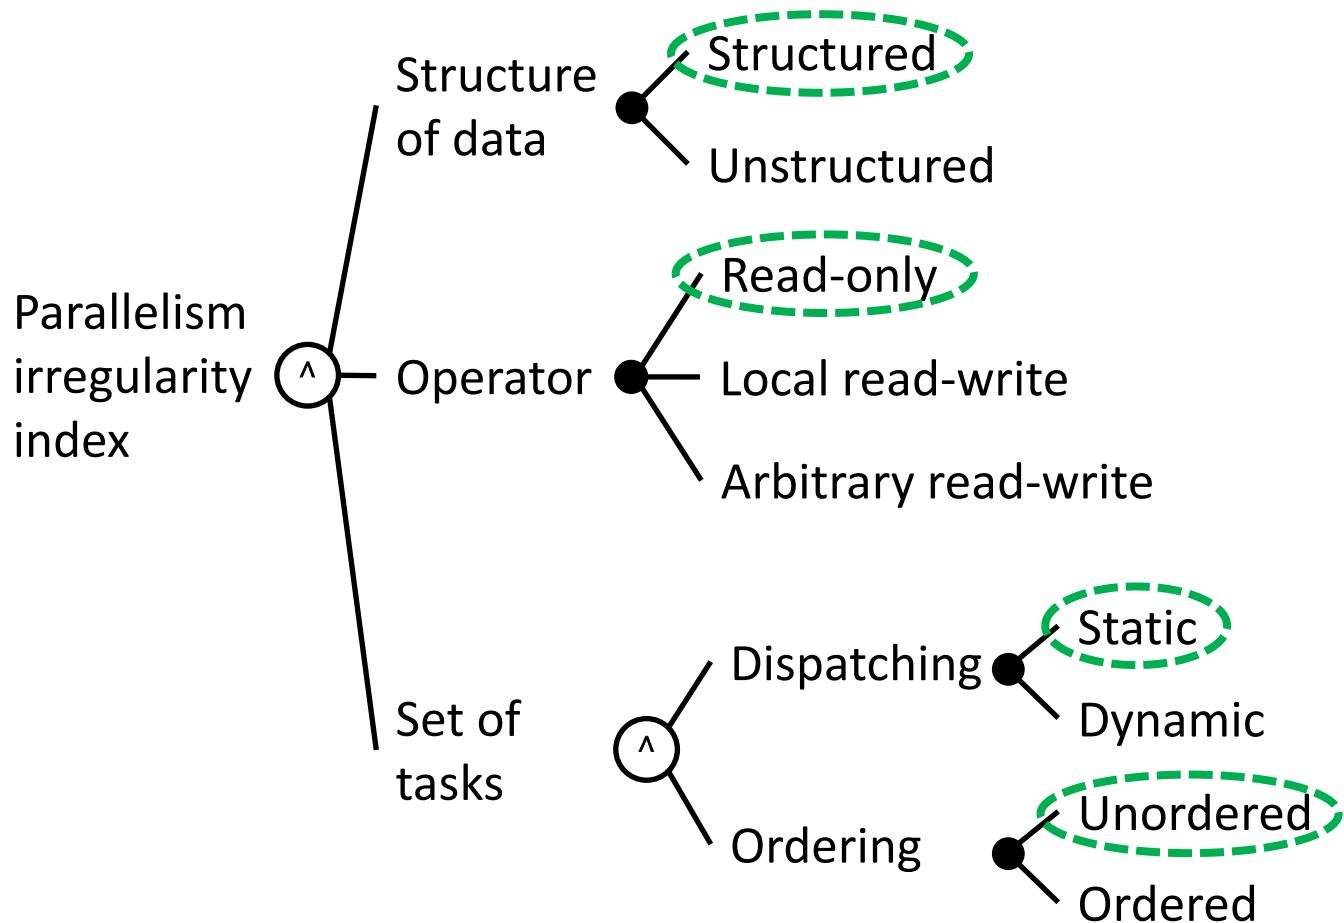
# The regularity of algorithms affects ease of parallelization

[Pingali et al., PLDI 2011]



# The regularity of algorithms affects ease of parallelization

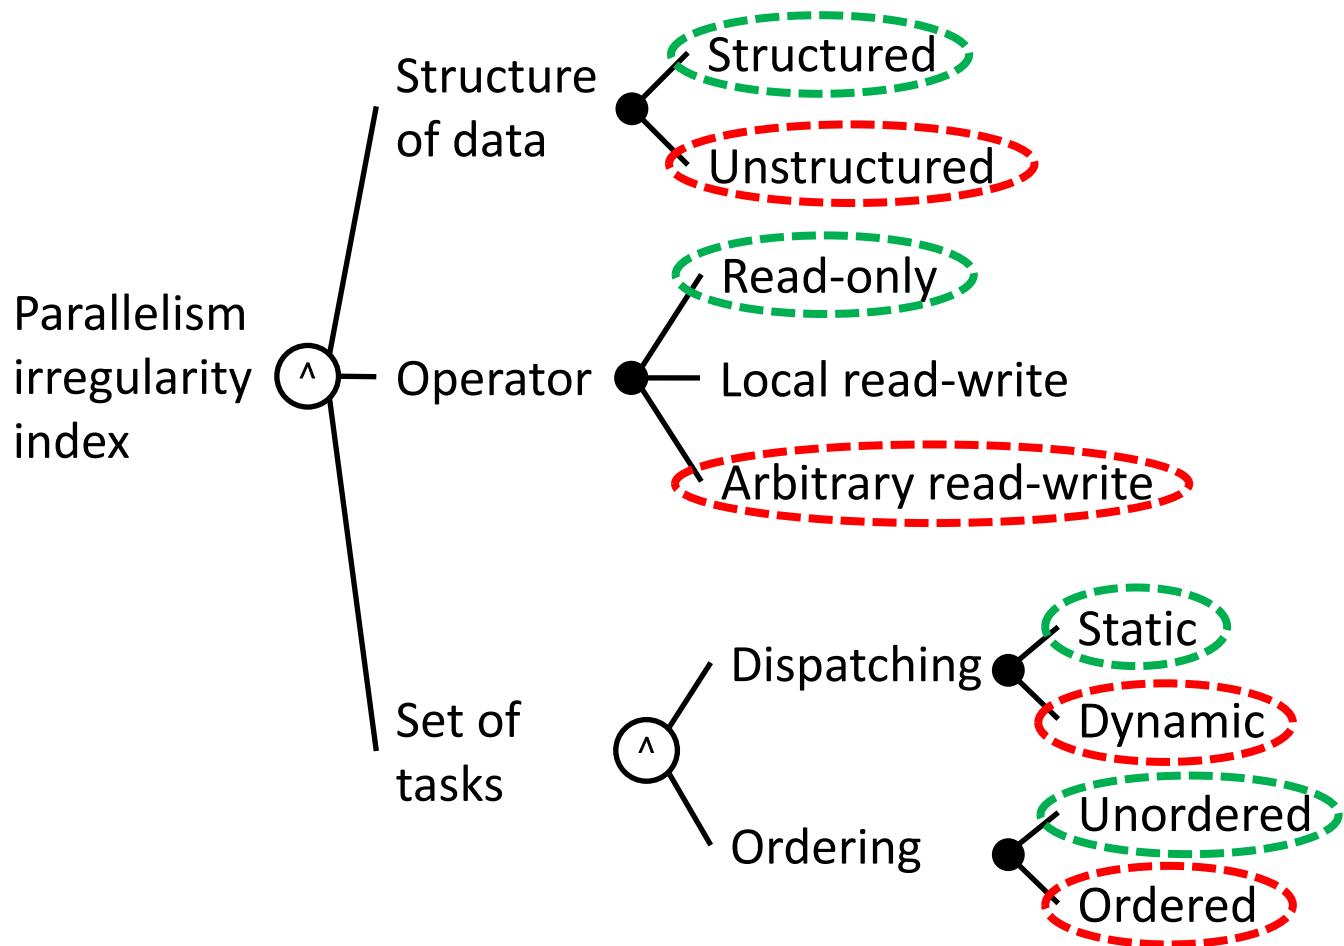
[Pingali et al., PLDI 2011]



Parallel array summation

# The regularity of algorithms affects ease of parallelization

[Pingali et al., PLDI 2011]



Parallel array summation

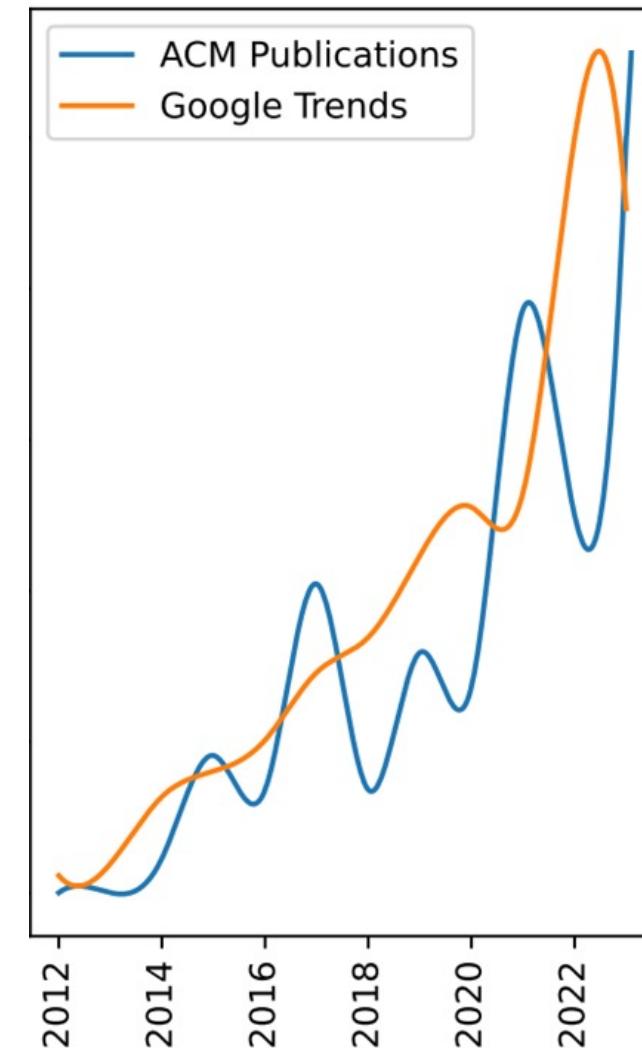
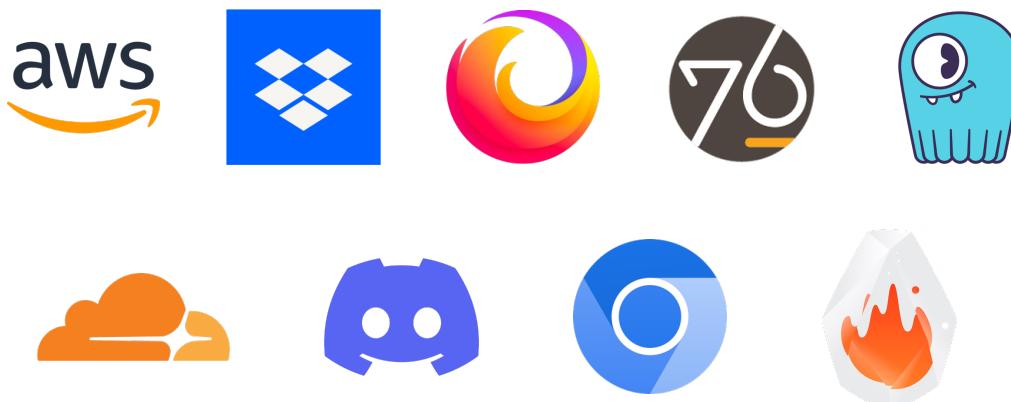
Delta-stepping SSSP

# Background

---

Rust

# Rust is gaining popularity



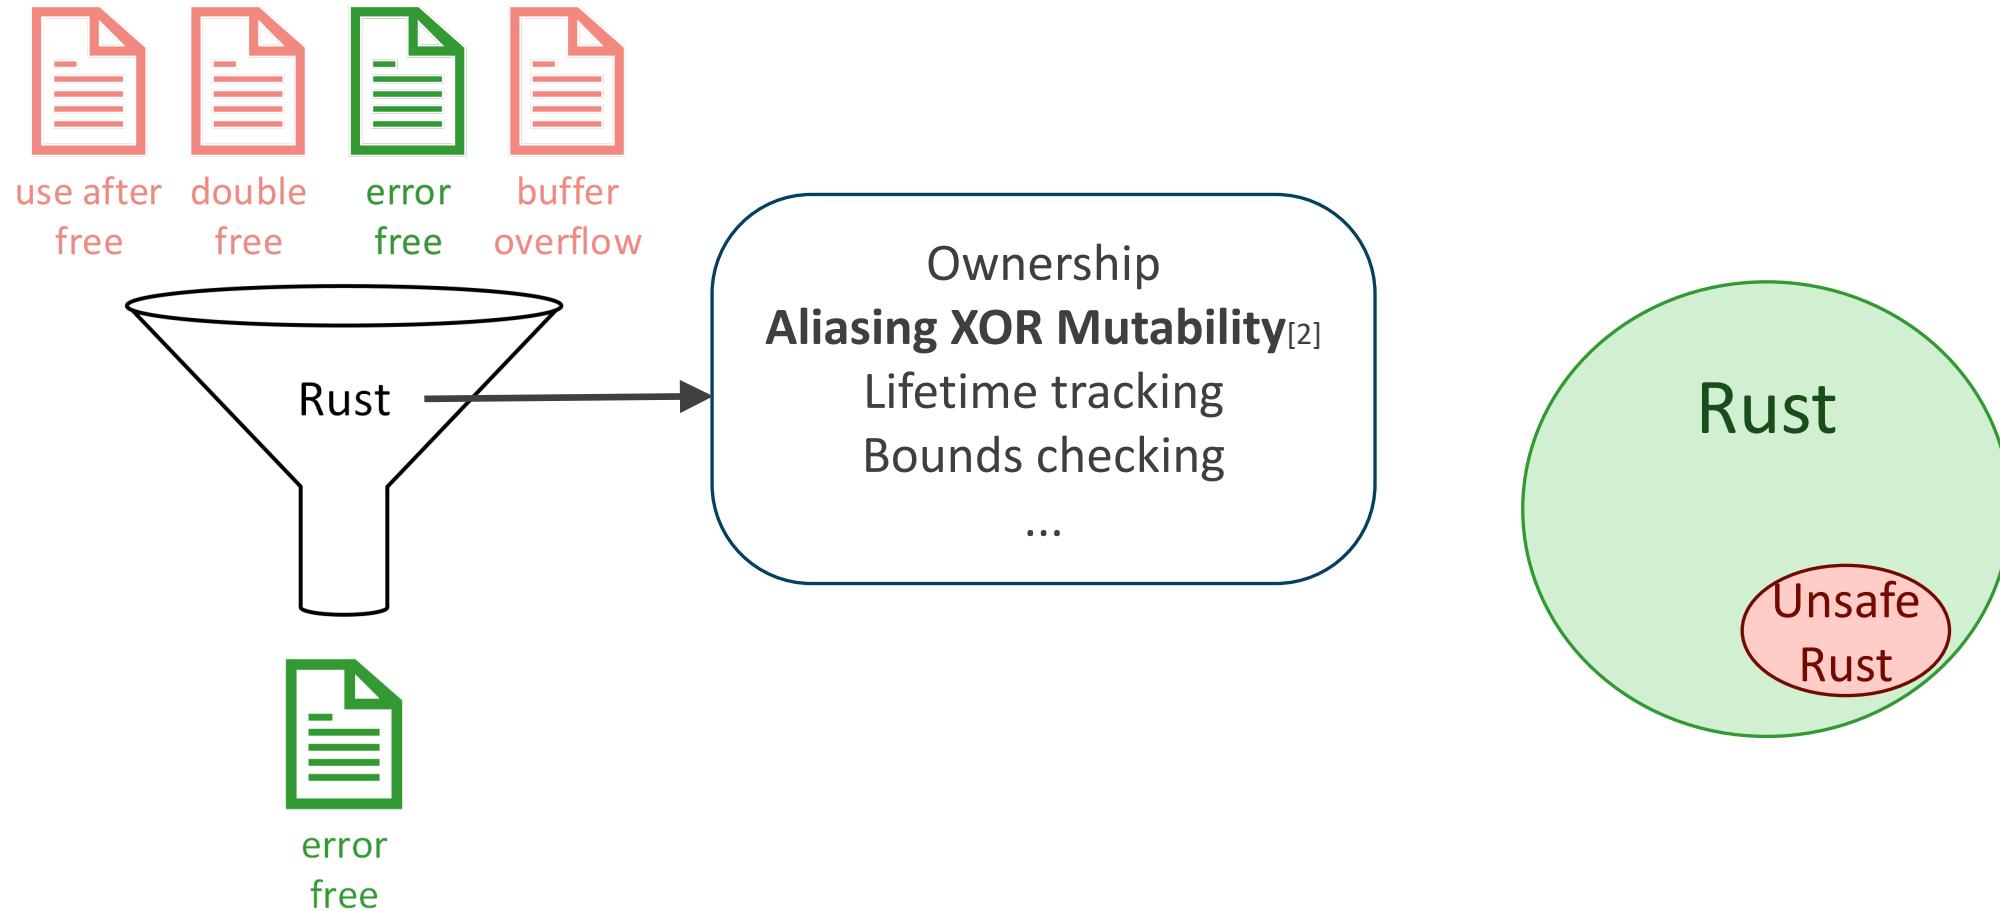
# Rust is gaining popularity

Program with Rust!



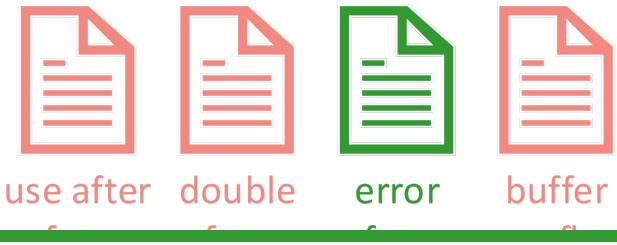
# Rust is gaining popularity because of its safety guarantees

---

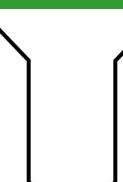


[2] Yanovski et al., ICFP'21, GhostCell: separating permissions from data in Rust.

# Rust is gaining popularity because of its safety guarantees



**Rust catches all type and memory safety errors**



Bounds checking

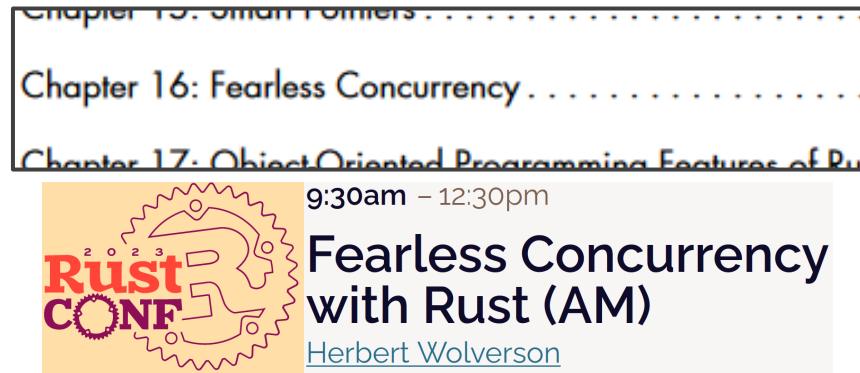
...

Unsafe  
Rust

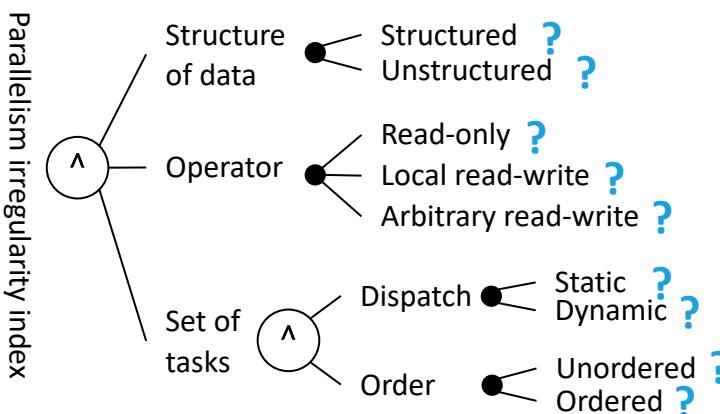
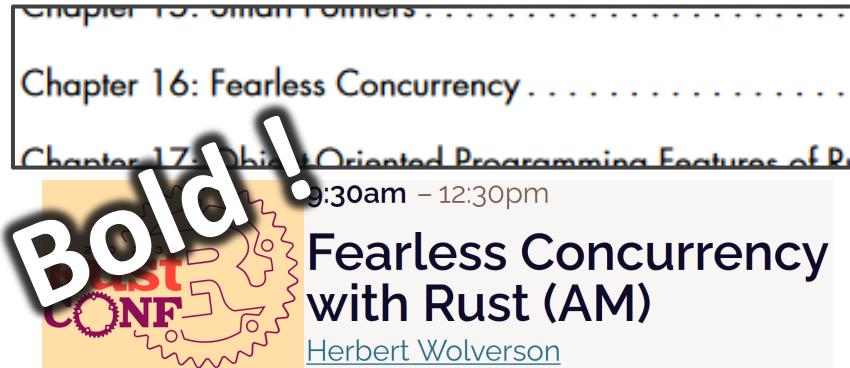
[2] Yanovski et al., ICFP'21, GhostCell: separating permissions from data in Rust.

# Rust claims to enable *fearless concurrency*!!!

---

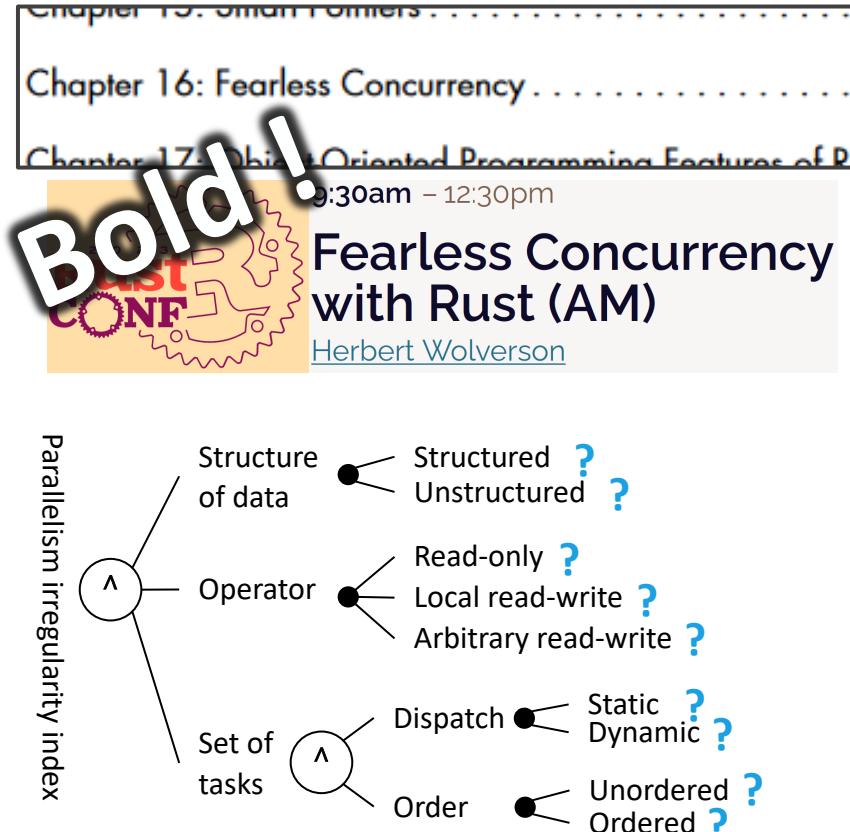


# Rust claims to enable *fearless concurrency*!!!



Does Rust cover all these branches?

# Rust claims to enable *fearless concurrency*!!!

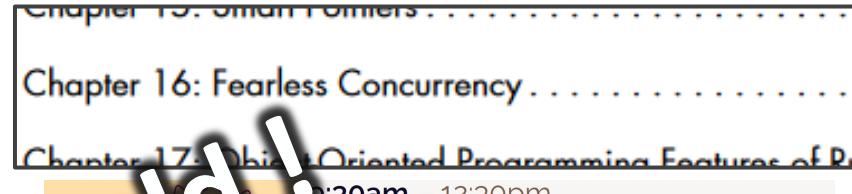


# Does Rust cover all these branches?

tions that *may* conflict, even if they rarely or never do. This is an alternative that can significantly improve the performance and transactional safety of current programs, as well as making them easier to transact, understand, and maintain.

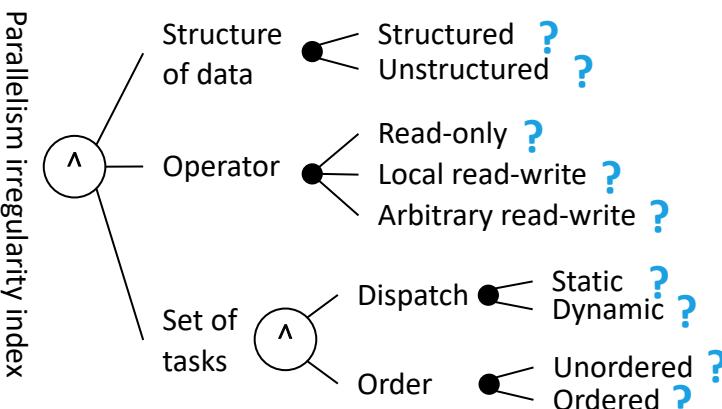
transactional memory (TM) [12] simplifies concurrent programming by providing atomic execution for a block of code. A TM database. The main goal is to make concurrent applications, which is an alternative to traditional sequential programming by providing atomic execution for making concurrent modifications to shared memory (TM) [13, 15]. This is a significant improvement of concurrent programming because it allows for more widespread use of concurrent programming. The Transactional Memory paradigm has raised a lot of hope for mastering the art of concurrent programming. The aim is to provide the developer with a simple and intuitive way to handle concurrency in their applications.

# Rust claims to enable *fearless concurrency*!!!

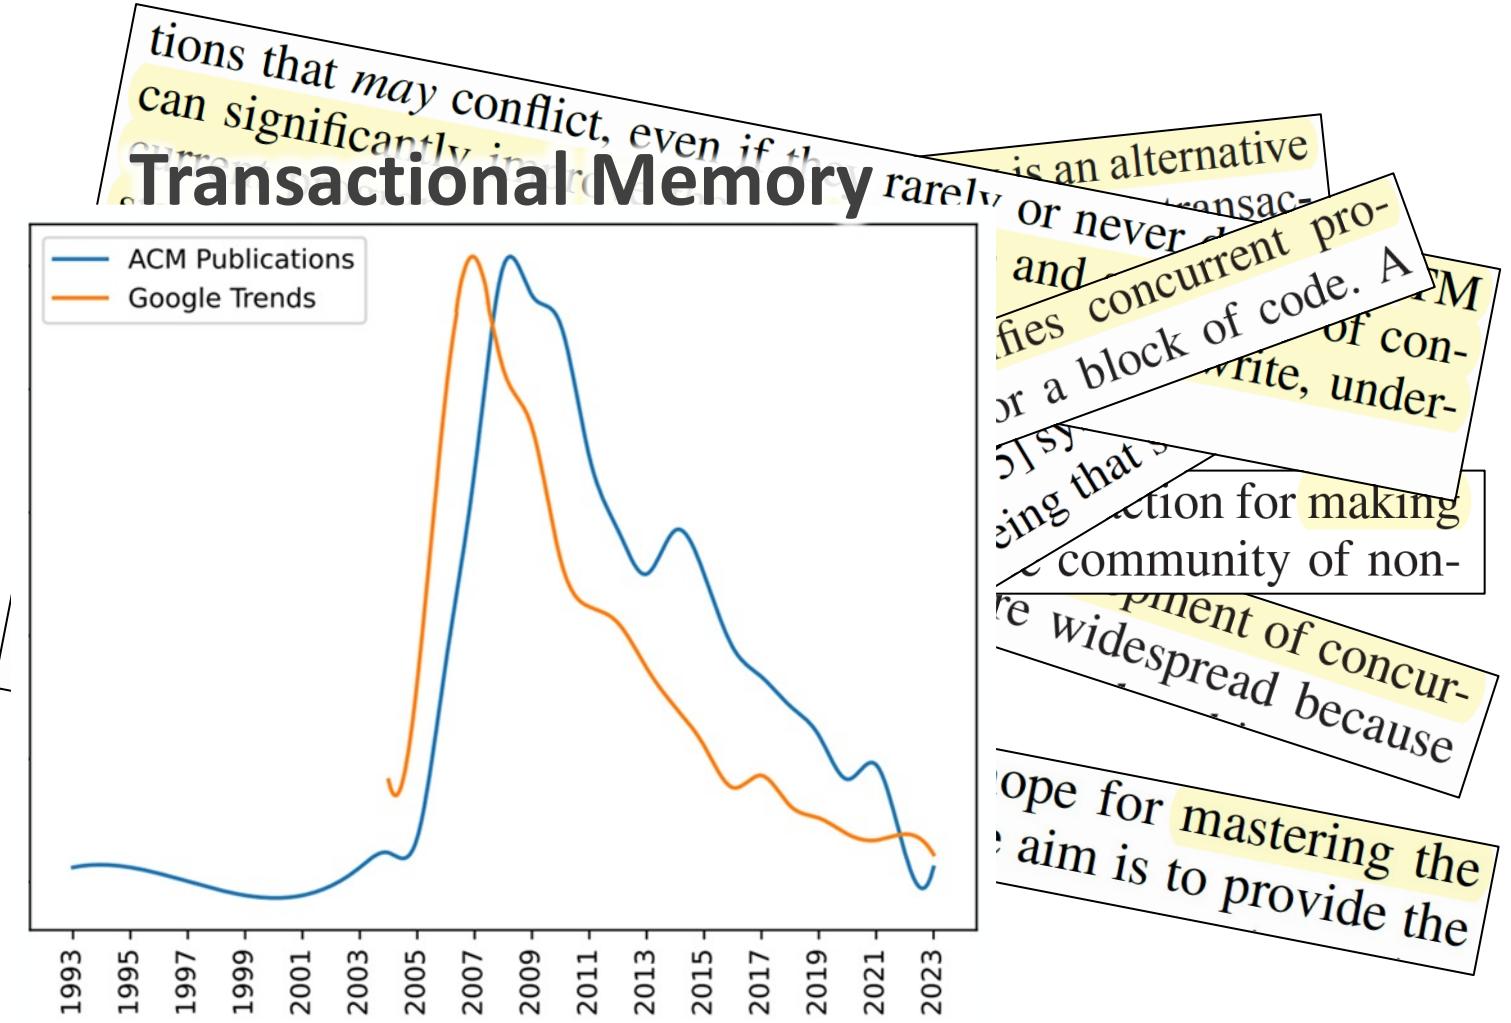


## BOLD! Fearless Concurrency with Rust (AM)

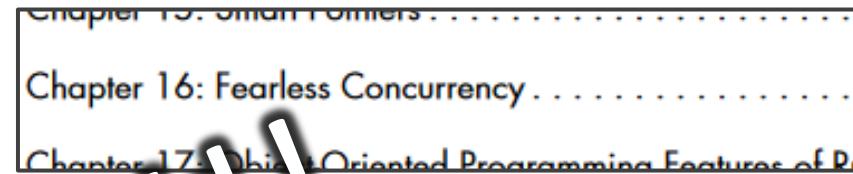
[Herbert Wolverson](#)



Does Rust cover all these branches?



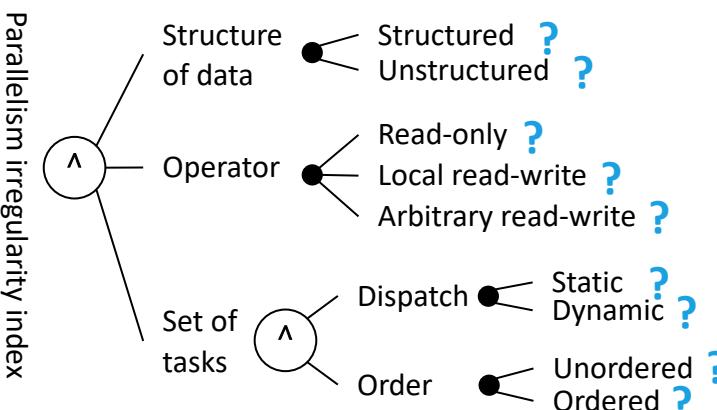
# Rust claims to enable *fearless concurrency*!!!



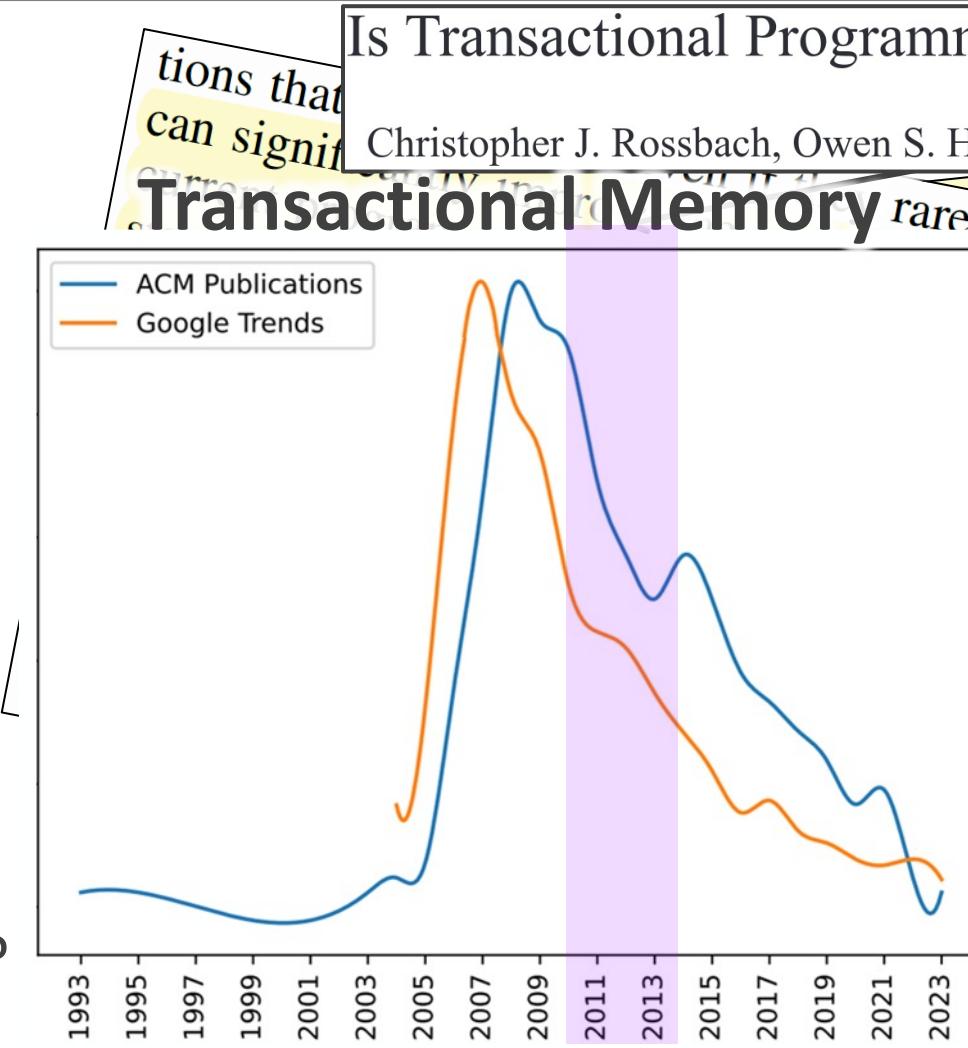
# **Bold!**

## Fearless Concurrency with Rust (AM)

## Herbert Wolverson



# Does Rust cover all these branches?



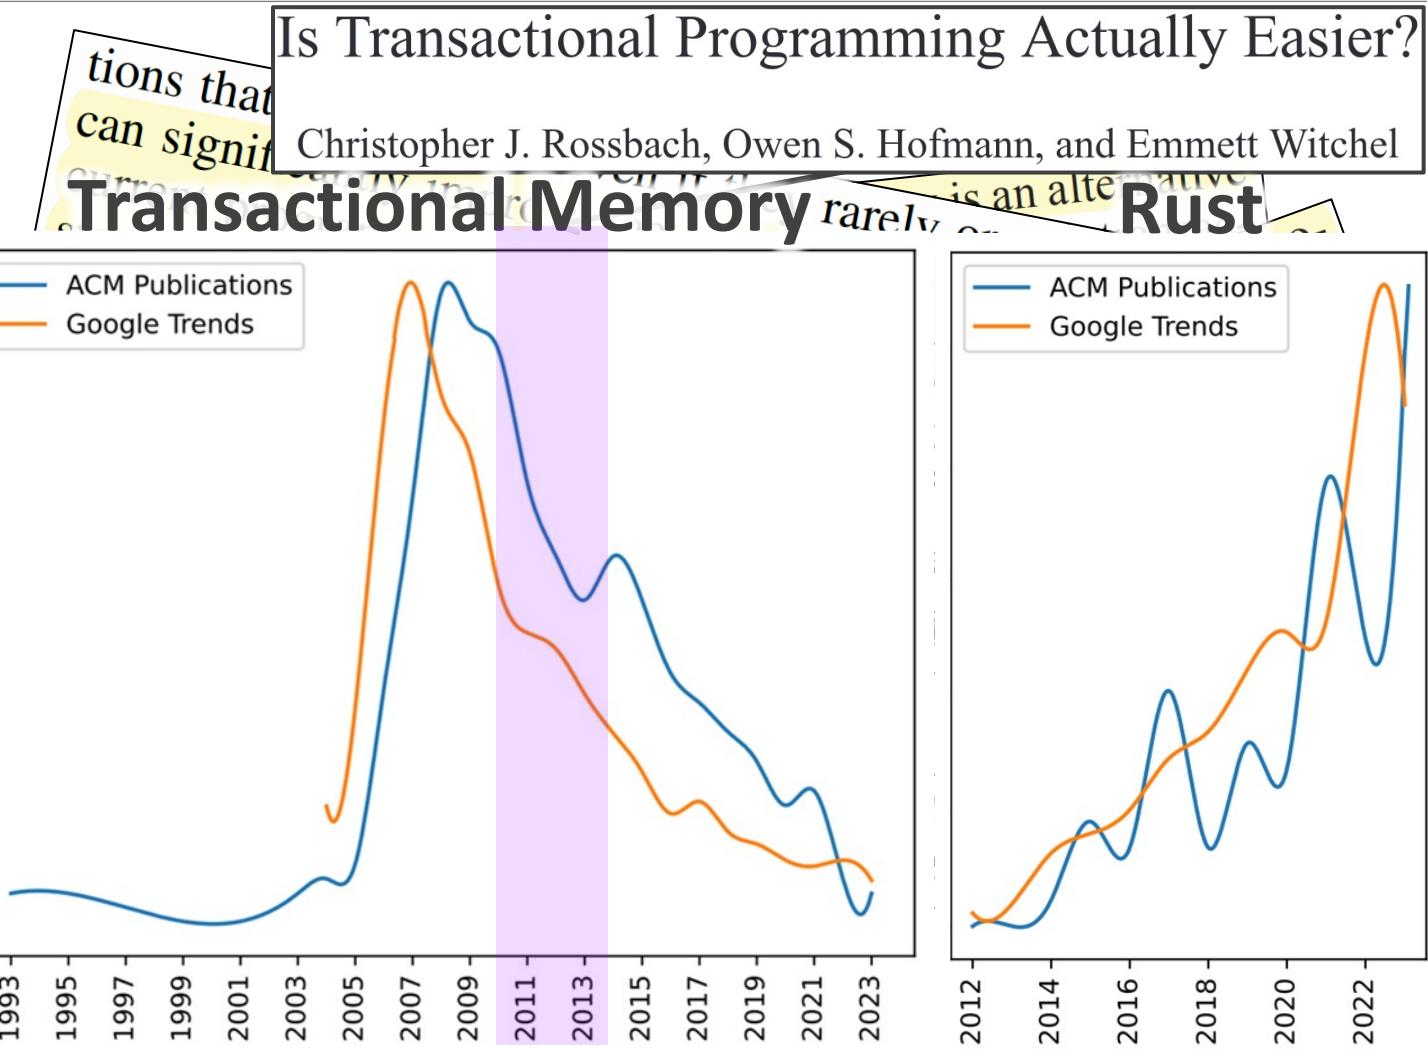
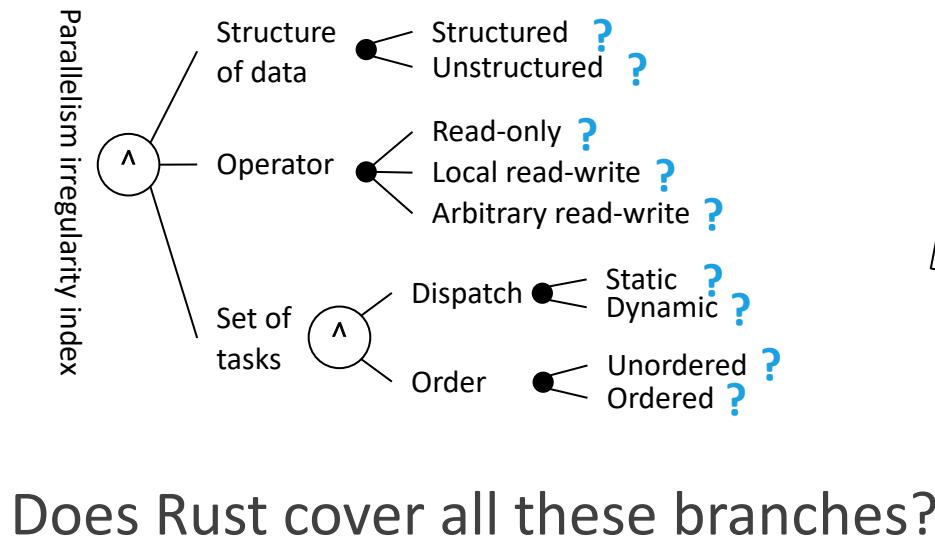
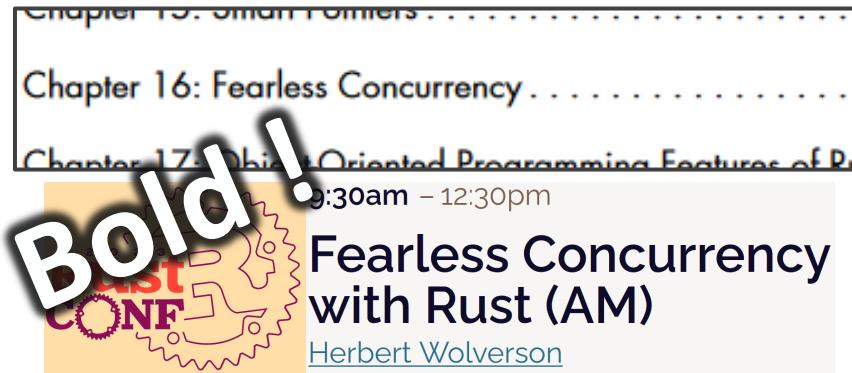
## Is Transactional Programming Actually Easier?

Christopher J. Rossbach, Owen S. Hofmann, and Emmett Witchel

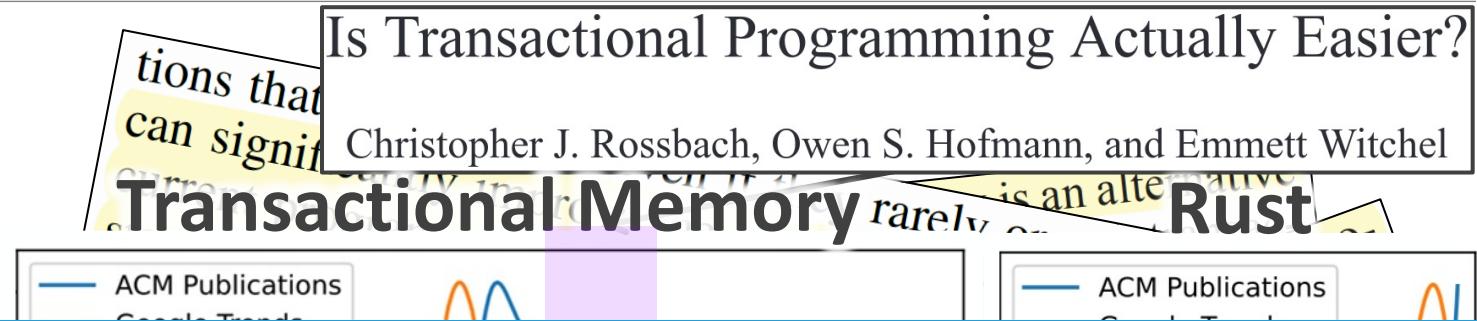
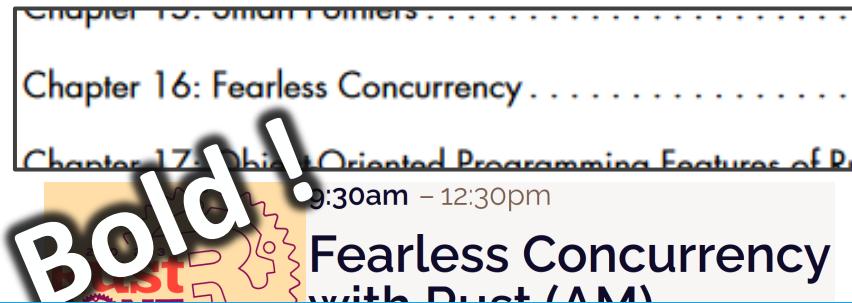
# Transactional Memory

— ACM Publication  
— Google Trends

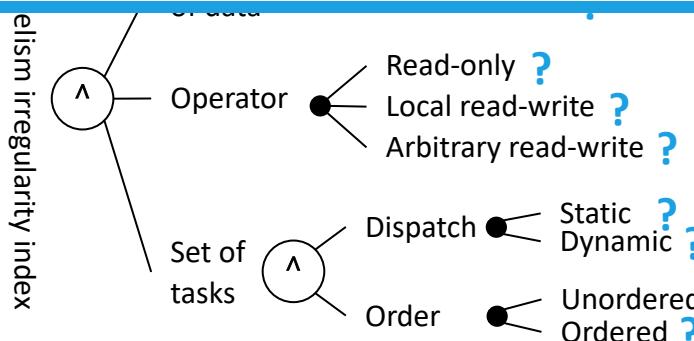
# Rust claims to enable *fearless concurrency*!!!



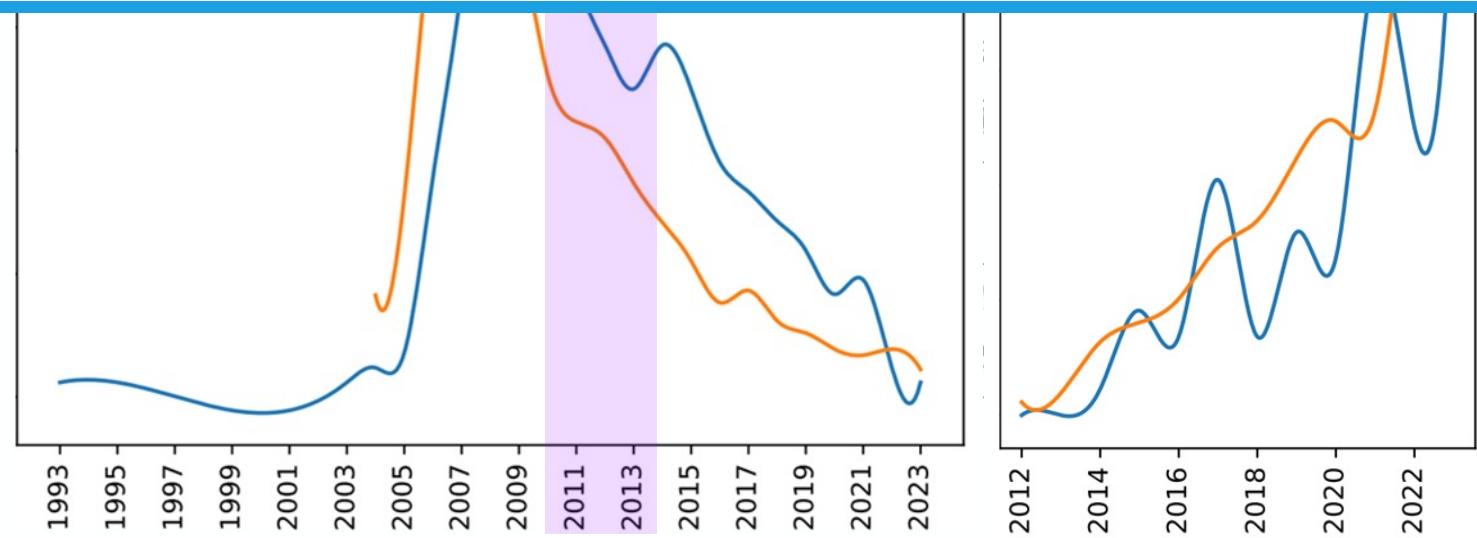
# Rust claims to enable *fearless concurrency*!!!



**It's time to study Rust's capabilities and limitations for parallelism**



Does Rust cover all these branches?



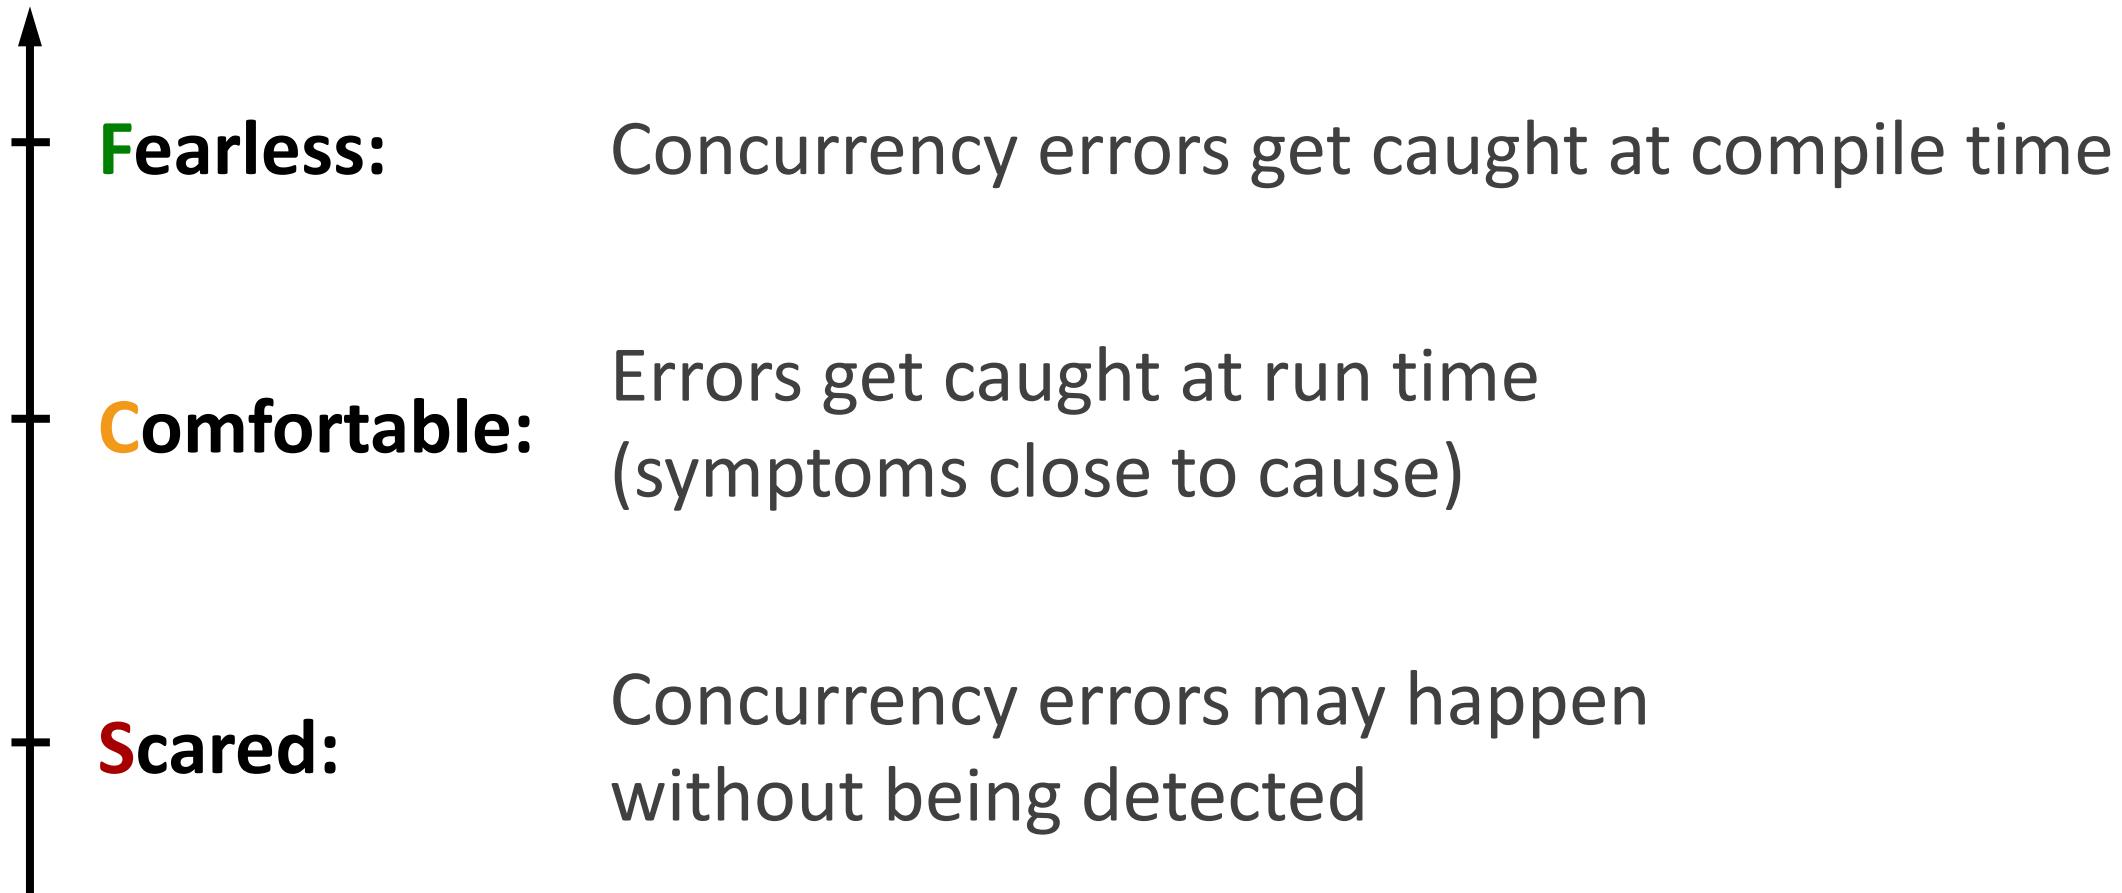
# Case study

---

# Clarifying “fearless concurrency”

---

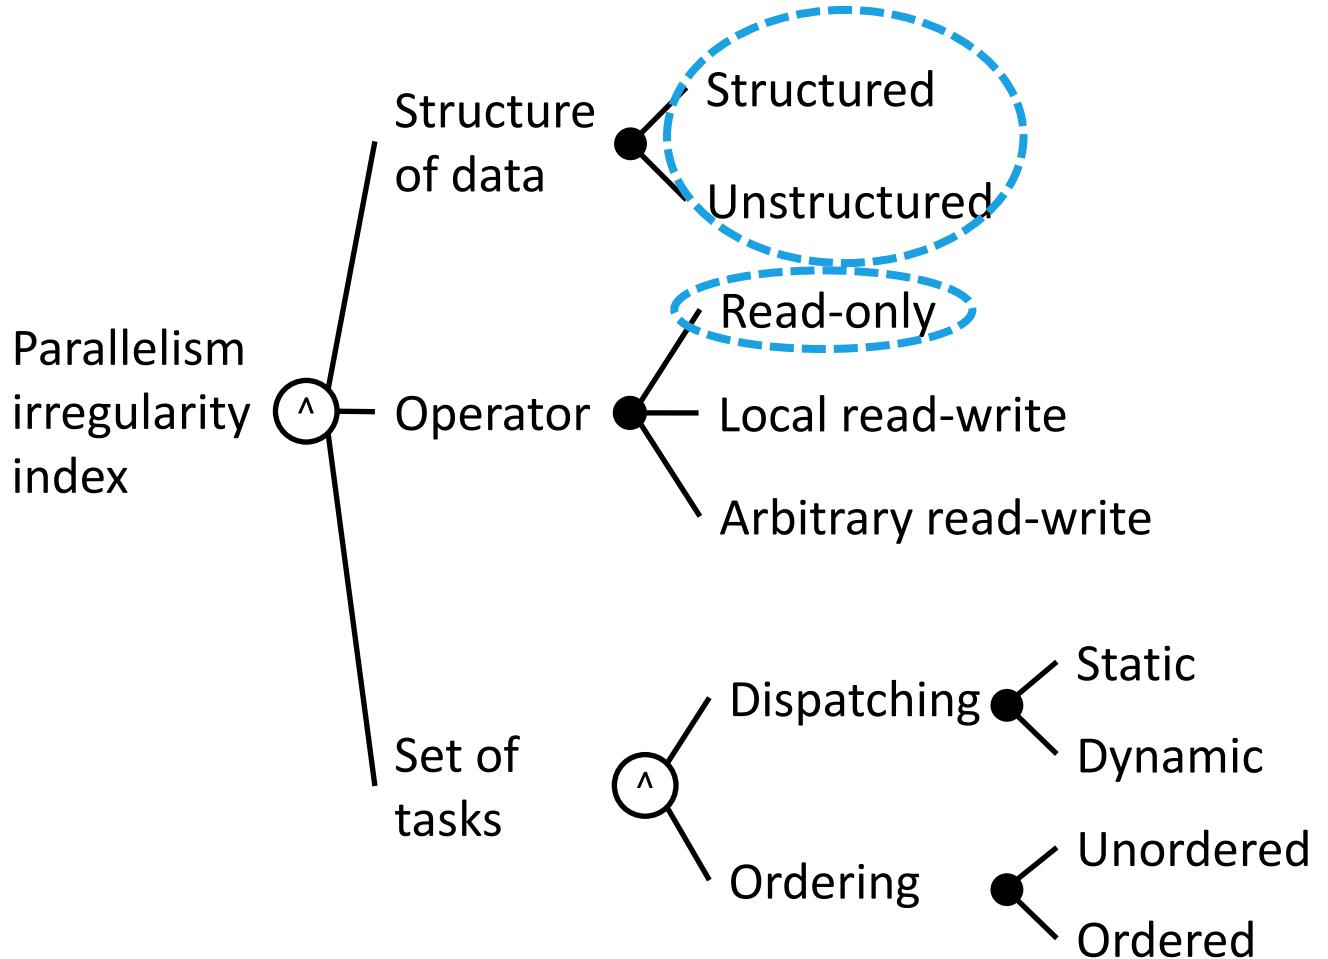
**Fear** : Anticipation of concurrency errors that manifest at run time.



# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

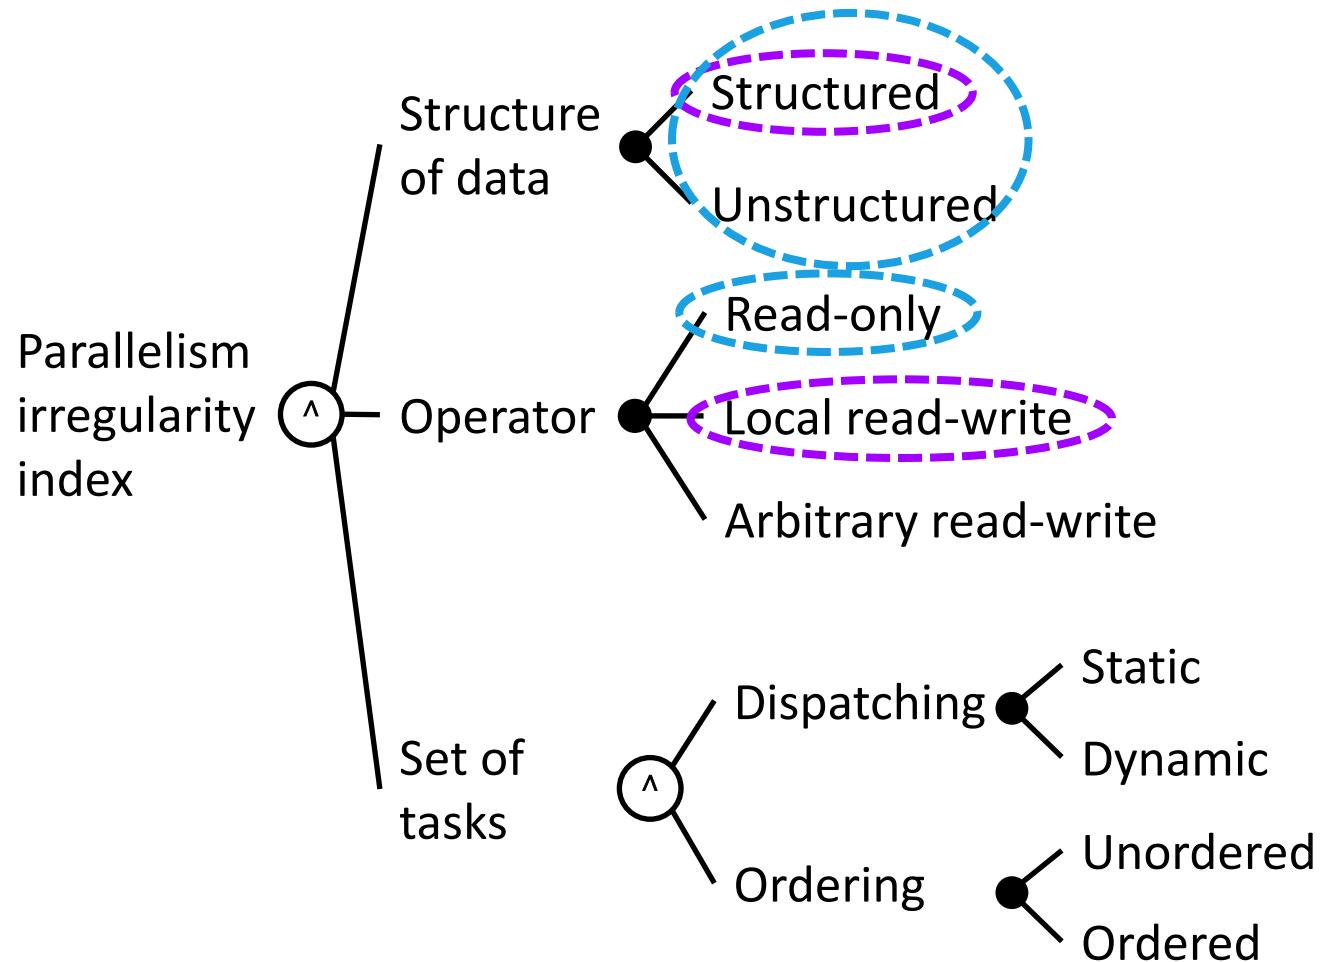
- **Read-only accesses**



# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

- Read-only accesses
- Local read-writes on structured data

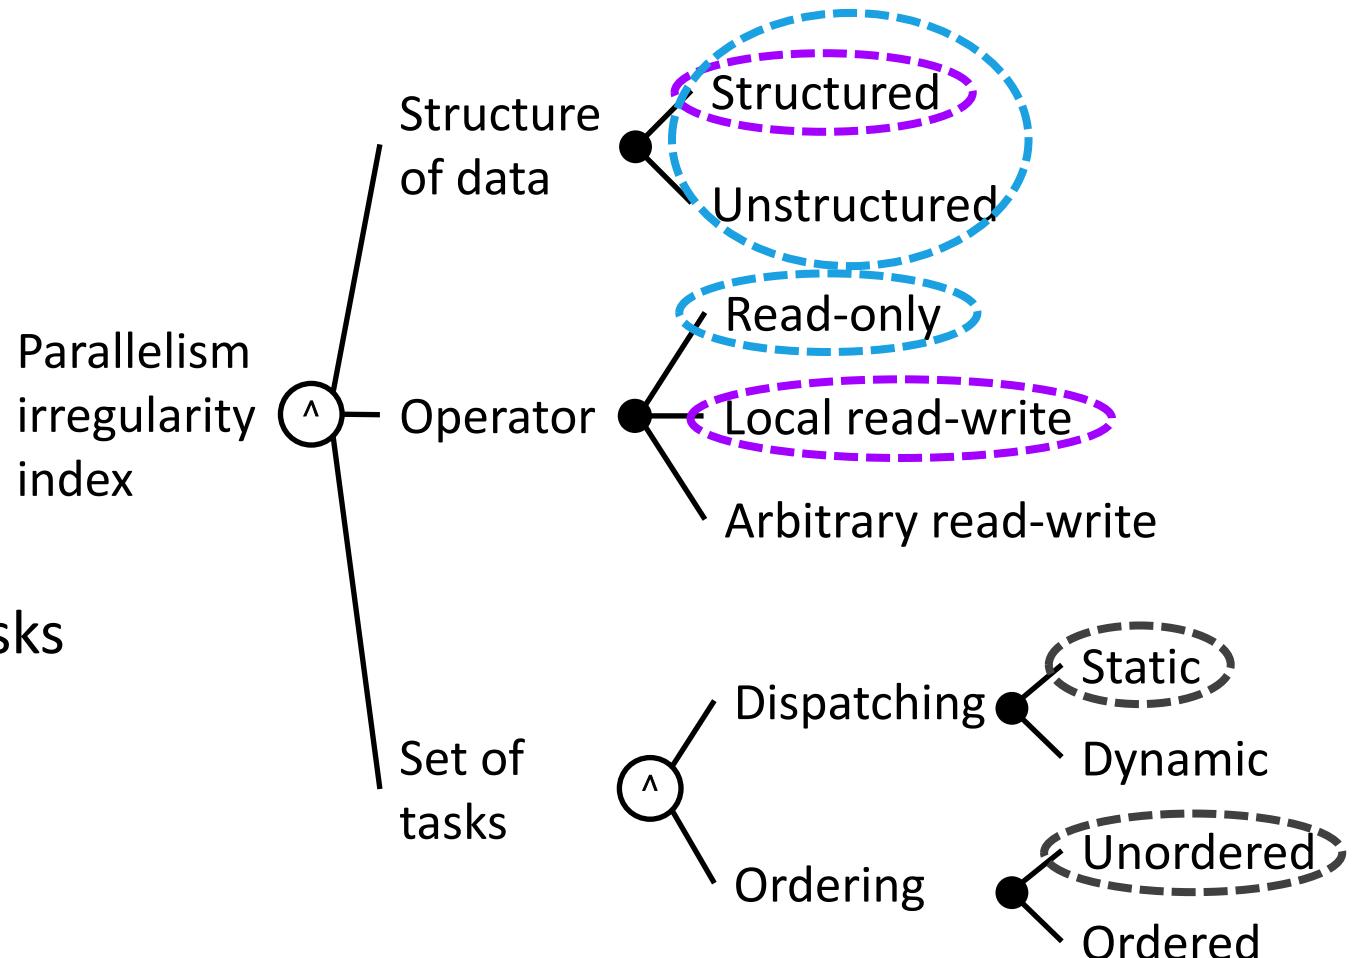


# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

- Read-only accesses
- Local read-writes on structured data

Assume a static unordered set of tasks



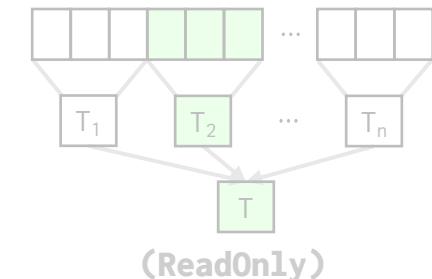
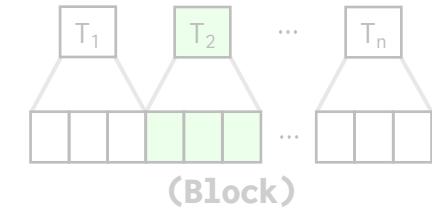
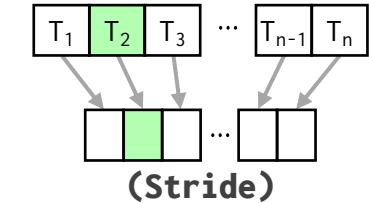
# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

Read-only accesses

Local read-writes on structured data

```
fn par_increment(v: &mut [u32])  
{  
    v.par_iter_mut() // stride pattern on v  
        .for_each(|vi| *vi+=1);  
}
```



# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

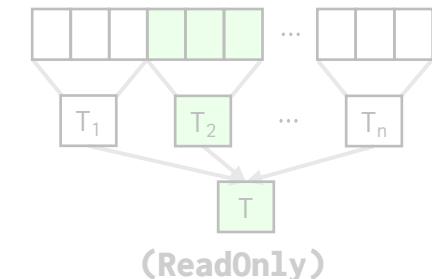
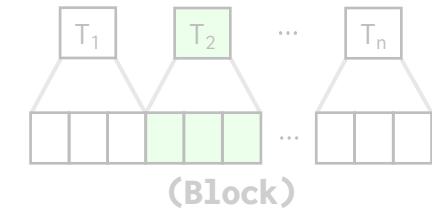
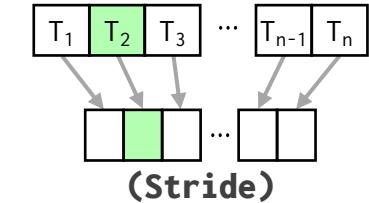
Read-only accesses

Local read-writes on structured data

```
fn par_increment(v: &mut [u32])  
{  
    v.par_iter_mut()  
        .for_each(|vi| *vi+=1);  
}
```

stride pattern on **v**

task



# Expressing tasks with regular accesses in Rust (+Rayon)

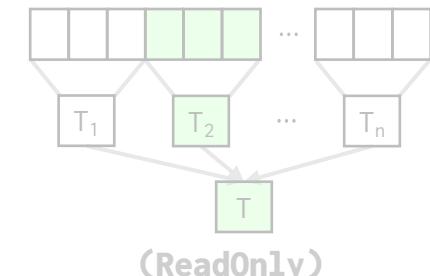
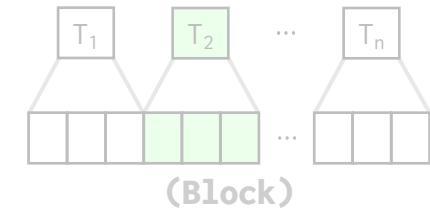
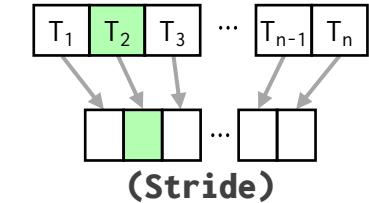
Regular accesses:

Read-only accesses

Local read-writes on structured data

```
fn par_increment(v: &mut [u32])  
{  
    v.par_iter_mut()  
        .for_each(|vi| *vi+=1);  
}
```

: Can only  
mutate vi



# Expressing tasks with regular accesses in Rust (+Rayon)

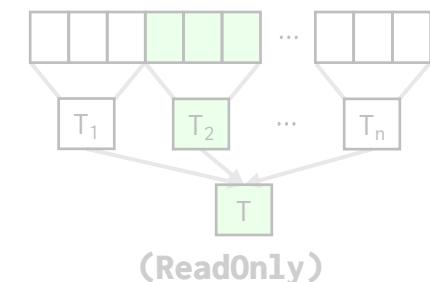
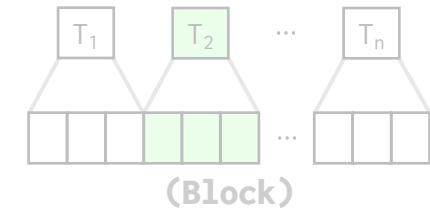
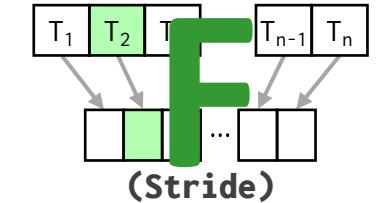
Regular accesses:

Read-only accesses

Local read-writes on structured data

```
fn par_increment(v: &mut [u32])  
{  
    v.par_iter_mut()  
        .for_each(|vi| *vi+=1);  
}
```

stride pattern on v  
task : Can only  
mutate vi → No  
errors



# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

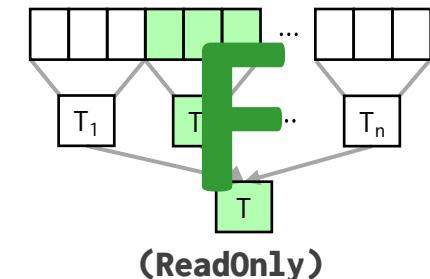
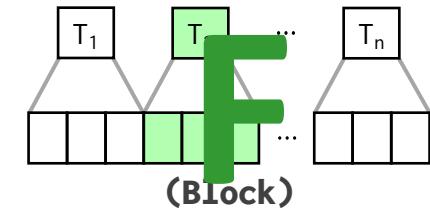
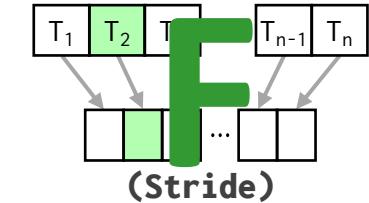
Read-only accesses

Local read-writes on structured data

```
fn par_increment(v: &mut [u32])  
{  
    v.par_iter_mut()  
        .for_each(|vi| *vi+=1);  
}
```

stride pattern on v

task : Can only mutate vi → No errors

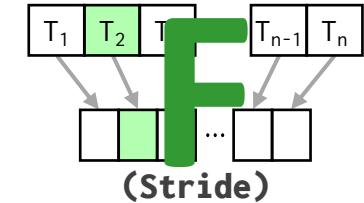


# Expressing tasks with regular accesses in Rust (+Rayon)

Regular accesses:

Read-only accesses

Local read-writes on structured data



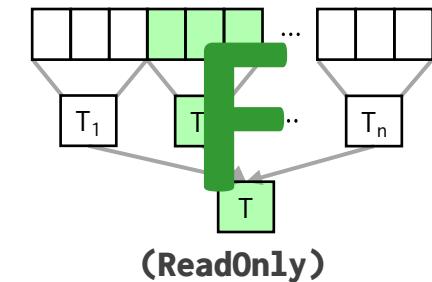
**Rust provides fearlessness for regular accesses**

```
i
v.par_iter_mut()
    .for_each(|vi| *vi+=1);
}
```

stride pattern on v

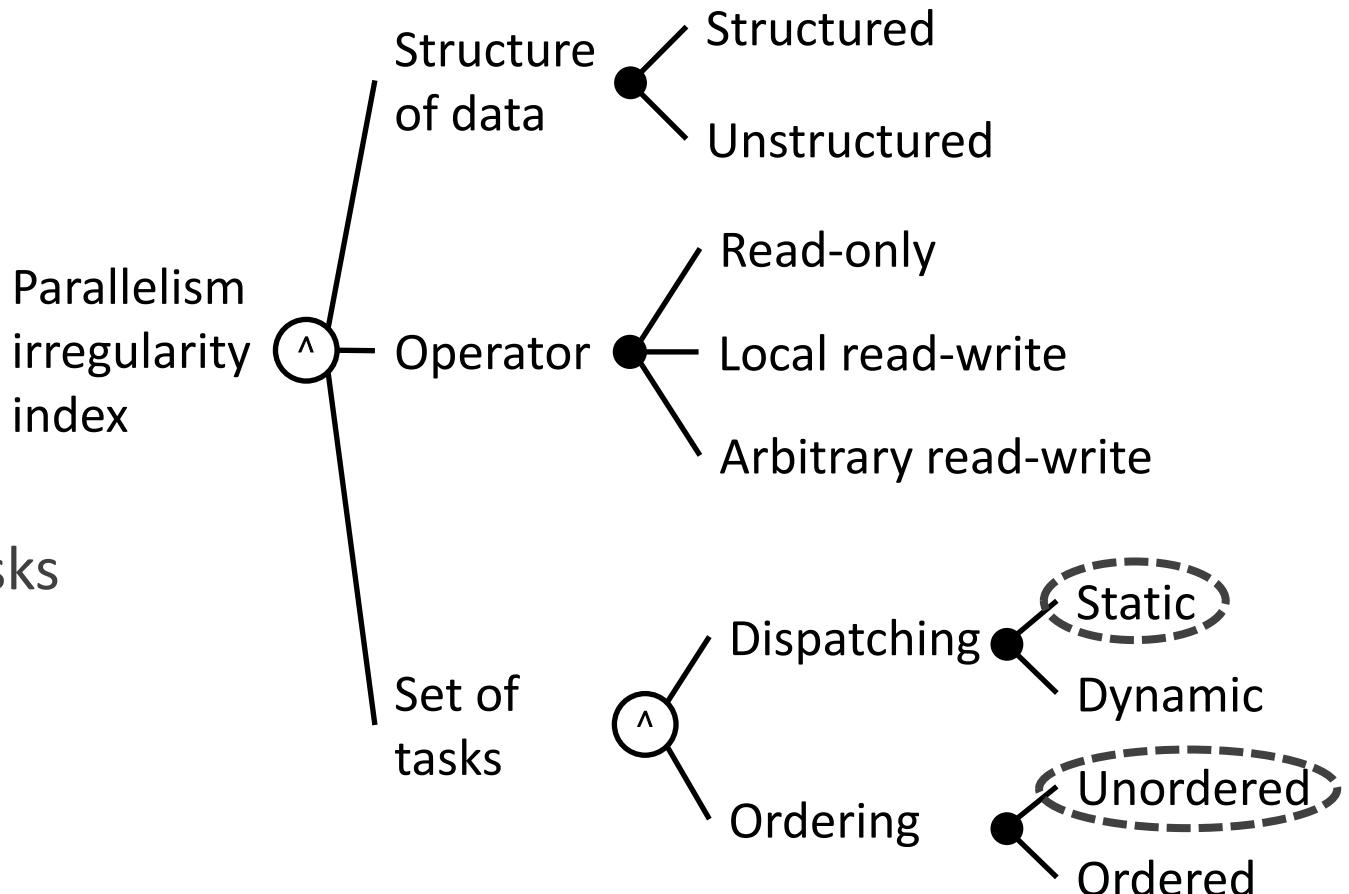
task

: Can only  
mutate vi ➡ No  
errors



# Rust rejects data-race-prone irregular accesses

Irregular accesses:



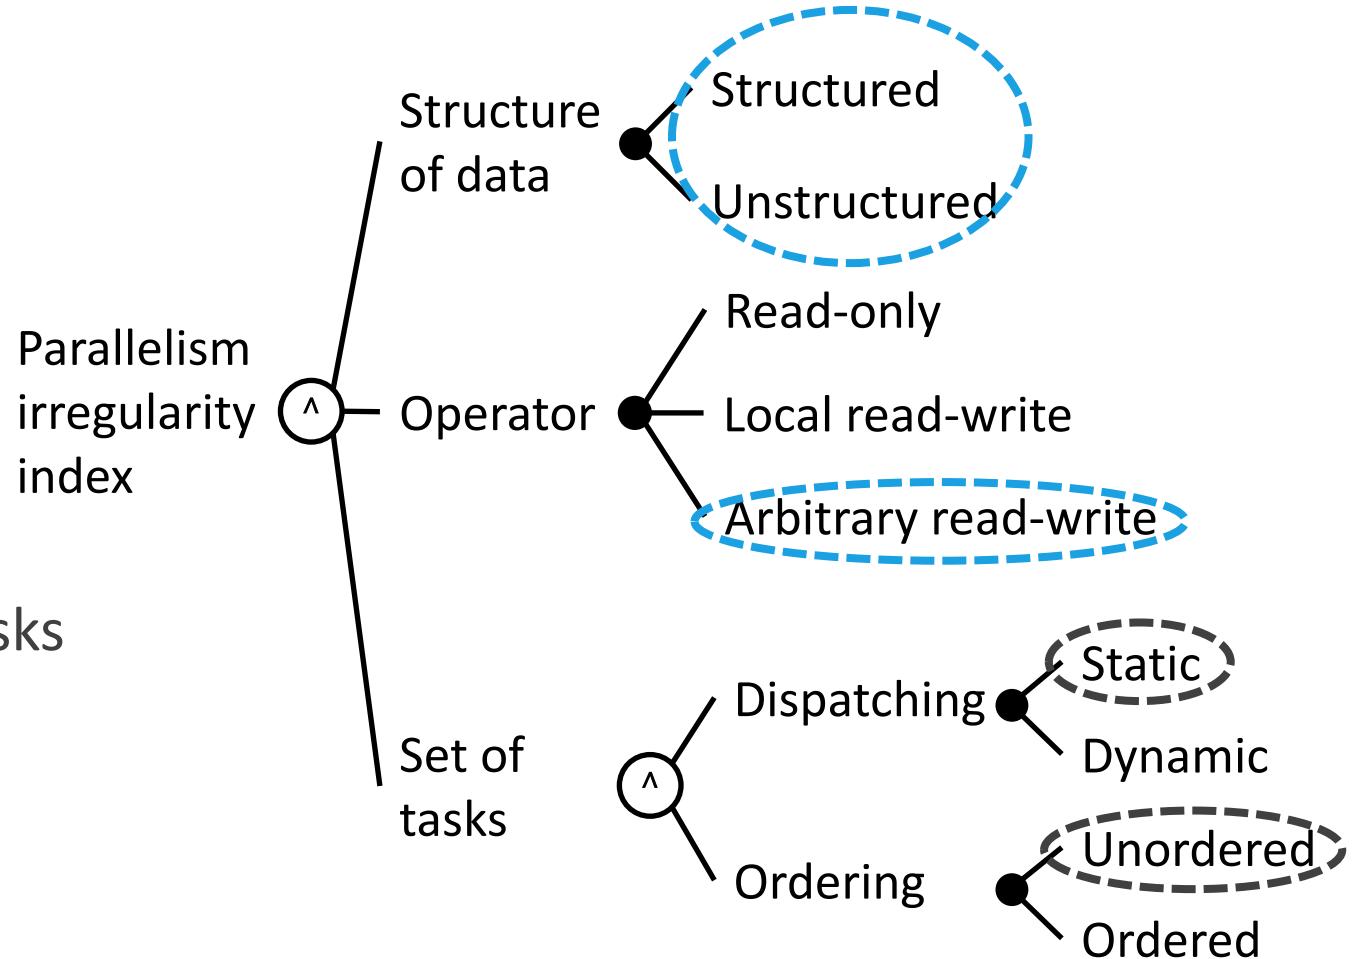
Assume a static unordered set of tasks

# Rust rejects data-race-prone irregular accesses

Irregular accesses:

- **Arbitrary read-writes**

Assume a static unordered set of tasks

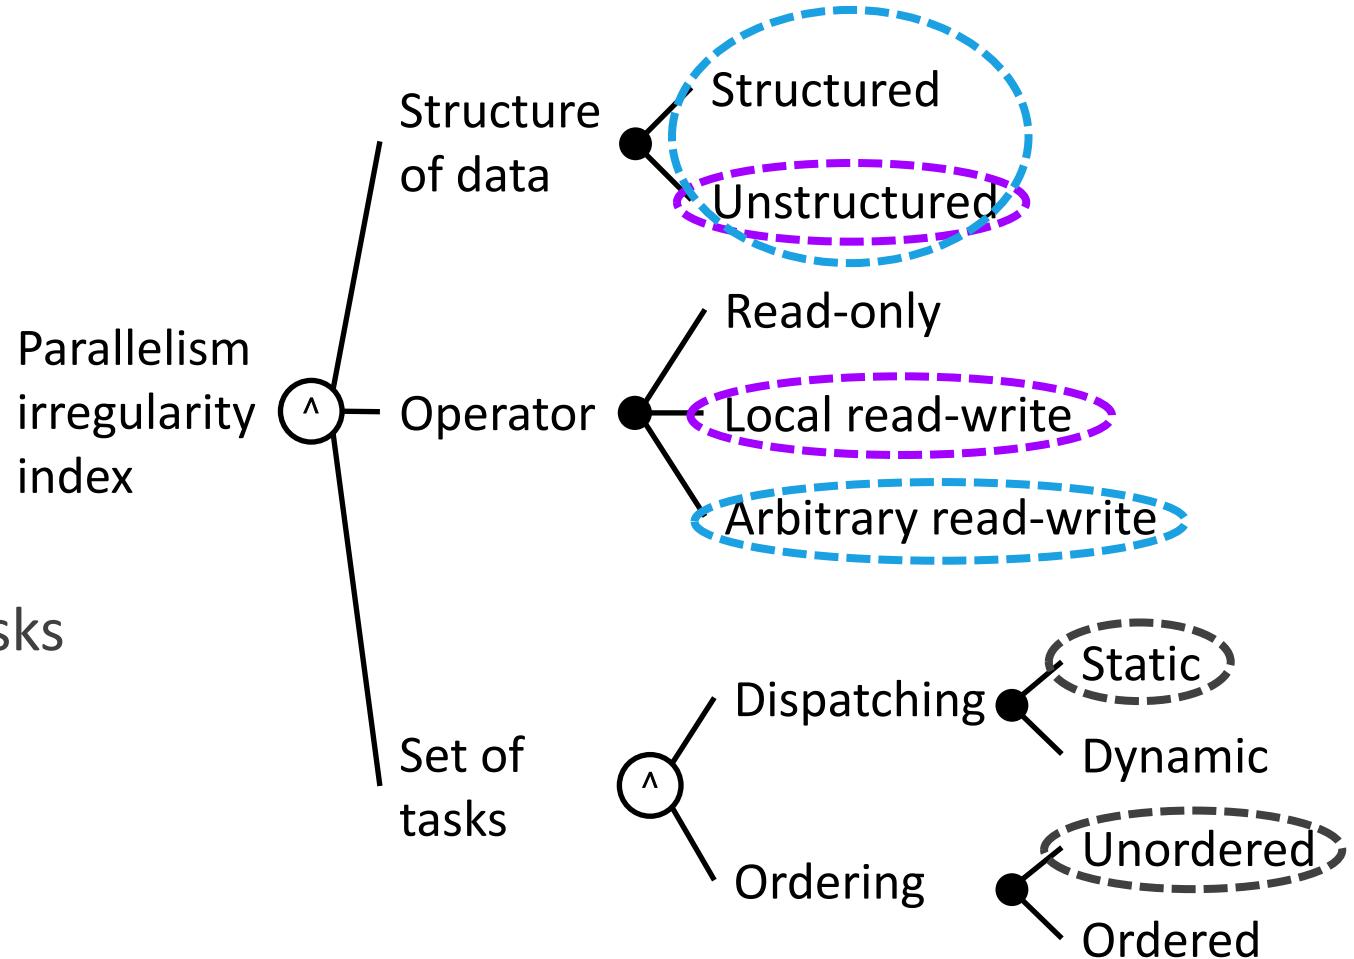


# Rust rejects data-race-prone irregular accesses

Irregular accesses:

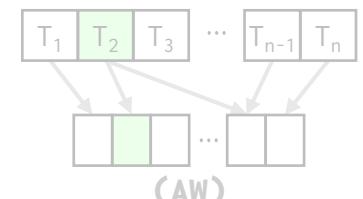
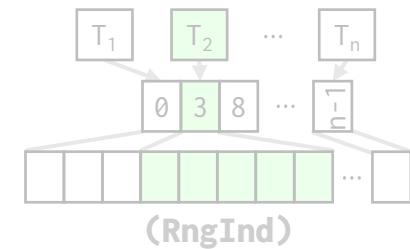
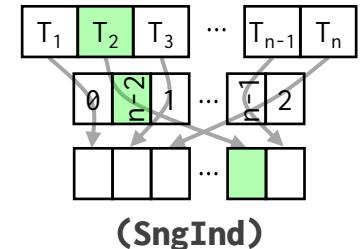
- Arbitrary read-writes
- Local read-writes on unstructured data

Assume a static unordered set of tasks



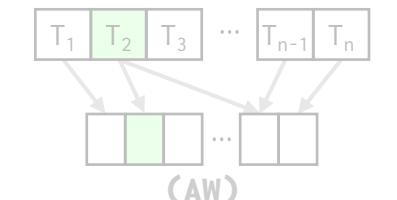
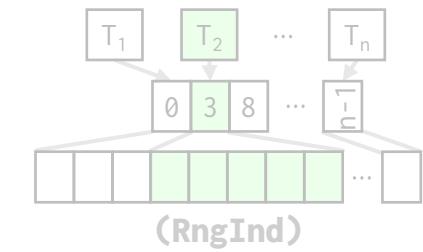
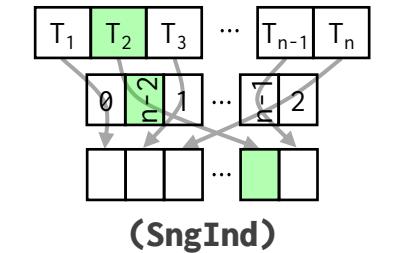
# Rust rejects data-race-prone irregular accesses

```
// serial indirect increment  
(0..n)  
    .into_iter()  
    .for_each(|i| {  
        v[offsets[i]] += 1;  
    })
```



# Rust rejects data-race-prone irregular accesses

```
// serial indirect increment          // parallel  
(0..n)                          (0..n)  
    .into_iter()  
    .for_each(|i| {  
        v[offsets[i]] += 1;  
    })  
    .into_par_iter()  
    .for_each(|i| {  
        v[offsets[i]] += 1;  
    })
```



# Rust rejects data-race-prone irregular accesses

```
// serial indirect increment  
(0..n)  
    .into_iter()  
    .for_each(|i| {  
        v[offsets[i]] += 1;  
    })
```

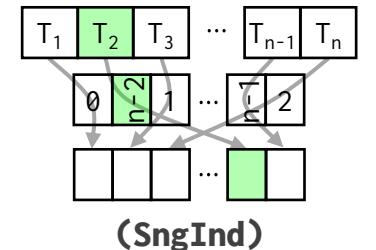
Risky

```
// parallel  
(0..n)
```

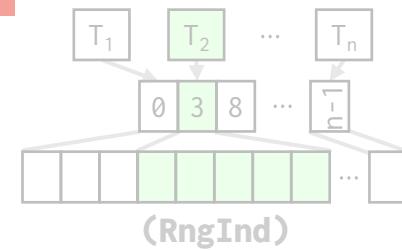
```
.into_par_iter()  
.for_each(|i| {
```

```
v[offsets[i]] += 1;
```

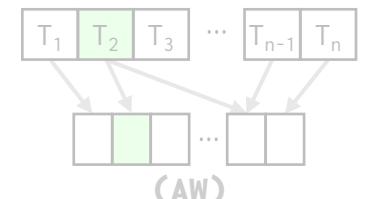
)



(SngInd)



(RngInd)



(AW)

# Rust rejects data-race-prone irregular accesses

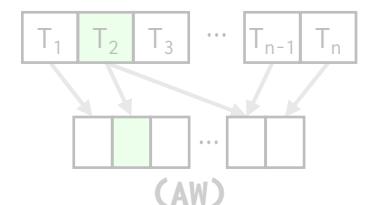
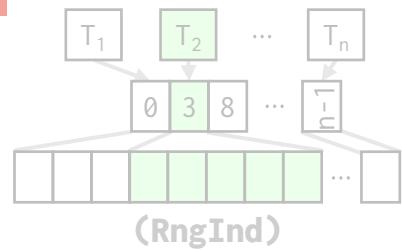
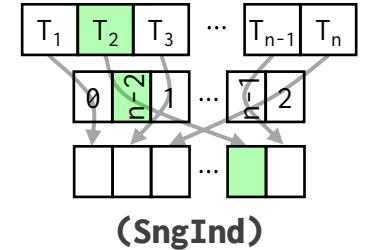
```
// serial indirect increment  
(0..n)  
.into_iter()  
.for_each(|i| {  
    v[offsets[i]] += 1;  
})
```

Risky

```
// parallel  
(0..n)  
.into_par_iter()  
.for_each(|i| {  
    v[offsets[i]] += 1;  
})
```

Compile error

**error[E0596]: cannot borrow `v`  
as mutable, as it is a captured  
variable in a `Fn` closure**



# Rust rejects data-race-prone irregular accesses

```
// serial indirect increment  
(0..n)  
.into_iter()  
.for_each(|i| {  
    v[offsets[i]] += 1;  
})
```

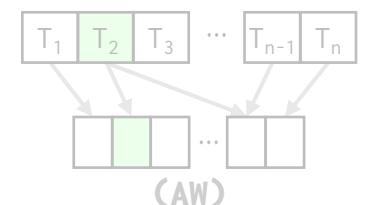
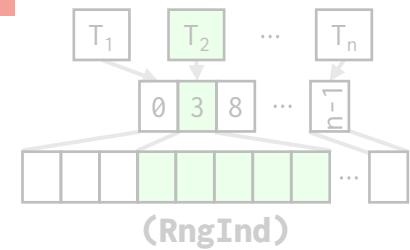
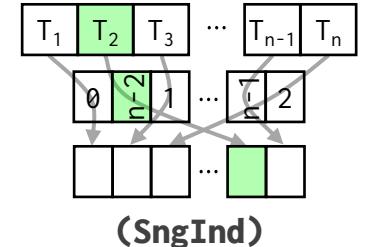
Offsets {  
 Unique  
 Duplicates

```
// parallel  
(0..n)  
.into_par_iter()  
.for_each(|i| {  
    v[offsets[i]] += 1;  
})
```

Risky

Compile error

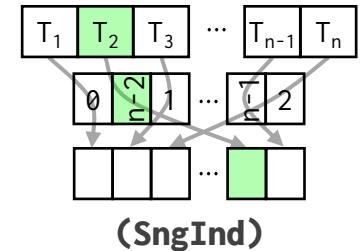
**error[E0596]: cannot borrow `v`  
as mutable, as it is a captured  
variable in a `Fn` closure**



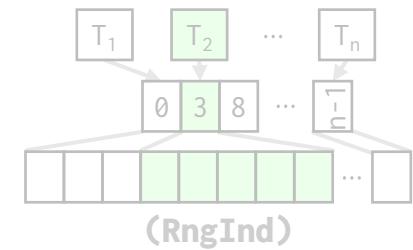
# Programmers face a conundrum expressing irregular accesses

## Synchronization

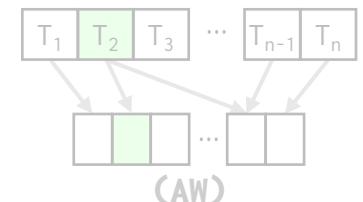
```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        })
    .into_iter()
    .map(|i| {
```



(SngInd)



(RngInd)



(AW)

# Programmers face a conundrum expressing irregular accesses

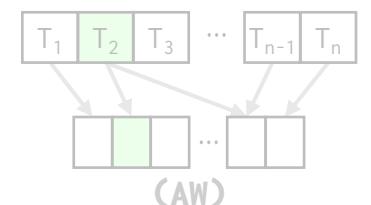
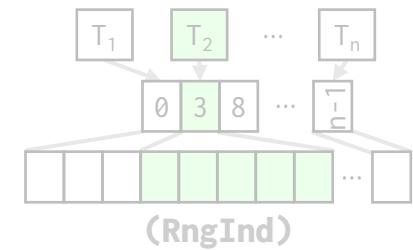
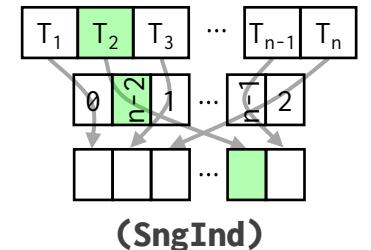
## Synchronization

### Locks

#### Coarse

500x

```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        v.lock().unwrap()[offsets[i]]+=1;
    })
}
```



# Programmers face a conundrum expressing irregular accesses

## Synchronization

### Locks

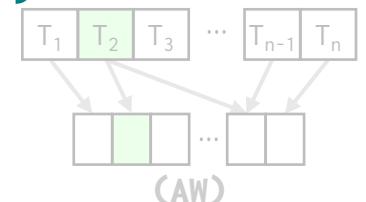
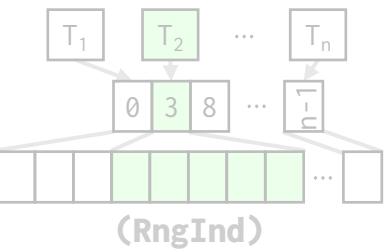
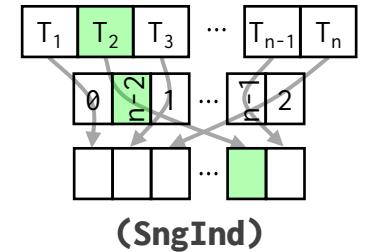
#### Coarse

500x

#### Fine

5x

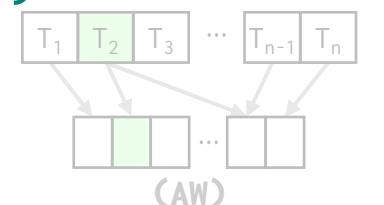
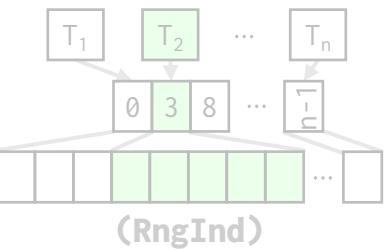
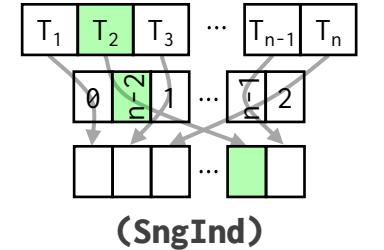
```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        *v[offsets[i]].lock().unwrap()+=1;
    })
}
```



# Programmers face a conundrum expressing irregular accesses

Synchronization  
└ Locks  
  └ Coarse      500x  
  └ Fine        5x  
Atoms

```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        v[offsets[i]].fetch_add(1, Order);
    })
}
```



# Programmers face a conundrum expressing irregular accesses

## Synchronization

### Locks

#### Coarse

500x

#### Fine

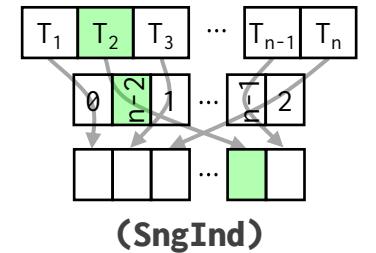
5x

### Atomics

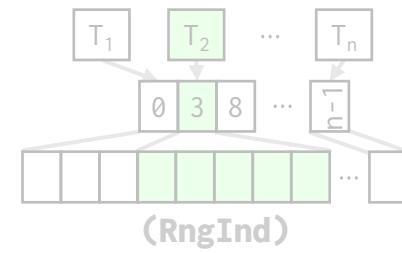
#### Relaxed

1x

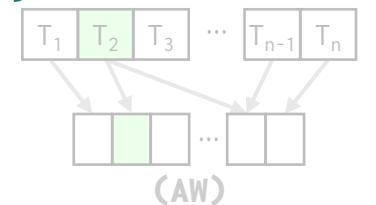
```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        v[offsets[i]].fetch_add(1, Order);
    })
}
```



(SngInd)



(RngInd)



(AW)

# Programmers face a conundrum expressing irregular accesses

## Synchronization

### Locks

#### Coarse

500x

#### Fine

5x

### Atomics

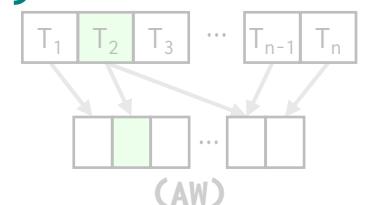
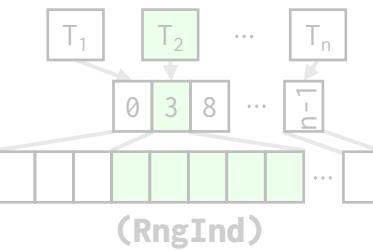
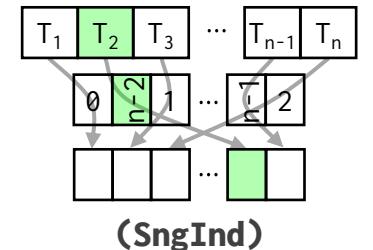
#### Relaxed

1x

#### SeqCst

2x

```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        v[offsets[i]].fetch_add(1, Order);
    })
}
```



# Programmers face a conundrum expressing irregular accesses

## Synchronization **S**

Locks

Coarse      500x

Fine      5x

Atomics

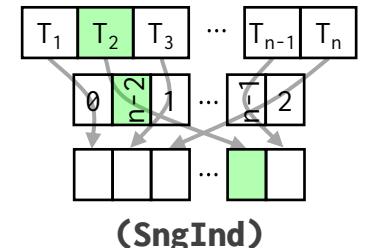
Relaxed      1x

SeqCst      2x

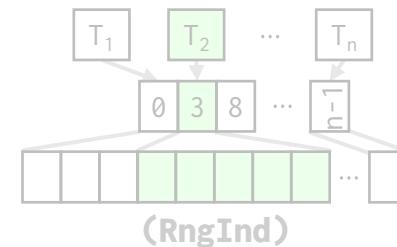
```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        Old = v[offsets[i]].load(Order);
        v[offsets[i]].store(old+1, Order);
    })

```

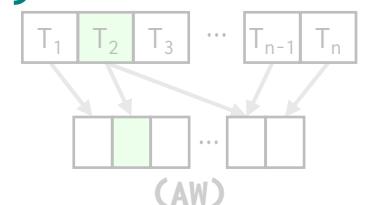
Atomicity violation -> compiles!!



(SngInd)



(RngInd)



(AW)

# Programmers face a conundrum expressing irregular accesses

## Synchronization **S**

Locks

Coarse      500x

Fine      5x

Atomics

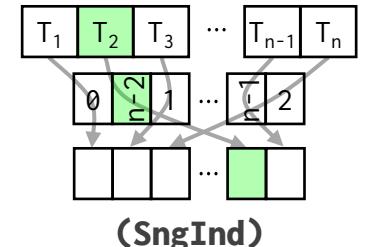
Relaxed      1x

SeqCst      2x

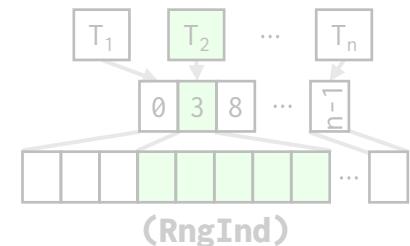
## Unsafe

```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        Old = v[offsets[i]].load(Order);
        v[offsets[i]].store(old+1, Order);
    })
}
```

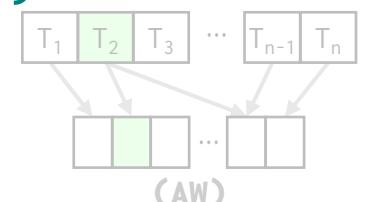
Atomicity violation -> compiles!!



(SngInd)



(RngInd)

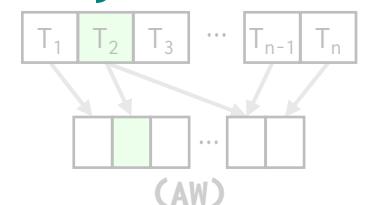
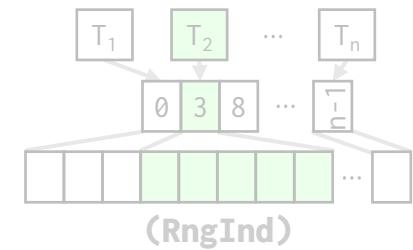
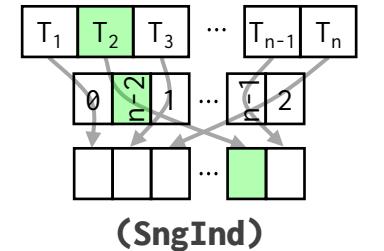


(AW)

# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x

```
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T) += 1;
        }
    })
})
```



# Programmers face a conundrum expressing irregular accesses

## Synchronization **S**

Locks

Coarse      500x

Fine      5x

Atomics

Relaxed      1x

SeqCst      2x

## Unsafe

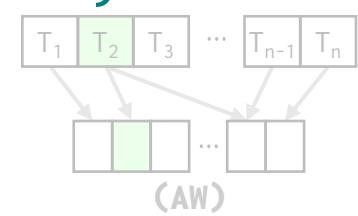
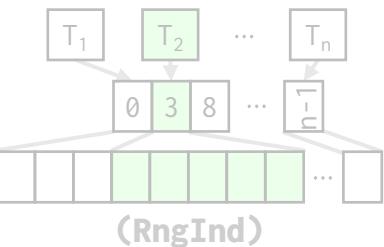
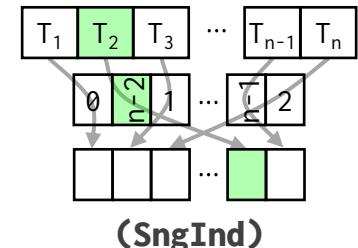
unchecked 1x

with check 3x

**S**

**C**

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T) += 1;
        }
    })
}
```

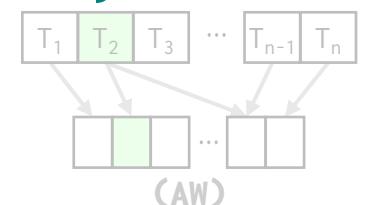
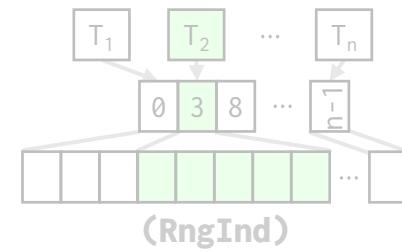
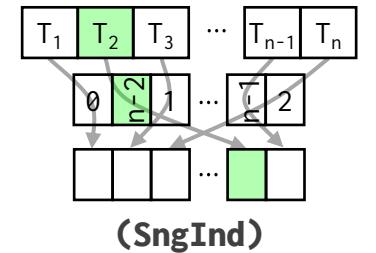


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T)+=1;
        }
    })
}
```

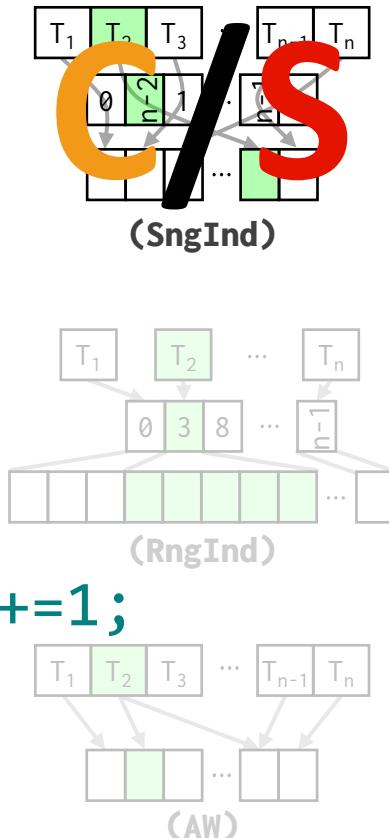


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S }

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T)+=1;
        }
    })
}
```

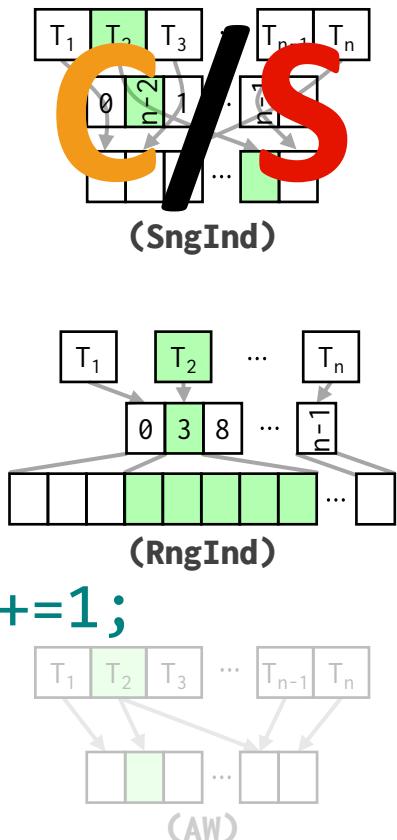


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S }

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T)+=1;
        }
    })
}
```

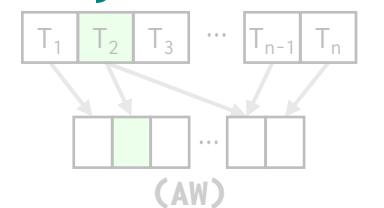
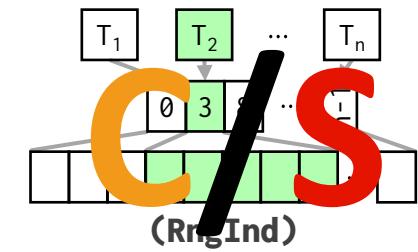
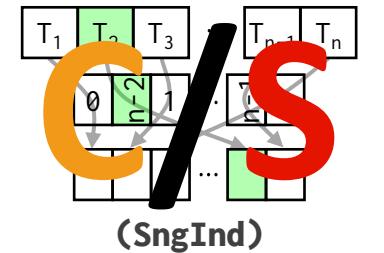


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S }

```
check_duplicates(offsets);
(0..n)
  .into_par_iter()
  .for_each(|i| {
    unsafe {
      *(v.as_ptr().add(i) as *mut T) += 1;
    }
  })
}
```

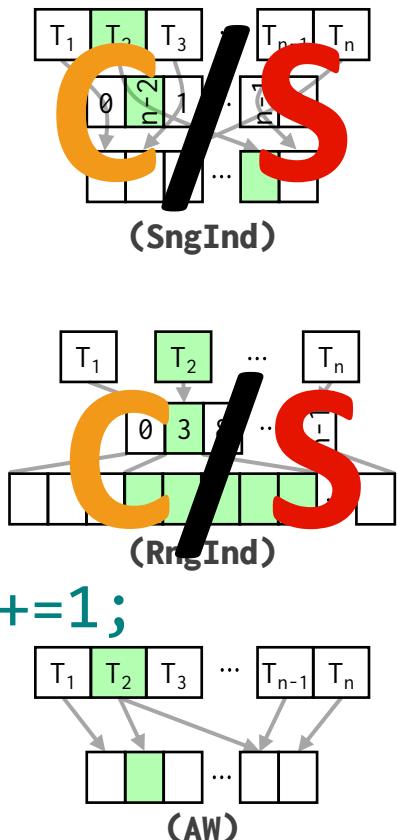


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S }

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T)+=1;
        }
    })
}
```

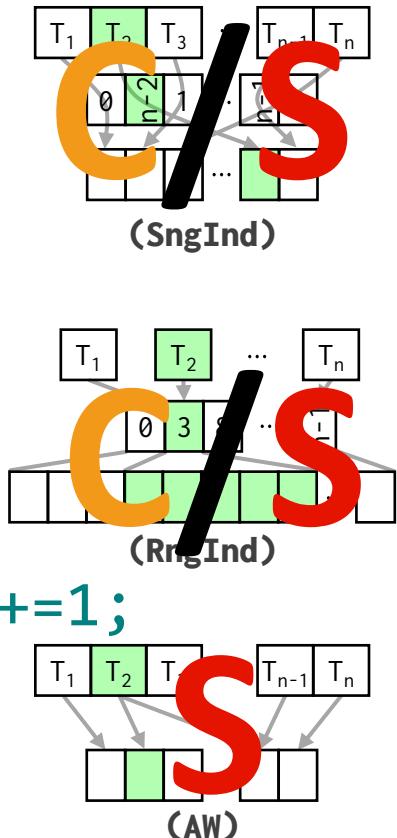


# Programmers face a conundrum expressing irregular accesses

Synchronization	S
Locks	
Coarse	500x
Fine	5x
Atomics	
Relaxed	1x
SeqCst	2x
Unsafe	
unchecked	1x
with check	3x

Offsets { Unique      C/S  
              Duplicates S }

```
check_duplicates(offsets);
(0..n)
    .into_par_iter()
    .for_each(|i| {
        unsafe {
            *(v.as_ptr().add(i) as *mut T)+=1;
        }
    })
}
```



# Programmers face a conundrum expressing irregular accesses

Synchronization **S**

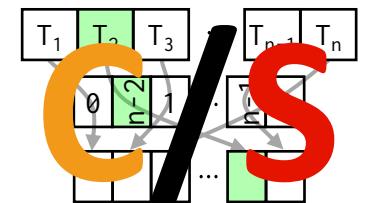
Locks

Coarse

500x

Offsets { Unique **C/S**  
Duplicates **S**

`check_duplicates(offsets);`



**Rust solutions for irregular accesses are not fearless**

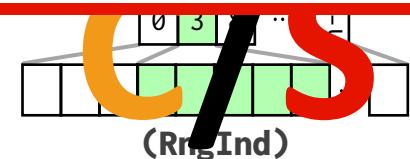
Relaxed

1x

SeqCst

2x

```
.for_each(|i| {
    unsafe {
        *(v.as_ptr().add(i)) as *mut T) += 1;
    }
})
```

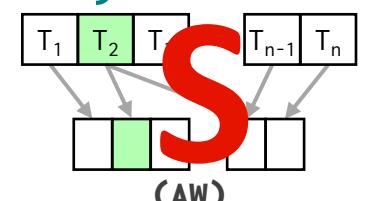


Unsafe

unchecked 1x

with check 3x

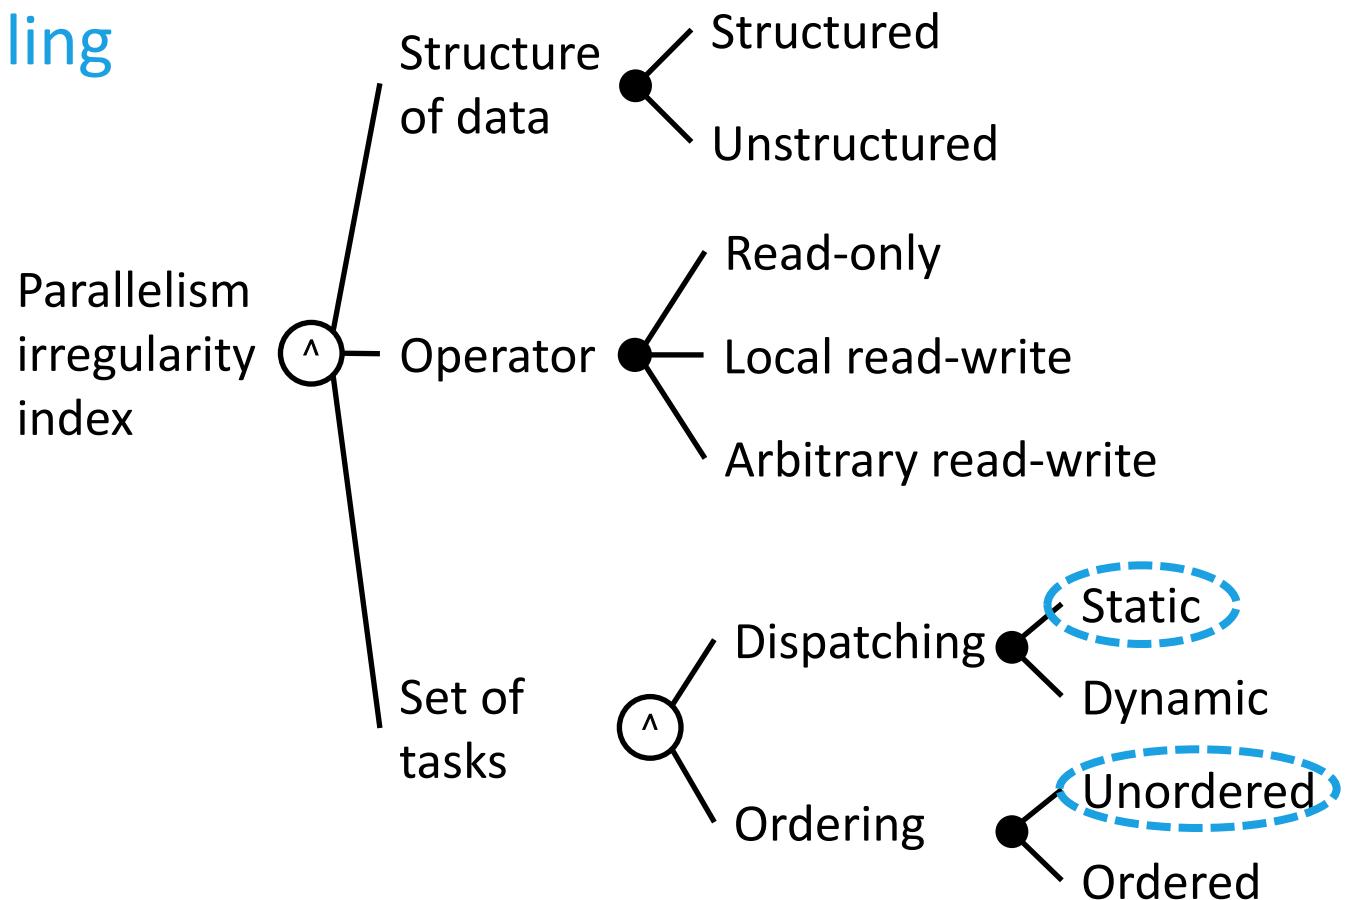
**S**  
**C**



# Task scheduling

So far, static unordered scheduling

Let's consider other cases



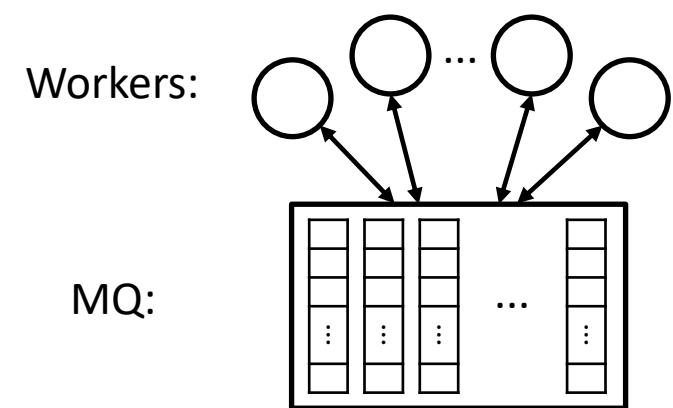
# Task scheduling

---

**Static:** E.g., bulk-synchronous parallelism (split-merge)

**Dynamic:**

- Tasks spawn new tasks
- Workers pop a task, execute, push new tasks (if exist)



E.g., MultiQueue

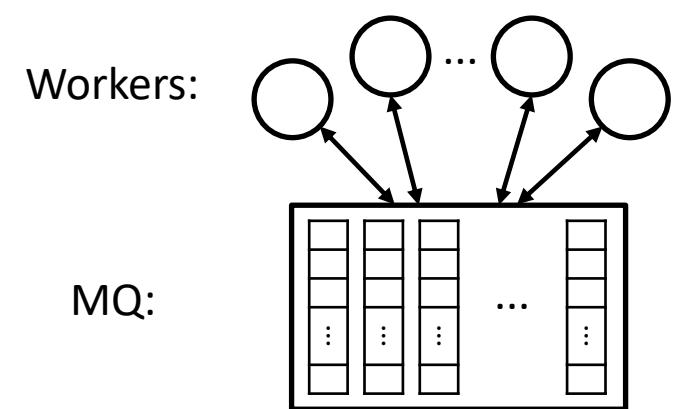
# Task scheduling

---

**Static:** E.g., bulk-synchronous parallelism (split-merge)

**Dynamic:**

- Tasks spawn new tasks
- Workers pop a task, execute, push new tasks (if exist)



E.g., MultiQueue

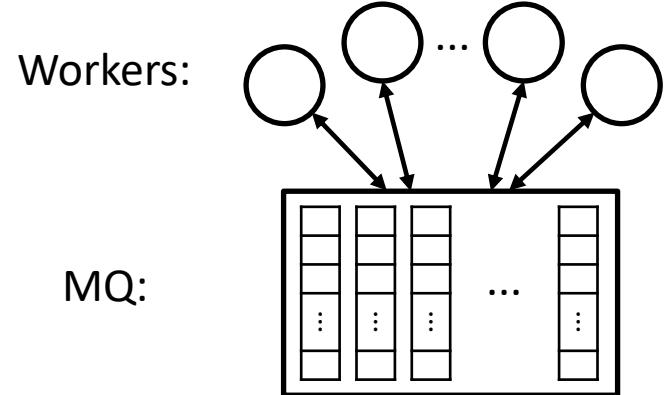
User: depends on access types **F/C/S**

# Task scheduling

**Static:** E.g., bulk-synchronous parallelism (split-merge)

**Dynamic:**

- Tasks spawn new tasks
- Workers pop a task, execute, push new tasks (if exist)



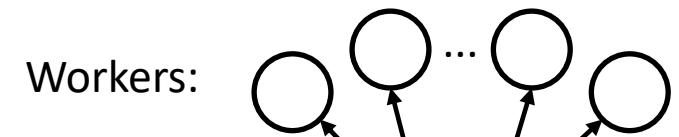
E.g., MultiQueue

- User: depends on access types **F/C/S**
- Implementer: Workers and Queue **S**

# Task scheduling

**Static:** E.g., bulk-synchronous parallelism (split-merge)

**Dynamic:**



**The dynamism of task scheduling does not affect fearlessness**



E.g., MultiQueue

- └ User: depends on access types **F/C/S**
- └ Implementer: Workers and Queue **S**

# Evaluation of RPB

---

# Irregular accesses are common in RPB

Bench-mark	Tasks' Accesses				Task dispatch	
	Regular		Irregular		static	dynamic
	RO	Stride	Block	D&C		
bw	✓	✓	✓		✓	✓
dedup	✓	✓	✓		✓	✓
dr	✓	✓	✓		✓	✓
hist	✓	✓	✓		✓	✓
isort	✓	✓	✓		✓	✓
lrs		✓	✓	✓	✓	✓
mis		✓	✓	✓	✓	✓
mm		✓	✓	✓	✓	✓
msf		✓	✓	✓	✓	✓
sa	✓	✓	✓	✓	✓	✓
sf		✓	✓	✓	✓	✓
sort	✓		✓		✓	✓
bfs					✓	✓
sssp					✓	✓

# Irregular accesses are common in RPB

Bench-mark	Tasks' Accesses				Task dispatch		
	Regular	Irregular					
	RO	Stride	Block	D&C	Sng Ind	Rng Ind	AW
bw	✓	✓			✓	✓	✓
dedup	✓	✓	✓		✓	✓	✓
dr	✓	✓	✓		✓	✓	✓
hist	✓	✓	✓		✓	✓	✓
isort	✓	✓	✓		✓	✓	✓
lrs	✓	✓	✓	✓	✓	✓	✓
mis	✓	✓	✓		✓	✓	✓
mm	✓	✓	✓		✓	✓	✓
msf	✓	✓	✓		✓	✓	✓
sa	✓	✓			✓	✓	✓
sf		✓			✓	✓	✓
sort	✓		✓	✓	✓		
bfs					✓	✓	
sssp					✓		

Regular accesses      F  
Irregular accesses    C/S

# Irregular accesses are common in RPB

Bench-mark	Tasks' Accesses				Task dispatch		
	Regular	Irregular					
	RO	Stride	Block	D&C	Sng Ind	Rng Ind	AW
bw	✓	✓			✓	✓	✓
dedup	✓	✓			✓	✓	✓
dr	✓	✓			✓	✓	✓
hist	✓	✓			✓	✓	✓
isort	✓	✓			✓	✓	✓
..	.	.			✓	✓	✓

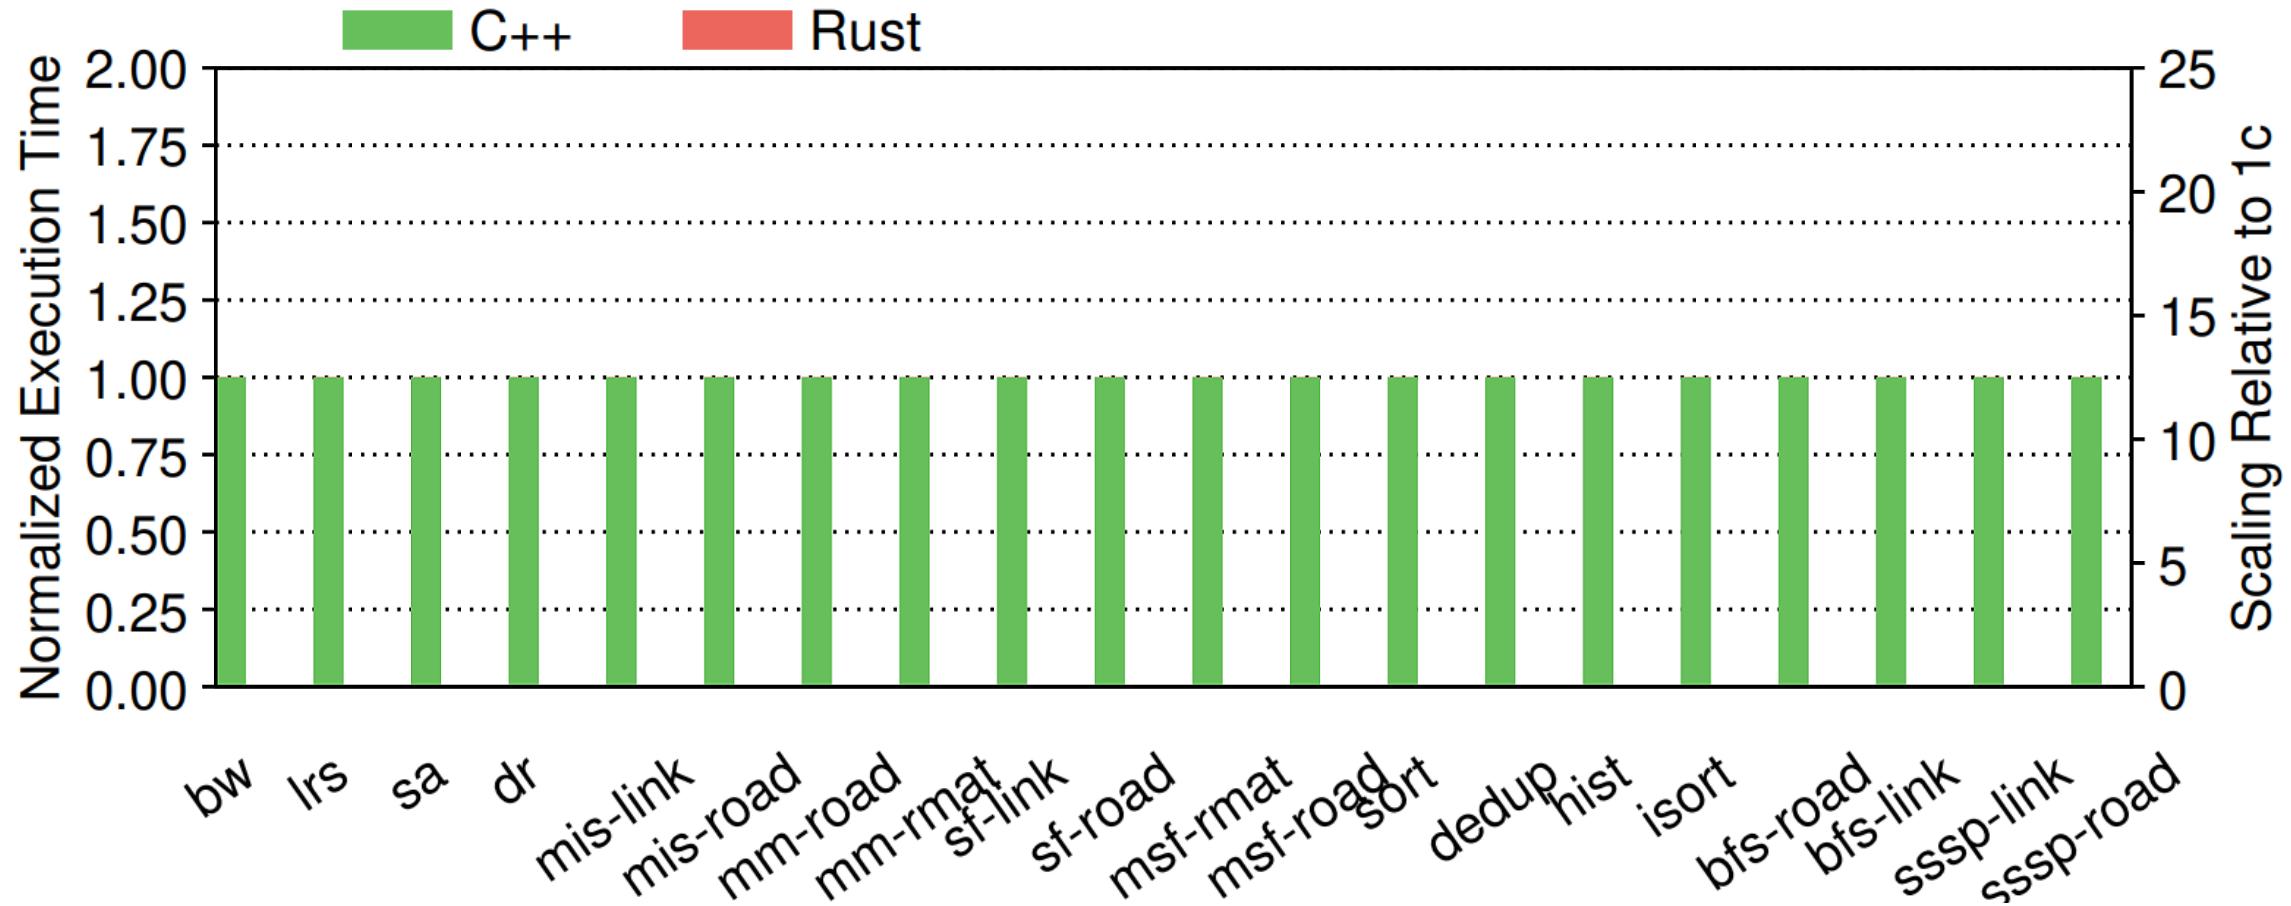
Regular accesses F  
Irregular accesses C/S

Expressing RPB in Rust is not fearless

	Regular	Irregular	Task dispatch
sa	✓	✓	✓
sf	✓	✓	✓
sort	✓	✓	✓
bfs			✓
sssp	✓	✓	✓

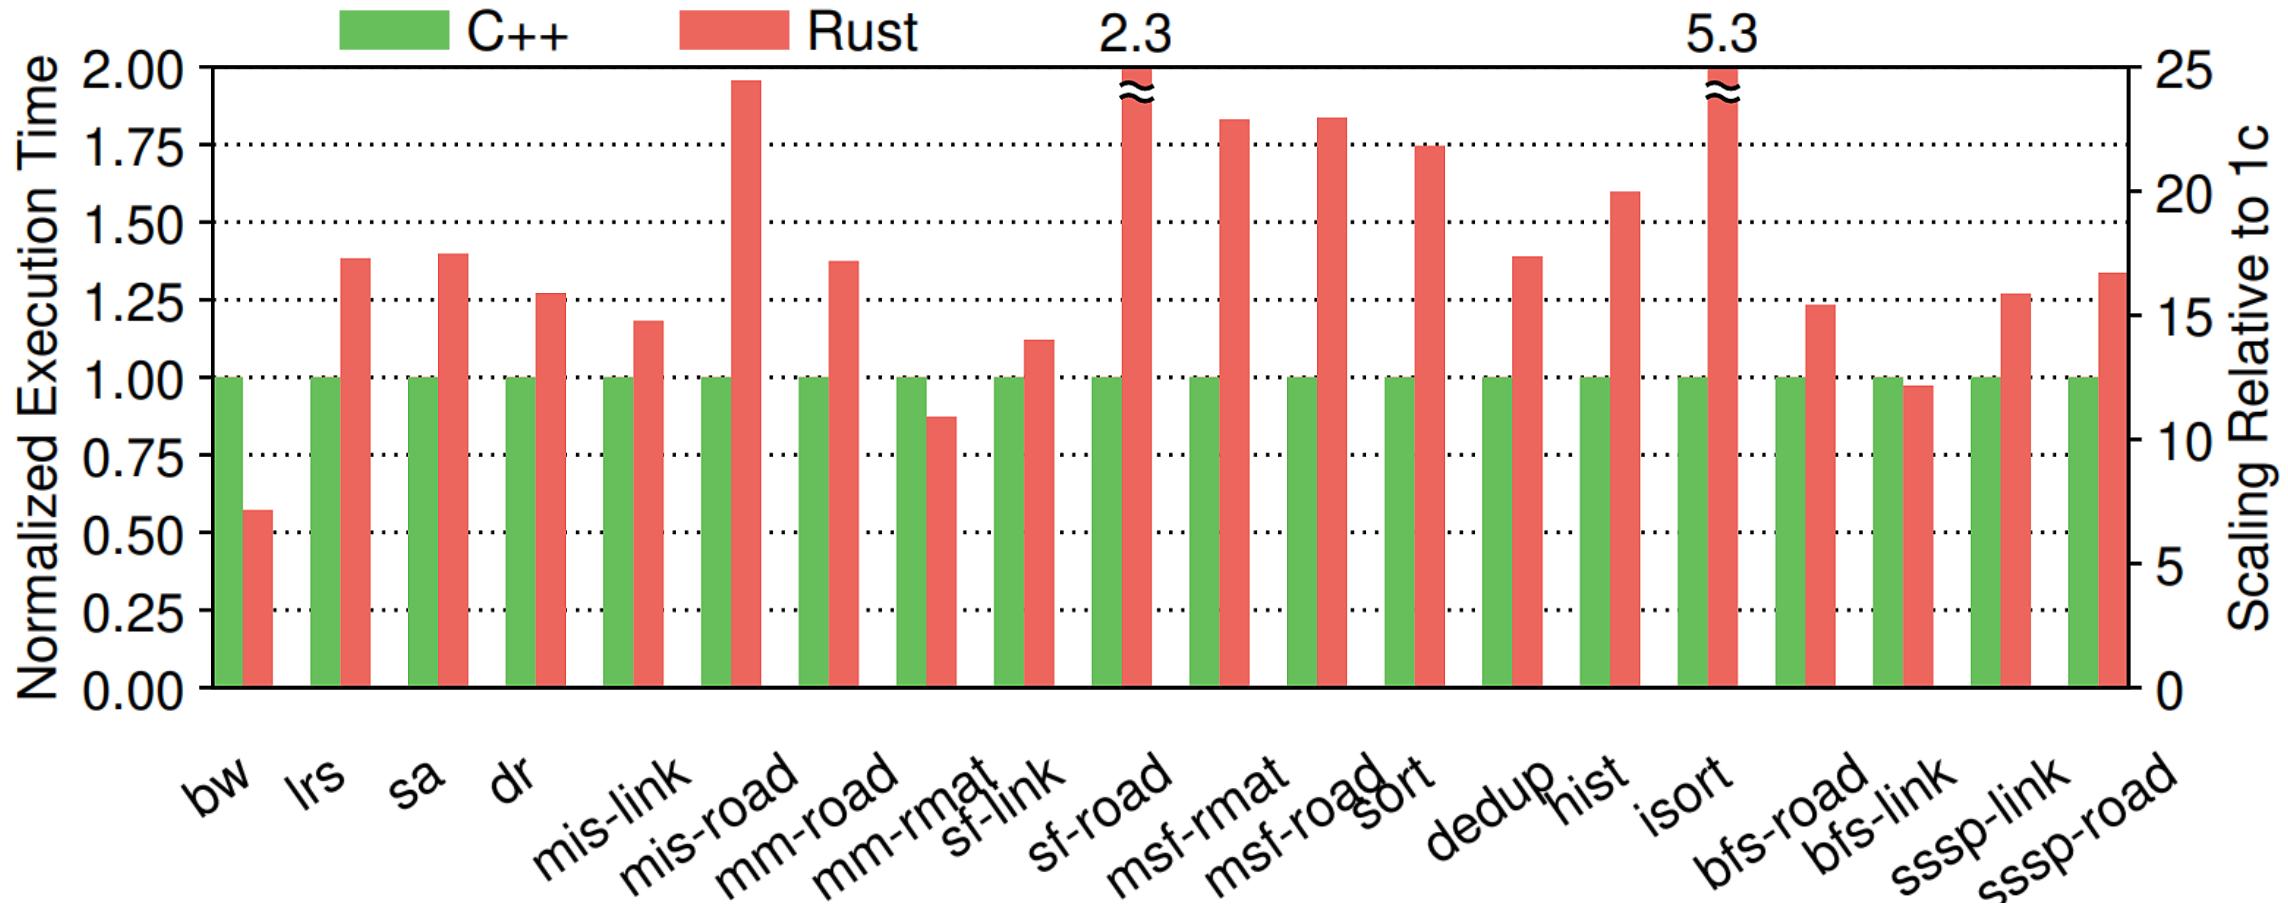
# Performance of RPB vs C++ across benchmarks

\* using unsafe for irregular accesses (24 cores)



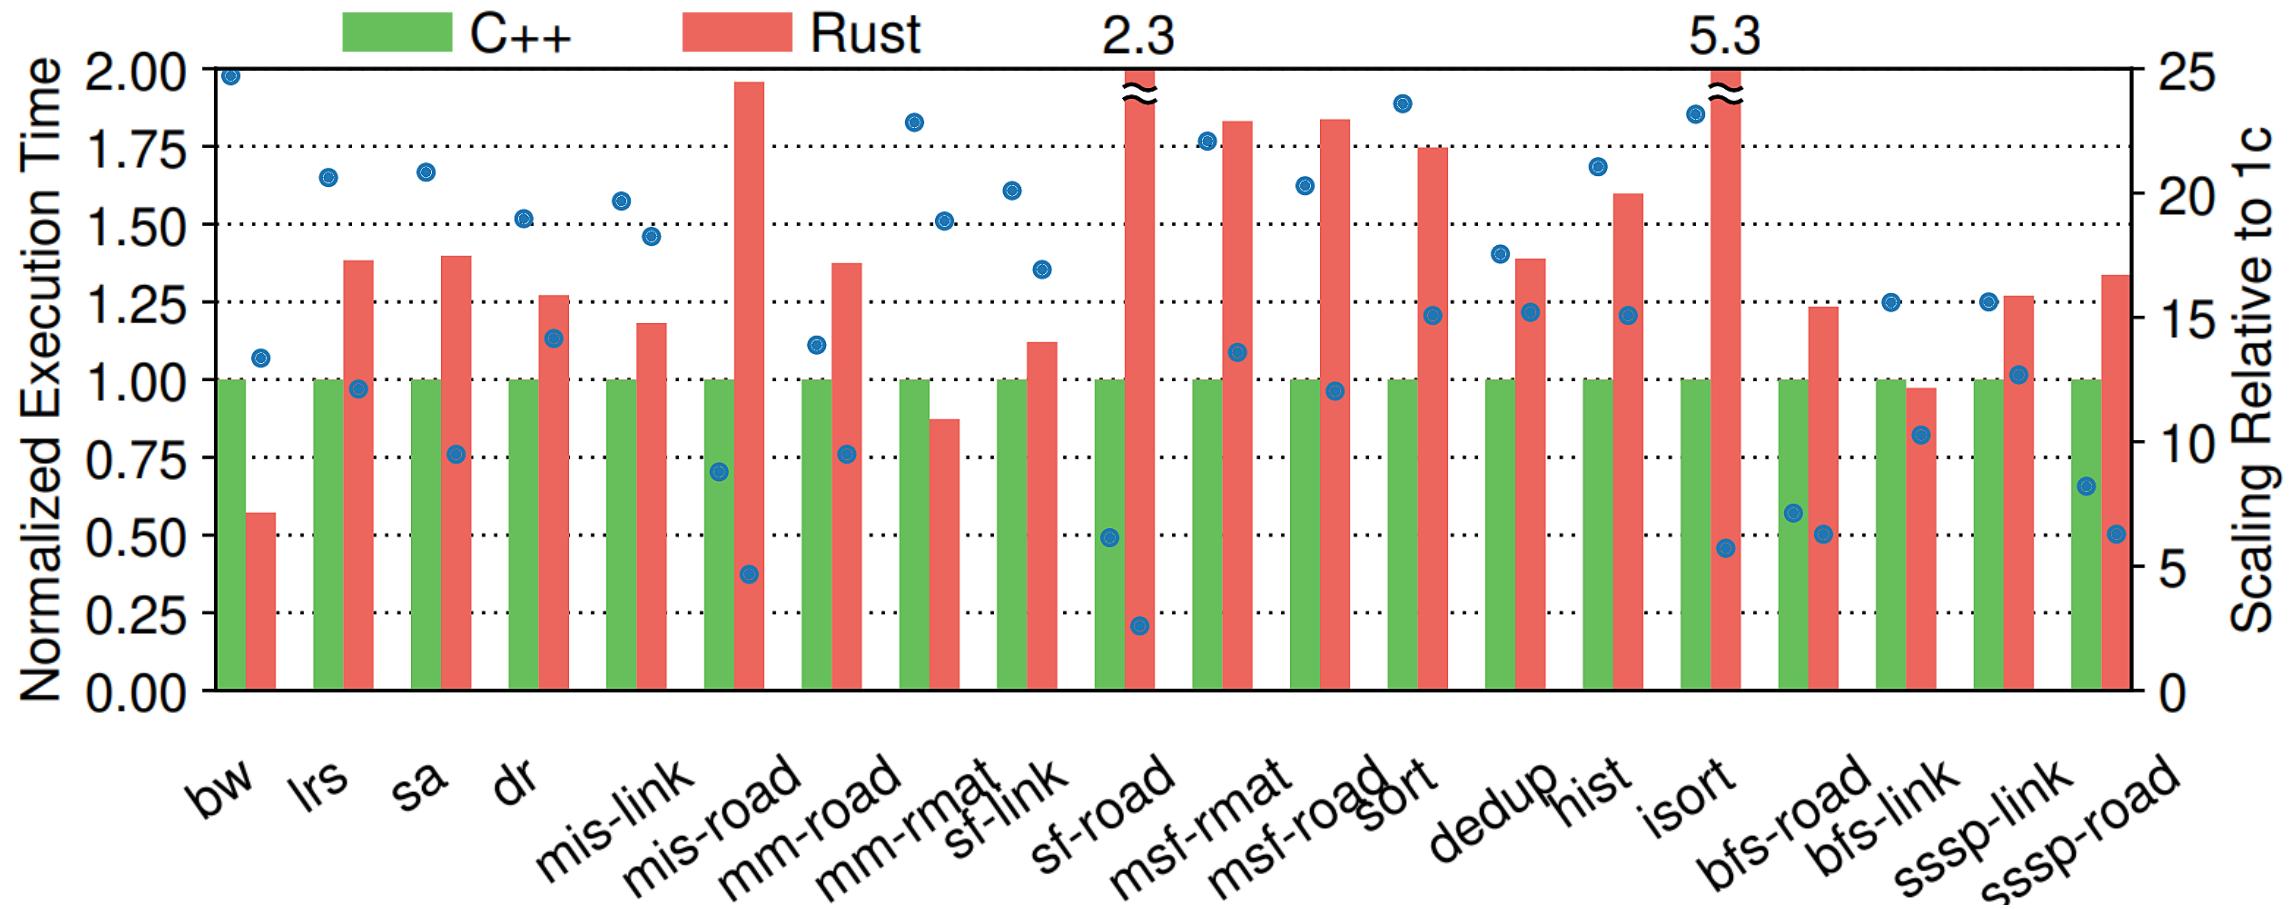
# Performance of RPB vs C++ across benchmarks

\* using unsafe for irregular accesses (24 cores)



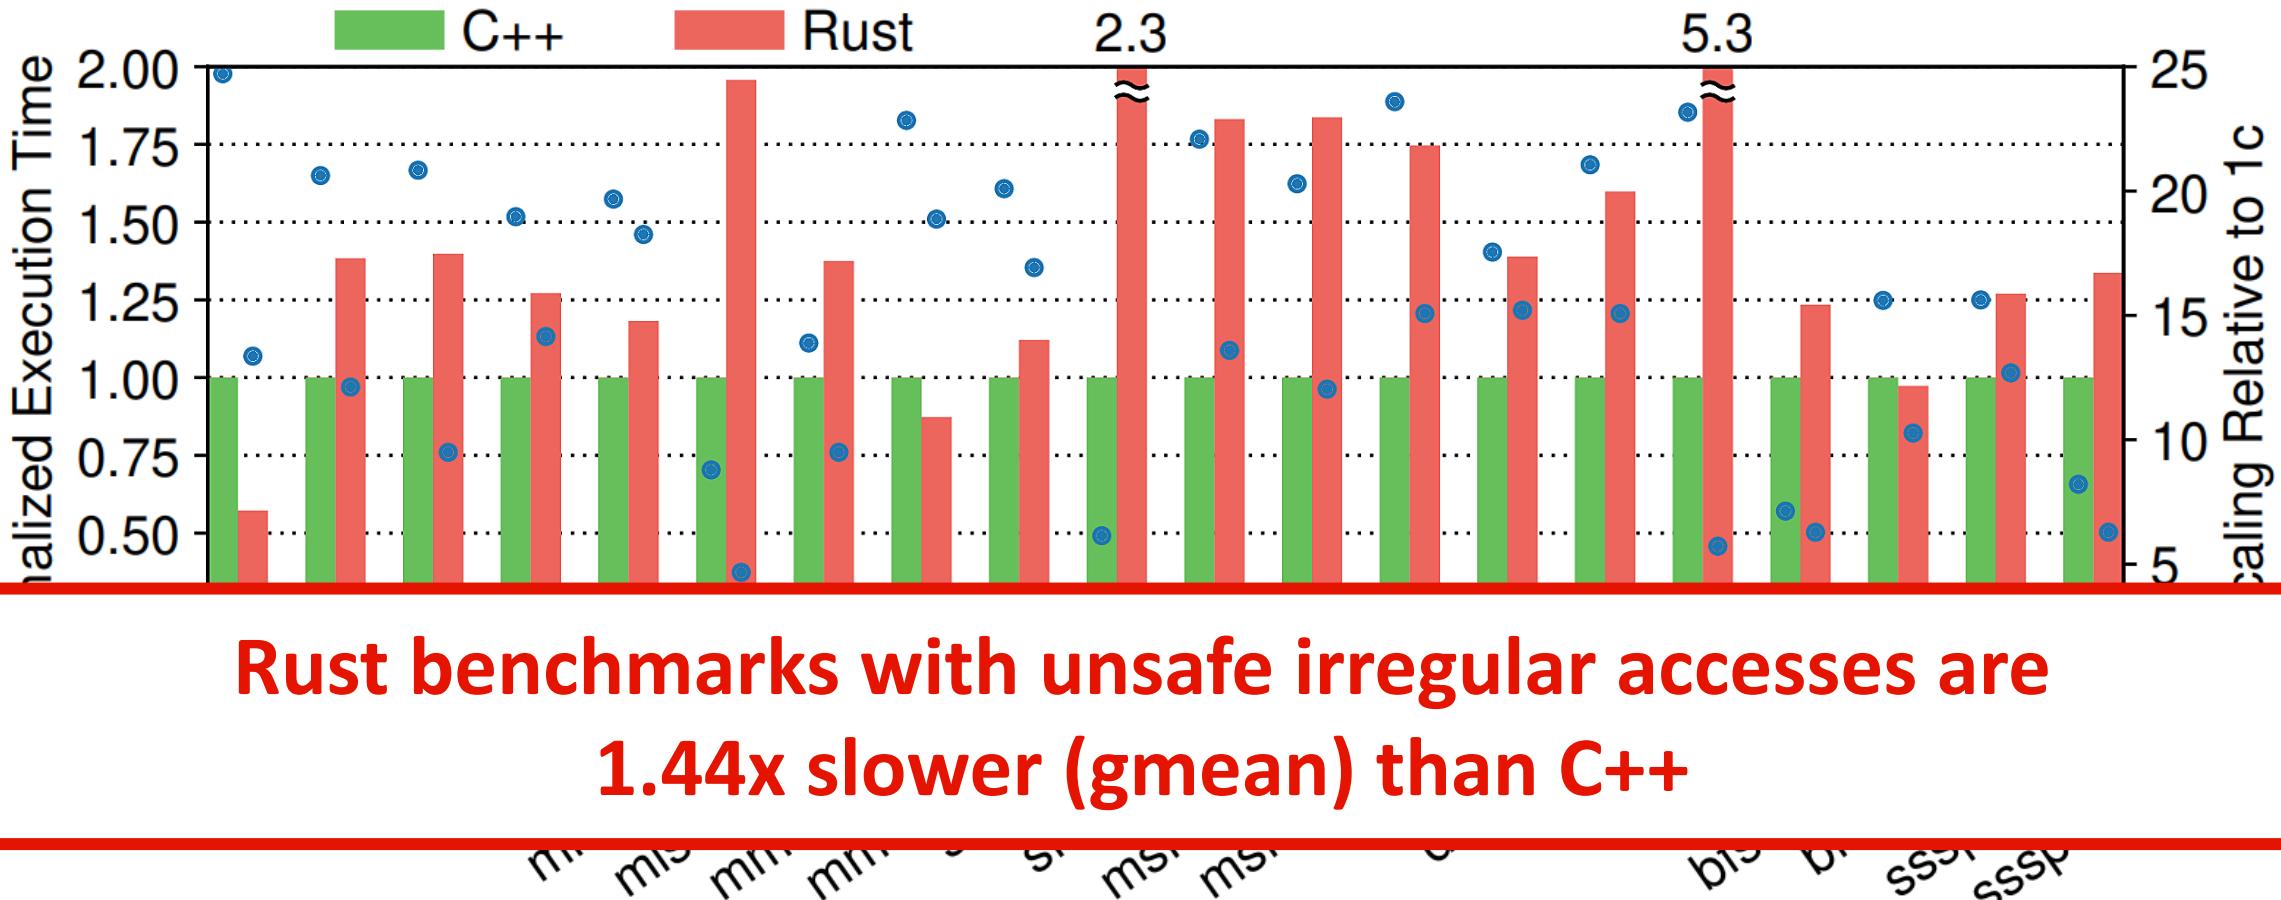
# Performance of RPB vs C++ across benchmarks

\* using unsafe for irregular accesses (24 cores)



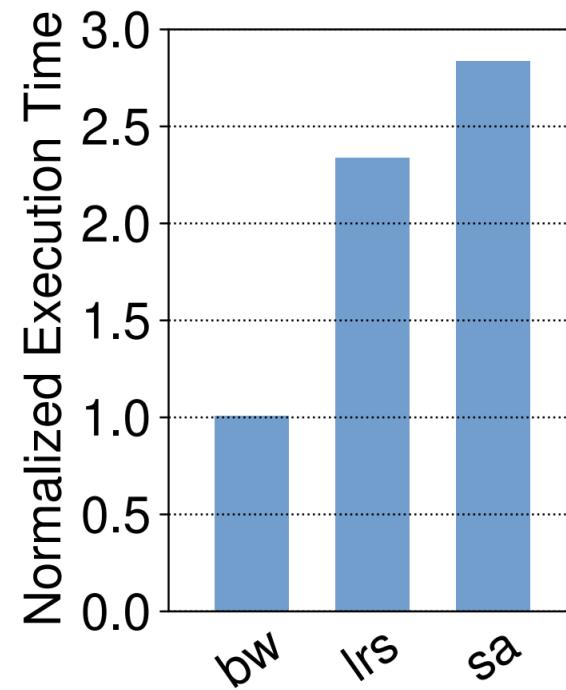
# Performance of RPB vs C++ across benchmarks

\* using unsafe for irregular accesses (24 cores)

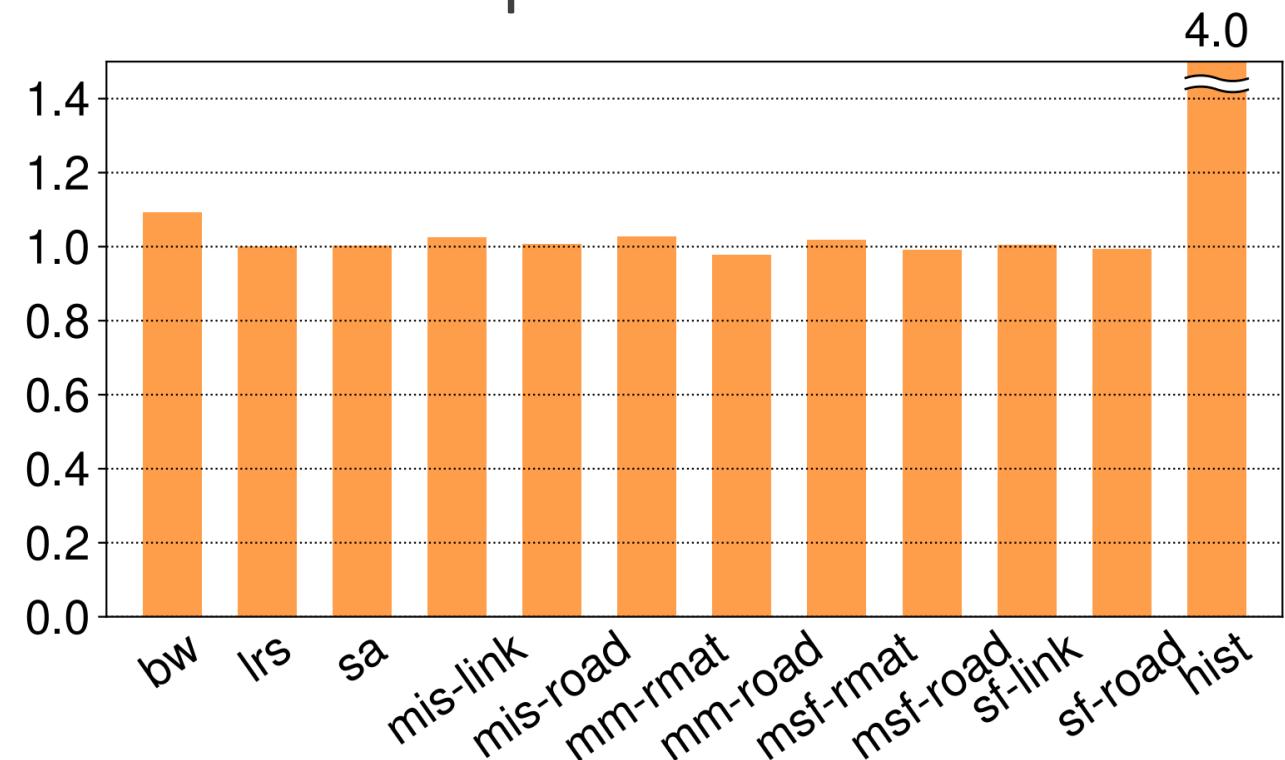


See paper for more evaluations on ...

Runtime checks with unsafe blocks



Unnecessary synchronization  
to placate rustc



# Conclusions

---

To replace C++, Rust's claims should be verified, including fearless concurrency

Rust grants fearless concurrency for regular accesses but not irregular ones

The task scheduling scheme does not affect fearlessness

Using (current) Rust can have a performance penalty



Find RPB on github!



## Conclusions

---

To replace C++, Rust's claims should be verified, including fearless concurrency

Rust grants fearless concurrency for regular accesses but not irregular ones

The task scheduling scheme does not affect fearlessness

Using (current) Rust can have a performance penalty