

PARALLEL PROGRAMMING IN HASKELL

David Walker

Thanks to Kathleen Fisher and recursively to
Simon Peyton Jones for much of the content of these slides.

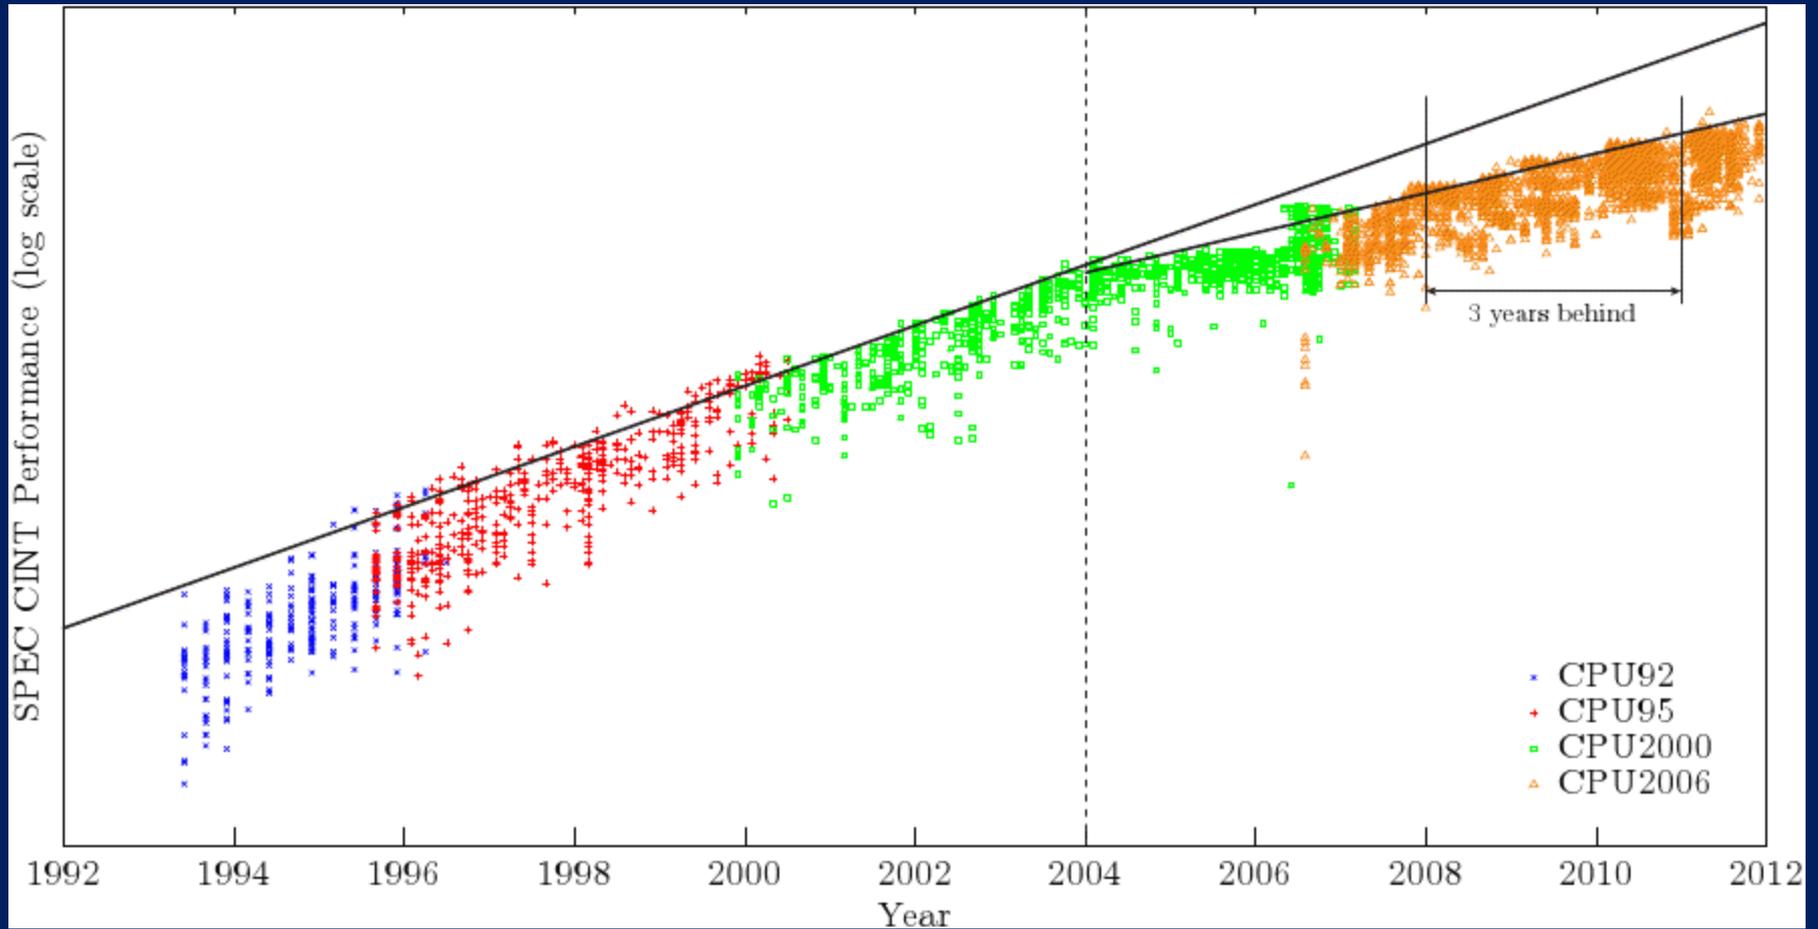
Optional Reading:

“Beautiful Concurrency”,

“The Transactional Memory / Garbage Collection Analogy”

“A Tutorial on Parallel and Concurrent Programming in Haskell”

Intel's Nightmare



Produced by Arun Raman, Ph.D., Princeton 2011.

The Multi-Cores are Coming

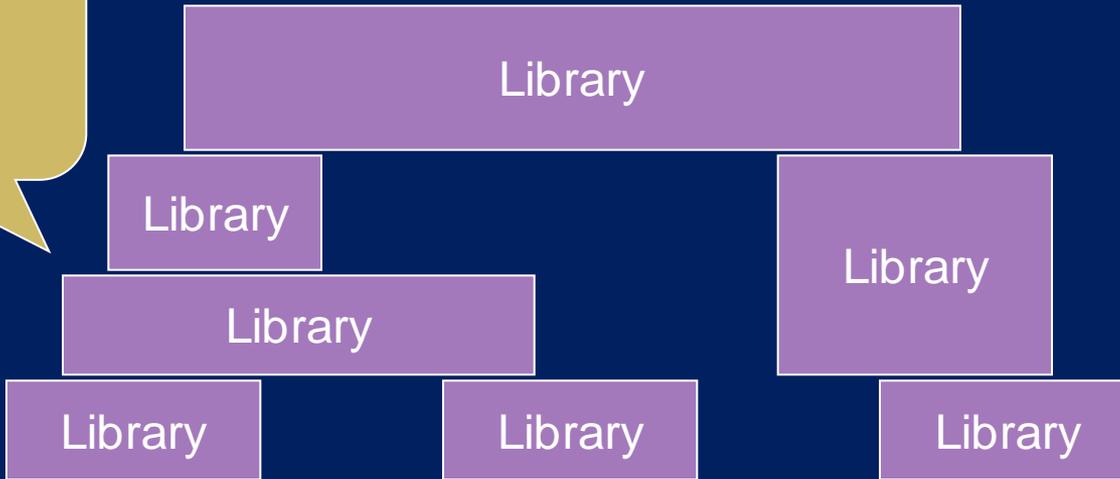
- Multi-cores are coming!
 - For 50 years, hardware designers delivered **40-50% increases per year** in sequential program performance.
 - Around 2004, this **pattern failed** because power and cooling issues made it impossible to increase clock frequencies.
 - Now hardware designers are using the extra transistors that Moore's law is still delivering to put more processors on a single chip.
- *If we want to improve performance, parallelism is no longer optional.*

Parallelism

- Parallelism is essential to improve performance on a multi-core machine.
- Unfortunately, parallel programming is seen as immensely more error-prone than traditional sequential programming, and it often is
 - If you have taken COS 318, you'll know already that using locks and condition variables is unbelievably hard relative to using if statements and method calls

What we want

Libraries build layered concurrency abstractions

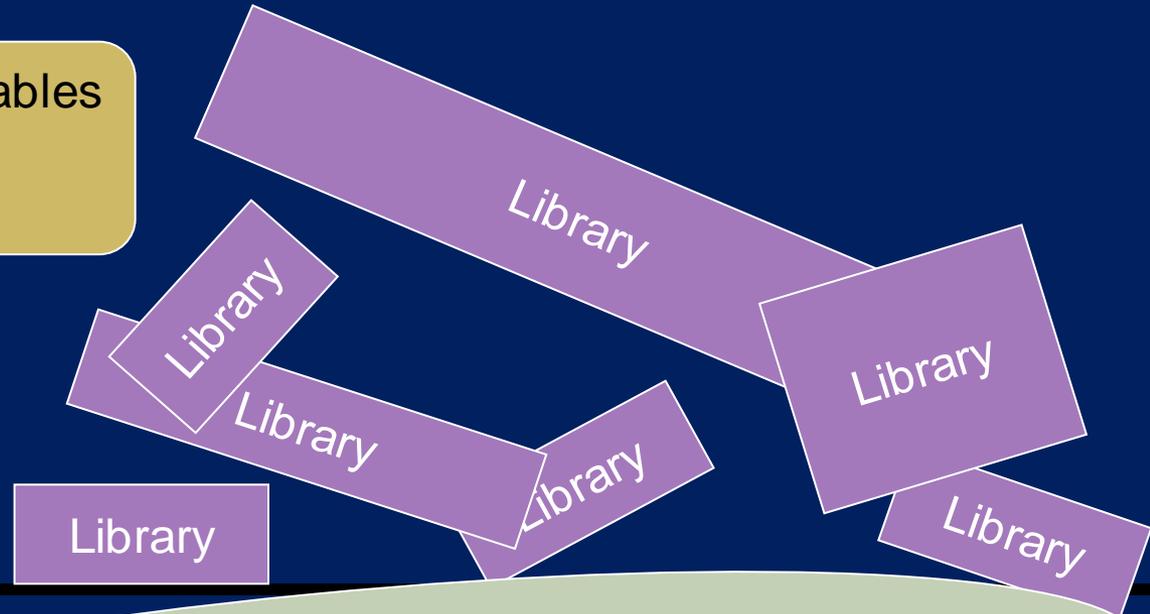


Concurrency primitives

Hardware

What we have using conventional techniques

Locks and condition variables
(a) are hard to use and
(b) do not compose



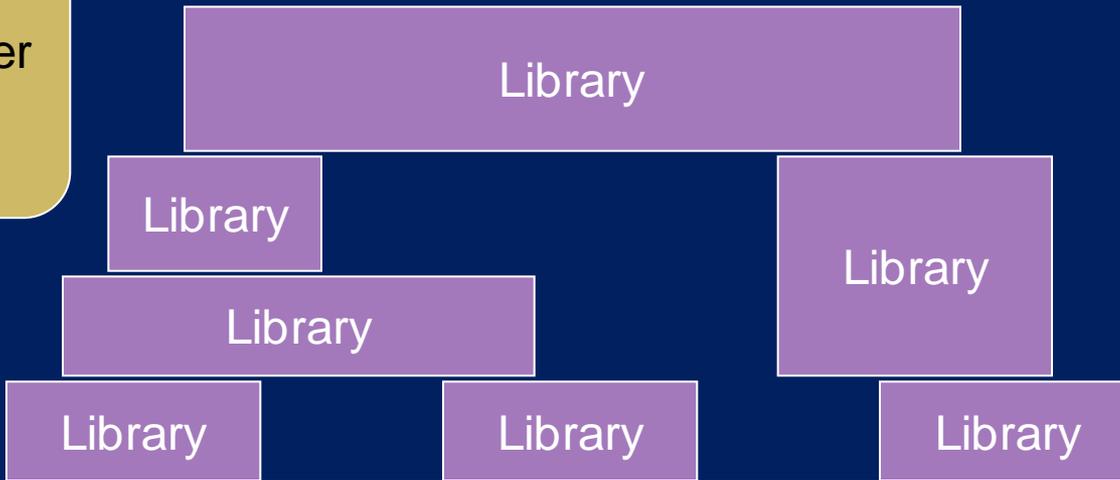
Locks and condition variables

Hardware

“Building complex parallel programs is like building a sky scraper out of bananas.” -- Simon Peyton Jones

Idea: Replace locks with atomic blocks

Atomic blocks are much easier to use, and do compose



Atomic blocks

Hardware

A Quick Primer on Locks and Concurrency

A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**

```
int x = 0;
```

global x
is initially 0



A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**

```
int x = 0;
```

global x
is initially 0

two imperative assignments
execute "at the same time":

```
x = x + 1;
```



```
x = x + 1;
```



A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**

```
int x = 0;
```

global x
is initially 0

two imperative assignments
execute "at the same time":

```
x = x + 1;
```



```
x = x + 1;
```



possible answers: 1 or 2

A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**

```
int x = 0;
```

```
tempA = x;
```

```
x = tempA + 1;
```

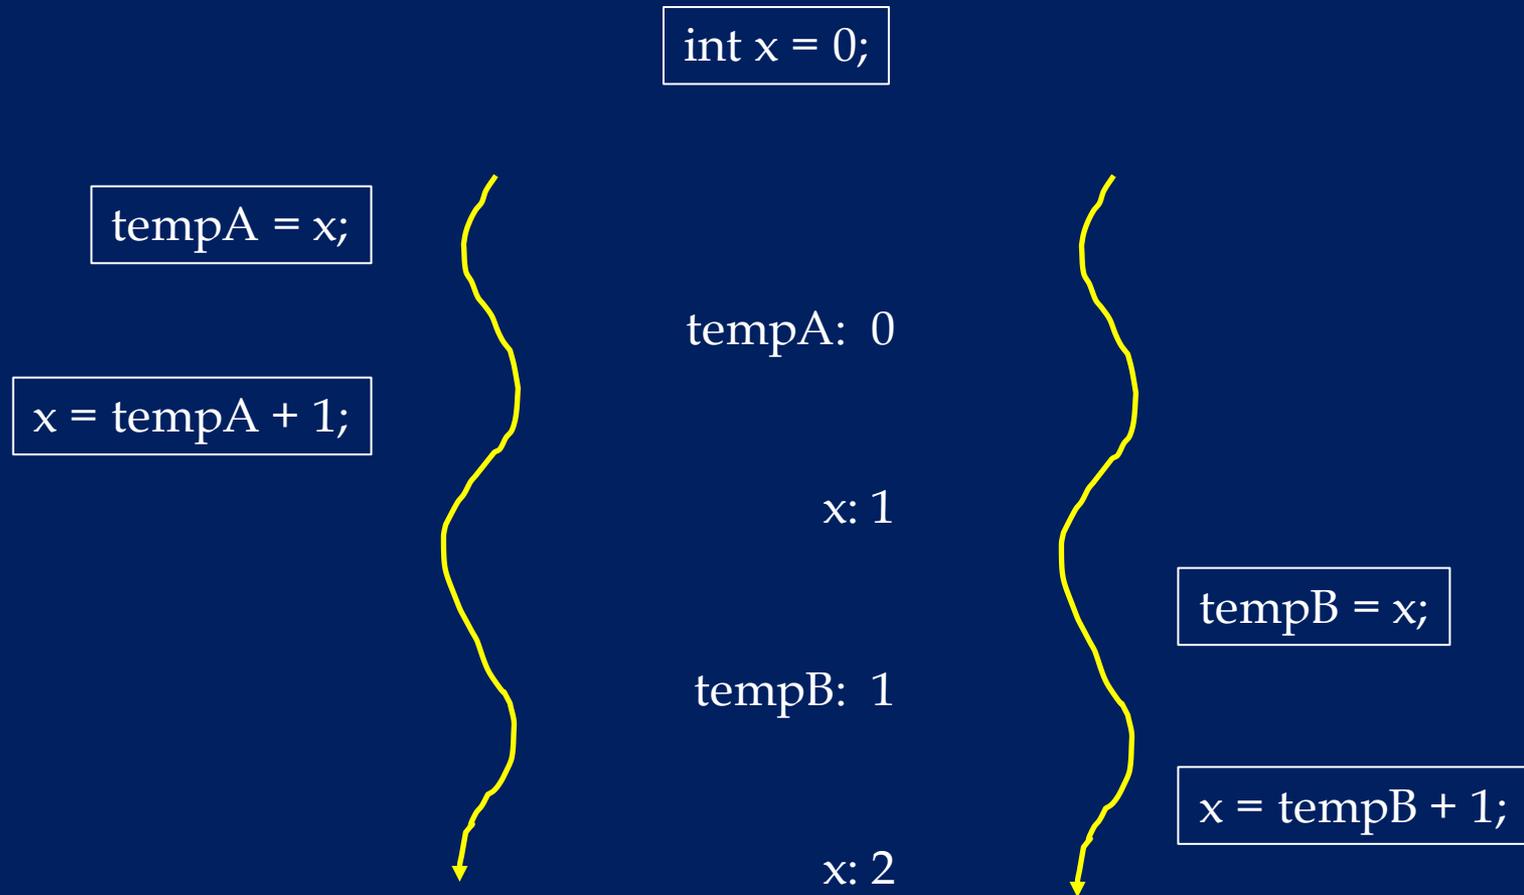


```
tempB = x;
```

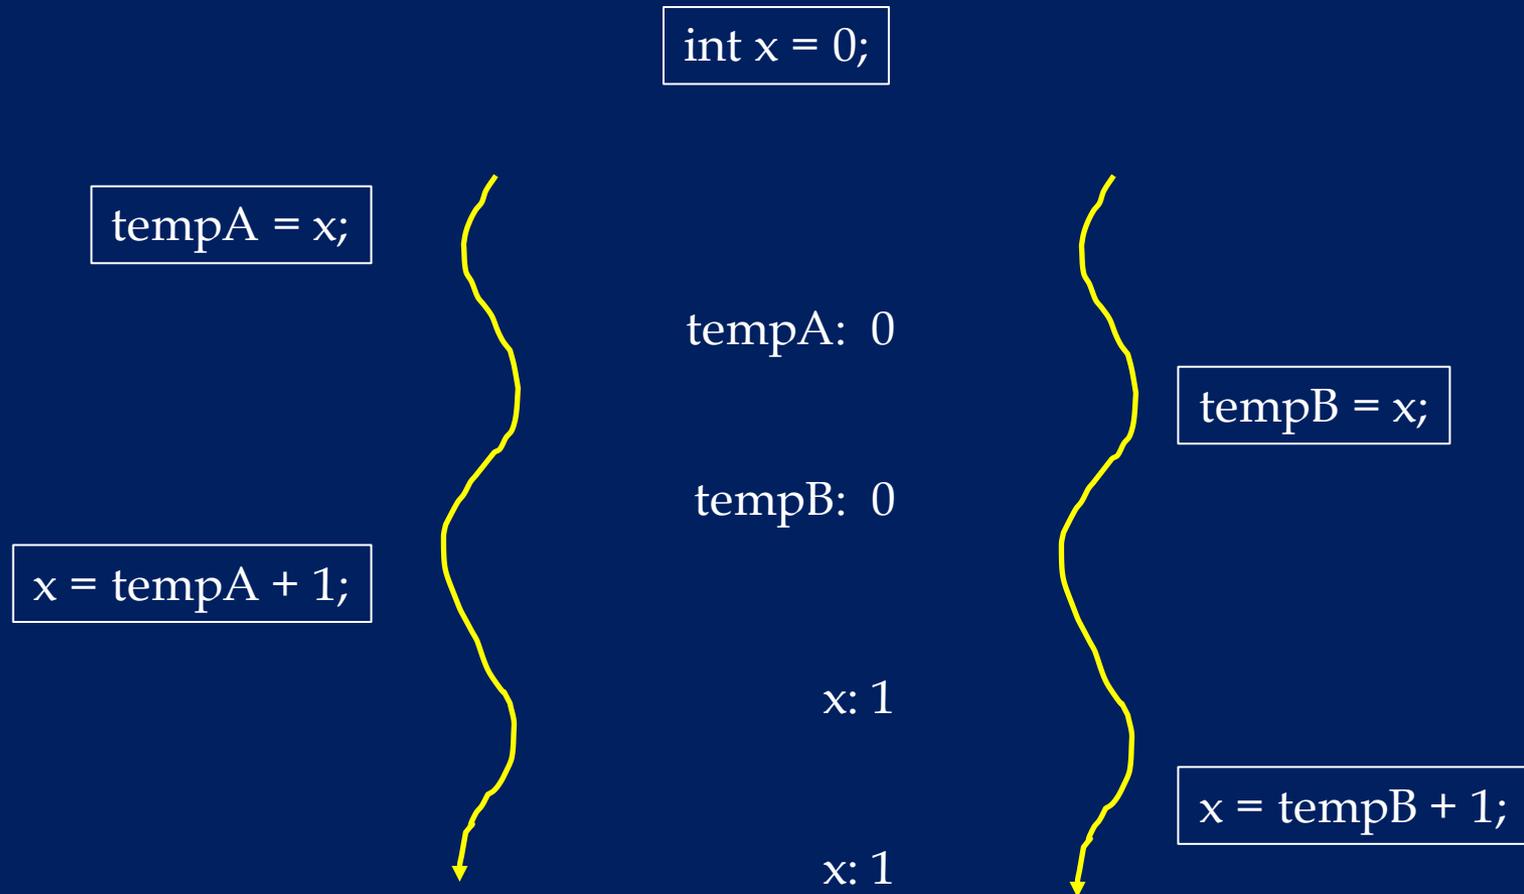
```
x = tempB + 1;
```



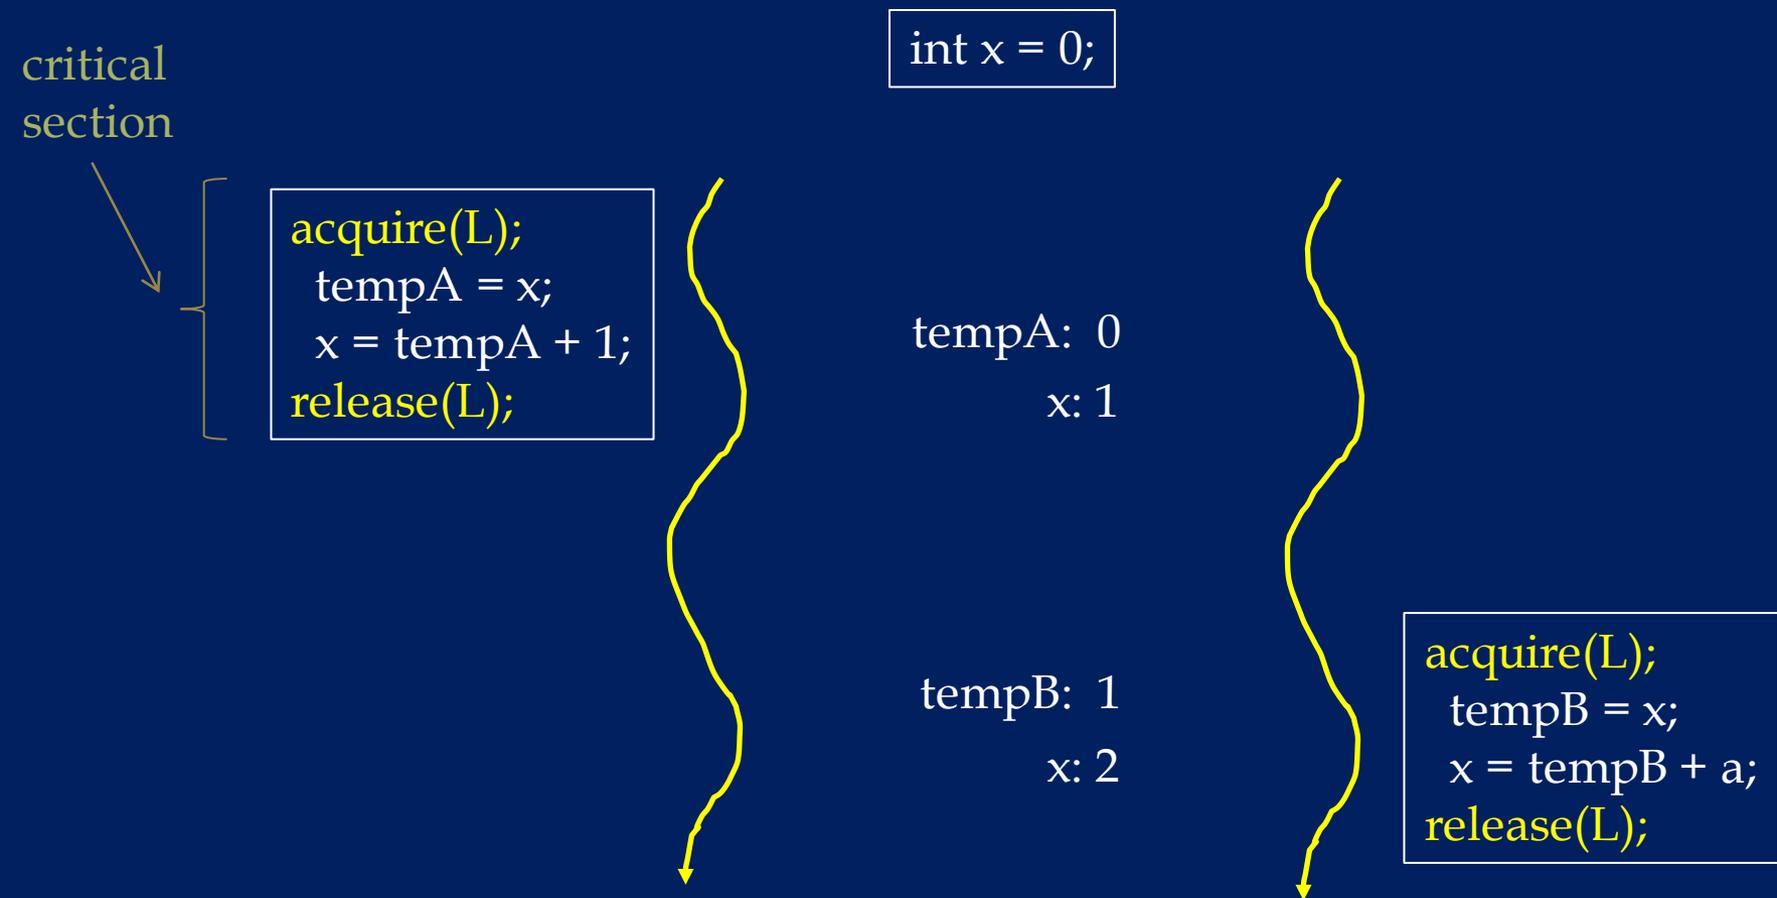
A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**



A Quick Primer on Why Imperative Parallel Programming is Hard: **1 + 1 ain't always 2!**



A Quick Primer on Why Imperative Parallel Programming is Hard: Locks and Critical Sections



A Quick Primer on Why Imperative Parallel Programming is Hard: Synchronized Methods

acquires and releases the lock
associated with the object (1 lock per object)

```
Adder a = new Adder(0);
```

Java Synchronized Methods:

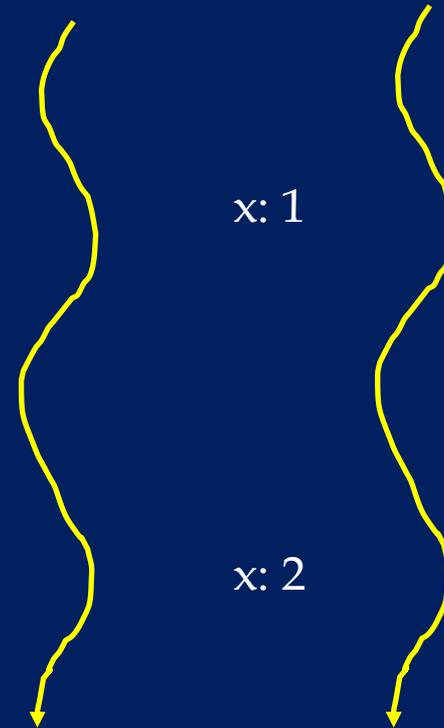
```
class Adder{  
    int x;  
  
    synchronized void add() {  
        x = x+1;  
    }  
}
```

a.add()

x: 1

x: 2

a.add()



What's wrong with locks?

Correct use of locks can solve concurrency problems, but locks are amazingly difficult to use correctly

- **Races**: forgotten locks (or synchronization commands) lead to inconsistent views
- **Deadlock**: locks acquired in “wrong” order
- **Lost wakeups**: forget to notify condition variables
- **Diabolical error recovery**: need to restore invariants and release locks in exception handlers. Yikes!
- These are serious problems. But even worse...

Locks are Non-Compositional

- Consider a (correct) Java bank **Account** class:

```
class Account{
    float balance;

    synchronized void deposit(float amt) {
        balance += amt;
    }

    synchronized void withdraw(float amt) {
        if (balance < amt)
            throw new OutOfMoneyError();
        balance -= amt;
    }
}
```

- Now suppose we want to add the ability to transfer funds from one account to another.

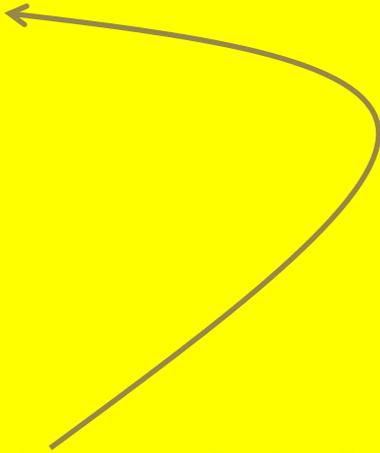
Locks are Non-Compositional

- Simply calling **withdraw** and **deposit** to implement **transfer** causes a race condition:

```
class Account{
    float balance;
    ...
    void badTransfer(Acct other, float amt) {
        other.withdraw(amt);

        this.deposit(amt);
    }
}
class Bank {
    Account[] accounts;
    float global_balance;

    checkBalances () {
        return (sum(Accounts) == global_balance);
    }
}
```



sees bad total balance value in between withdraw and deposit

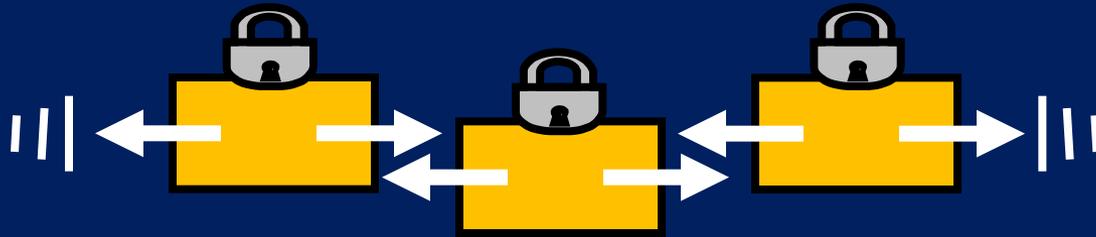
Locks are Non-Compositional

- Synchronizing **transfer** can cause deadlock:

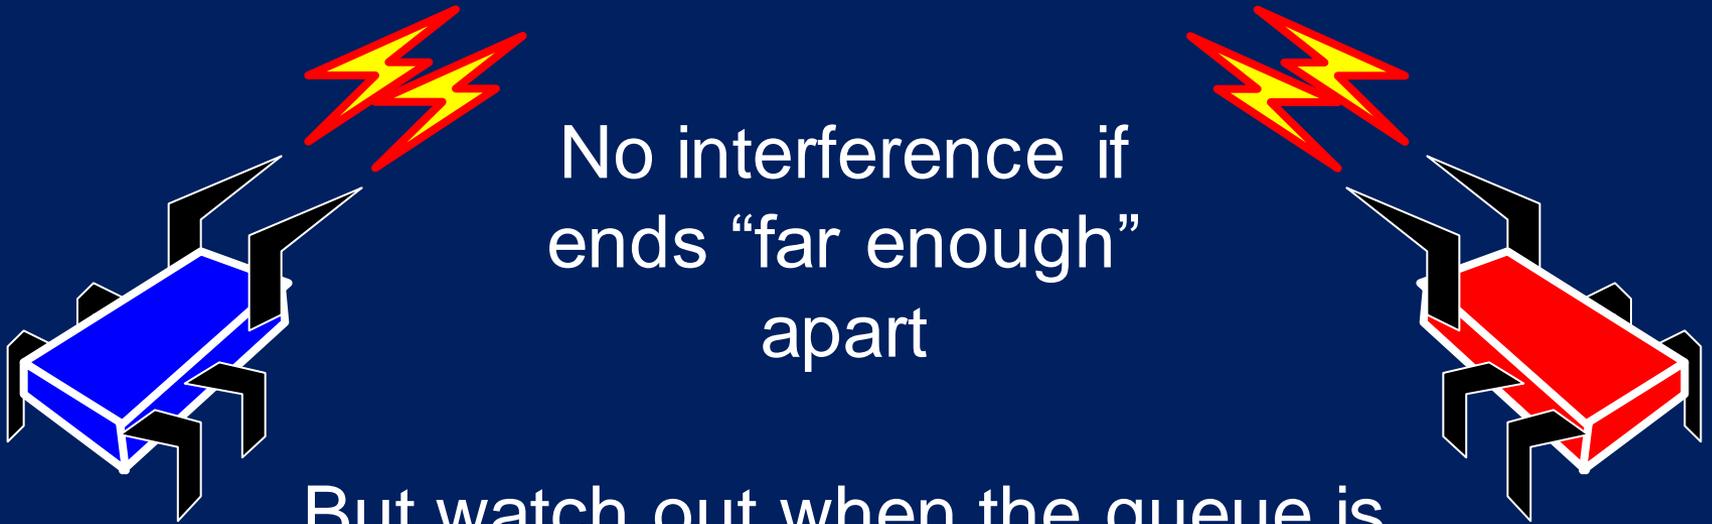
```
class Account{
    float balance;
    synchronized void deposit(float amt) {
        balance += amt;
    }
    synchronized void withdraw(float amt) {
        if(balance < amt)
            throw new OutOfMoneyError();
        balance -= amt;
    }
    synchronized void badTrans(Acct other, float amt) {
        // can deadlock with parallel reverse-transfer
        this.deposit(amt);
        other.withdraw(amt);
    }
}
```

Locks are absurdly hard to get right

Scalable double-ended queue: one lock per cell



No interference if
ends “far enough”
apart



But watch out when the queue is
0, 1, or 2 elements long!

Locks are absurdly hard to get right

Coding style	Difficulty of queue implementation
Sequential code	Undergraduate (COS 226)

Locks are absurdly hard to get right

Coding style	Difficulty of queue implementation
Sequential code	Undergraduate (COS 226)
Efficient parallel code with locks and condition variables	Publishable result at international conference ¹

¹ Simple, fast, and practical non-blocking and blocking concurrent queue algorithms.

What does Haskell do about this?

- Caveat: Nobody has the general answer. It is a huge area of current research.
- Haskell has an edge on Java and C because data structures are **immutable by default** and the problems appear when two parallel threads are actively mutating shared data or other resources
- Haskell provides a tractor-trailer's worth of options:
 - parallelism over immutable data via “sparks”
 - parallelism over immutable data via data-parallel operators like parallel map and parallel fold (aka reduce)
 - Google's map-reduce architecture borrows from older functional programming just like Haskell
 - software transactional memory
 - ordinary locks & threads (boo!)

What does Haskell do about this?

- Caveat: Nobody has the general answer. It is a huge area of current research.
- Haskell has an edge on Java and C because data structures are **immutable by default** and the problems appear when two parallel threads are actively mutating shared data or other resources
- Haskell provides a tractor-trailer's worth of options:
 - parallelism over immutable data via “sparks”
 - parallelism over immutable data via data-parallel operators like parallel map
 - Google's map-reduce architecture borrows from older functional programming just like Haskell
 - **software transactional memory**
 - ordinary locks & threads (boo!)

we will look at this today



Software Transactional Memory (STM)

The Punchline for STM

Coding style	Difficulty of queue implementation
Sequential code	Undergraduate (COS 226)
Efficient parallel code with locks and condition variables	Publishable result at international conference ¹
Parallel code with STM	<i>Undergraduate</i>

¹ Simple, fast, and practical non-blocking and blocking concurrent queue algorithms.

STM = Atomic Memory Transactions

Like database transactions

```
atomic { ...sequential code... }
```

- To a first approximation, just write the sequential code, and wrap **atomic** around it
- All-or-nothing semantics: **Atomic** commit
- Atomic block executes in **Isolation**
 - with automatic retry if another conflicting atomic block interferes
- Cannot deadlock (there are no locks!)
 - guarantees about progress on retry
- Atomicity makes error recovery easy (e.g. throw exception inside sequential code)

ACID

How does it work?

```
read y;  
read z;  
write 10 x;  
write 42  
z;  
...
```

```
atomic { ... <code> ... }
```

One possibility:

- Execute `<code>` *optimistically* without taking any locks.
- Log each read and write in `<code>` to a thread-local transaction log.
- Writes go to the log only, not to memory.
- At the end, the transaction validates the log.
 - Validation: Are the values I read the same now as when I read them?
 - If valid, atomically **commits changes** to memory.
 - If not valid, re-runs from the beginning, discarding changes.

Realizing STM in Haskell

Why STM in Haskell?

- Logging memory effects is **expensive**.
- Haskell already partitions the world into
 - immutable values (zillions and zillions)
 - mutable locations (some or none)

Only need to log the latter!
- **Type system**: Controls where I/O effects happen.
- **Monad infrastructure**: Ideal for constructing transactions & implicitly passing transaction log.
- **Already paid the bill**: Simply reading or writing a mutable location is expensive (involving a procedure call) so transaction overhead is not as large as in an imperative language.

Haskell programmers brutally trained from birth to use memory effects sparingly.

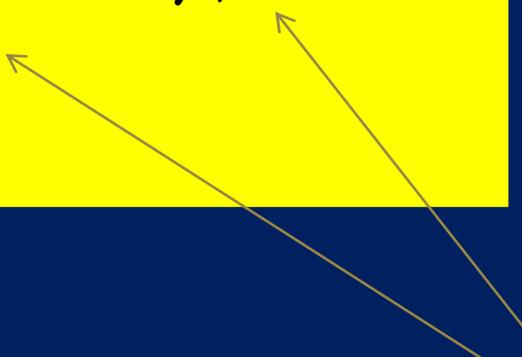
Concurrent Threads in Haskell

- The **fork** function spawns a thread.
- It takes an action as its argument.

```
fork :: IO a -> IO ThreadId
```

```
main = do { fork (action1);  
            action2  
            ... }
```

action 1 and
action 2 in
parallel



Atomic Blocks in Haskell

- **Idea:** add a function `atomic` that guarantees atomic execution of its argument computation atomically.

```
main = do { r <- new 0;
            fork (atomic (action1));
            atomic (action2);
            ... }
```

Atomic Details

- Introduce a type for imperative transaction variables (**TVar**) and a new Monad (**STM**) to track transactions.
- Ensure **TVars** can only be modified in transactions.

```
atomic      :: STM a -> IO a
newTVar     :: a -> STM (TVar a)
readTVar    :: TVar a -> STM a
writeTVar   :: TVar a -> a -> STM ()
```

```
-- inc adds 1 to the mutable reference r
```

```
inc :: TVar Int -> STM ()
```

```
inc r = do { v <- readTVar r; writeTVar r (v+1) }
```

```
main = do { r <- atomic (newTVar 0);
            fork (atomic (inc r))
            atomic (inc r);
            ... }
```

STM in Haskell

```
atomic      :: STM a -> IO a
newTVar     :: a -> STM (TVar a)
readTVar    :: TVar a -> STM a
writeTVar   :: TVar a -> a -> STM ()
```

The STM monad includes different ops than the IO monad:

- Can't use TVars outside atomic block
 - just like you can't use `printStrLn` outside of IO monad
- Can't do IO inside atomic block:

```
atomic (if x<y then launchMissiles)
```

- **atomic** is a function, not a syntactic construct
 - called *atomically* in the actual implementation
- ...and, best of all...

STM Computations Compose (unlike locks)

```
incT  :: TVar Int -> STM ()
incT r = do { v <- readTVar r;
             writeTVar r (v+1) }

incT2 :: TVar Int -> STM ()
incT2 r = do { incT r; incT r }

foo  :: IO ()
foo = ...atomic (incT2 r)...
```

Composition
is THE way
to build big
programs
that work

- The type guarantees that an **STM** computation is always executed atomically.
 - Glue **STMs** together arbitrarily as many times as you want
 - Then wrap with **atomic** to produce an IO action.

Exceptions

- The **STM** monad supports exceptions:

```
throw :: Exception -> STM a
catch :: STM a ->
       (Exception -> STM a) -> STM a
```

- In the call (**atomic s**), if **s** throws an exception, the transaction is aborted with no effect and the exception is propagated to the enclosing IO code.
- **No need to restore invariants, or release locks!**
- See “**Composable Memory Transactions**” for more information.

Three more combinators:
retry, orElse, always

Idea 1: Compositional Blocking

```
withdraw :: TVar Int -> Int -> STM ()
withdraw acc n =
    do { bal <- readTVar acc;
        if bal < n then retry;
        writeTVar acc (bal-n) }
```

```
retry :: STM ()
```

- **retry** means “abort the current transaction and re-execute it from the beginning”.
- Implementation avoids early retry using reads in the transaction log (i.e. **acc**) to wait on all read variables.
 - ie: retry only happens when one of the variables read on the path to the retry changes

Compositional Blocking

```
withdraw :: TVar Int -> Int -> STM ()
withdraw acc n =
    do { bal <- readTVar acc;
        if bal < n then retry;
        writeTVar acc (bal-n) }
```

- No condition variables!
- Retrying thread is woken up automatically when **acc** is written, so there is no danger of forgotten notifies.
- No danger of forgetting to test conditions again when woken up because the transaction runs from the beginning. For example:
`atomic (do { withdraw a1 3;
 withdraw a2 7 })`

What makes Retry Compositional?

- **retry** can appear anywhere inside an atomic block, including nested deep within a call. For example,

```
atomic (do { withdraw a1 3;  
           withdraw a2 7 })
```

waits for a1's balance >3 AND a2's balance >7, **without any change to withdraw function.**

Idea 2: Choice

- Suppose we want to transfer 3 dollars from either account a1 or a2 into account b.

```
atomic (do {  
  withdraw a1 3  
  `orelse`  
  withdraw a2 3;  
  deposit b 3 })
```

Try this

...and if it retries,
try this

...and then
do this

```
orElse :: STM a -> STM a -> STM a
```

Choice is composable, too!

```
transfer :: TVar Int ->  
          TVar Int ->  
          TVar Int ->  
          STM ()
```

```
transfer a1 a2 b = do  
  { withdraw a1 3  
    `orElse`  
    withdraw a2 3;  
  
    deposit b 3 }
```

```
atomic  
  (transfer a1 a2 b  
   `orElse`  
   transfer a3 a4 b)
```

- The function `transfer` calls `orElse`, but calls to `transfer` can still be composed with `orElse`.

Equational Reasoning

- STM supports nice equations for reasoning:
 - `orElse` is associative (but not commutative)
 - `retry `orElse` s = s`
 - `s `orElse` retry = s`
- These equations make STM an instance of the Haskell typeclass `MonadPlus`, a `Monad` with some extra operations and properties.

Idea 3: Invariants

- The route to sanity is to establish **invariants** that are **assumed on entry**, and **guaranteed on exit**, by *every atomic block*.
- We want to check these guarantees. But we don't want to test every invariant after every atomic block.
- Hmm.... Only test when something read by the invariant has changed.... rather like **retry**.

Invariants: One New Primitive

```
always :: STM Bool -> STM ()
```

```
newAccount :: STM (TVar Int)
```

```
newAccount =
```

```
  do { v <- newTVar 0;  
      always (accountInv);  
      return v }
```

An arbitrary boolean valued
STM computation

```
accountInv = do { cts <- readTVar v;  
                 return (cts >= 0) };
```

Any transaction that modifies the account will check the invariant (no forgotten checks). If the check fails, the transaction restarts.

What does it all mean?

- Everything so far is intuitive and arm-wavey.
- But what happens if it's raining, and you are inside an `orElse` and you throw an exception that contains a value that mentions...?
- We need a precise specification!

One
exists

IO transitions $P; \Theta \xrightarrow{c} Q; \Theta'$

$$\begin{array}{l} P[\text{putChar } c]; \Theta \xrightarrow{!c} P[\text{return } ()]; \Theta \quad (\text{PUTC}) \\ P[\text{getChar}]; \Theta \xrightarrow{?c} P[\text{return } c]; \Theta \quad (\text{GETC}) \\ P[\text{forkIO } M]; \Phi, \Delta \rightarrow (P[\text{return } r] \mid M_r); \Phi, \Delta \cup \{r\} \quad r \notin \Delta \quad (\text{FORK}) \end{array}$$

$$\frac{M \rightarrow N}{P[M]; \Theta \rightarrow P[N]; \Theta} \quad (\text{ADMIN})$$

$$\frac{M; \Theta \xrightarrow{\Delta} \text{return } N; \Theta'}{P[\text{atomically } M]; \Theta \rightarrow P[\text{return } N]; \Theta'} \quad (\text{ARET}) \quad \frac{M; \Phi, \Delta \xrightarrow{\Delta} \text{throw } N; \Phi, \Delta'}{P[\text{atomically } M]; \Phi, \Delta \rightarrow P[\text{throw } N]; \Phi, \Delta'} \quad (\text{ATHROW})$$

Administrative transitions $M \rightarrow N$

$$\begin{array}{l} M \rightarrow V \quad \text{if } \mathcal{E}[[M]] = V \text{ and } M \neq V \quad (\text{EVAL}) \\ \text{return } N \gg M \rightarrow MN \quad (\text{BIND}) \\ \text{throw } N \gg M \rightarrow \text{throw } N \quad (\text{THROW}) \\ \text{catch } (\text{throw } M) N \rightarrow NM \quad (\text{CATCH1}) \\ \text{catch } (\text{return } M) N \rightarrow \text{return } M \quad (\text{CATCH2}) \end{array}$$

STM transitions $M; \Theta \Rightarrow N; \Theta'$

$$\begin{array}{l} E[\text{readTVar } r]; \Phi, \Delta \Rightarrow E[\text{return } \Phi(r)]; \Phi, \Delta \quad \text{if } r \in \text{dom}(\Phi) \quad (\text{READ}) \\ E[\text{writeTVar } r N]; \Phi, \Delta \Rightarrow E[\text{return } ()]; \Phi[r \mapsto M], \Delta \quad \text{if } r \in \text{dom}(\Phi) \quad (\text{WRITE}) \\ E[\text{newTVar } M]; \Phi, \Delta \Rightarrow E[\text{return } r]; \Phi[r \mapsto M], \Delta \cup \{r\} \quad \text{if } r \notin \Delta \quad (\text{NEW}) \end{array}$$

$$\frac{M \rightarrow N}{E[M]; \Theta \Rightarrow E[N]; \Theta} \quad (\text{ADMIN})$$

$$\frac{E[M_1]; \Theta \xrightarrow{\Delta} E[\text{return } N]; \Theta'}{E[M_1 \text{ `orElse` } M_2]; \Theta \Rightarrow E[\text{return } N]; \Theta'} \quad (\text{OR1}) \quad \frac{E[M_1]; \Theta \xrightarrow{\Delta} E[\text{throw } N]; \Theta'}{E[M_1 \text{ `orElse` } M_2]; \Theta \Rightarrow E[\text{throw } N]; \Theta'} \quad (\text{OR2})$$

$$\frac{E[M_1]; \Theta \xrightarrow{\Delta} E[\text{retry}]; \Theta'}{E[M_1 \text{ `orElse` } M_2]; \Theta \Rightarrow E[M_2]; \Theta} \quad (\text{OR3})$$

See “[Composable Memory Transactions](#)” for details.

STM in Mainstream Languages

- There are similar proposals for adding STM to Java and other mainstream languages.

```
class Account {
    float balance;
    void deposit(float amt) {
        atomic { balance += amt; }
    }
    void withdraw(float amt) {
        atomic {
            if(balance < amt) throw new OutOfMoneyError();
            balance -= amt; }
    }
    void transfer(Acct other, float amt) {
        atomic { // Can compose withdraw and deposit.
            other.withdraw(amt);
            this.deposit(amt); }
    }
}
```

Weak vs Strong Atomicity

- Unlike Haskell, type systems in mainstream languages don't control where effects occur.
- What happens if code outside a transaction conflicts with code inside a transaction?
 - **Weak Atomicity**: Non-transactional code can see **inconsistent** memory states. Programmer should avoid such situations by placing all accesses to shared state in transaction.
 - **Strong Atomicity**: Non-transactional code is guaranteed to see a consistent view of shared state. This guarantee may cause a performance hit.

For more information: "[Enforcing Isolation and Ordering in STM](#)"

Even in Haskell: Easier, But Not Easy.

- The essence of shared-memory concurrency is *deciding where critical sections should begin and end*. This is still a **hard problem**.
 - **Too small**: application-specific data races (Eg, may see deposit but not withdraw if transfer is not atomic).
 - **Too large**: delay progress because deny other threads access to needed resources.
- In Haskell, we can compose STM subprograms but at some point, we must decide to wrap an STM in "atomic"
 - When and where to do it can be a hard decision

Conclusions

- Atomic blocks (**atomic**, **retry**, **orElse**) dramatically raise the level of abstraction for concurrent programming.
- It is like using a high-level language instead of assembly code. Whole classes of low-level errors are eliminated.
- Not a silver bullet:
 - you can still write buggy programs;
 - concurrent programs are still harder than sequential ones
 - aimed only at shared memory concurrency, not message passing
- There is a performance hit, but it is usually acceptable in Haskell (and things can only get better as the research community focuses on the question.)

Course Conclusions

- The study of STMs brings together multiple threads of interest in this course:
 - functional programming
 - high-level abstractions
 - operational semantics
 - equational reasoning & proofs about programs
- The development of STM is an example of modern programming language research
- If you are interested, talk with Andrew Appel or I about independent work opportunities, including work involving other new parallel programming paradigms

End

What **always** does

```
always :: STM Bool -> STM ()
```

- The function **always** adds a new invariant to a global pool of invariants.
- Conceptually, every invariant is checked as every transaction commits.
- But the implementation checks only invariants that read TVars that have been written by the transaction
- ...and garbage collects invariants that are checking dead Tvars.

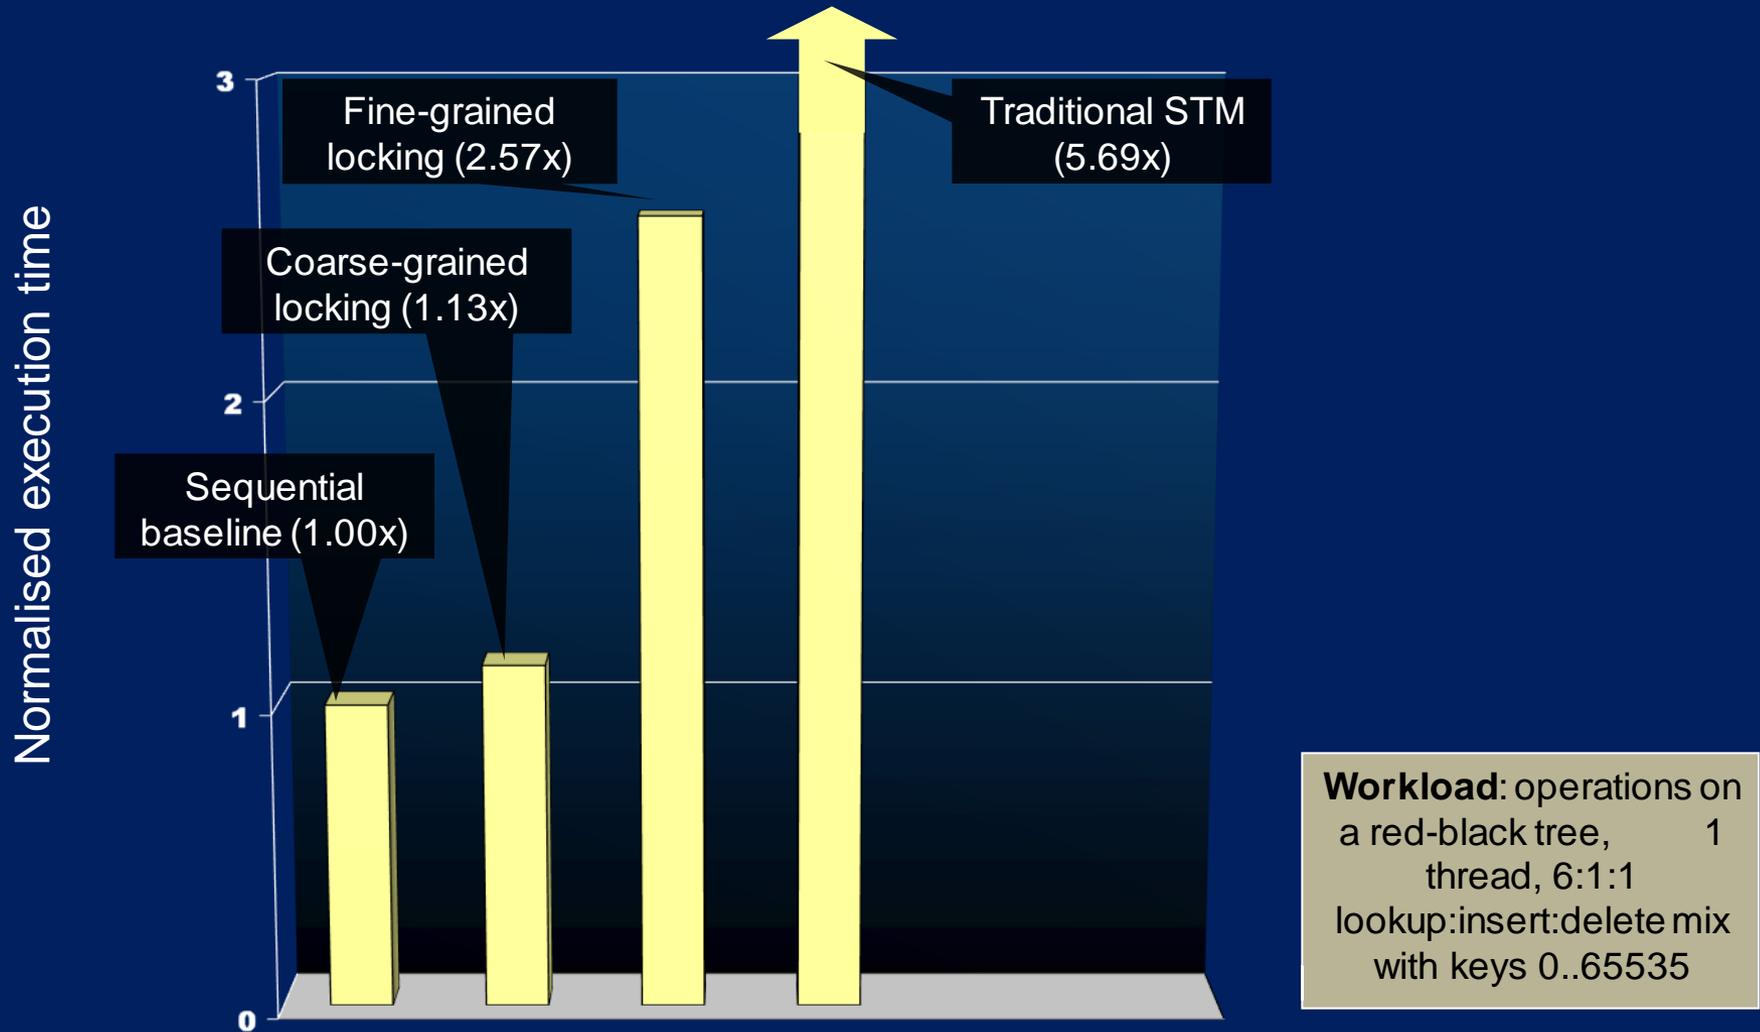
Haskell Implementation

- A complete, multiprocessor implementation of STM exists as of GHC 6.
- **Experience to date:** even for the most mutation-intensive program, the Haskell STM implementation is as fast as the previous MVar implementation.
 - The MVar version paid heavy costs for (usually unused) exception handlers.
- Need more experience using STM in practice, though!
- You can play with it. See the course website.

Performance

- At first, atomic blocks look insanely expensive. A naive implementation (c.f. databases):
 - Every load and store instruction logs information into a thread-local log.
 - A store instruction writes the log only.
 - A load instruction consults the log first.
 - Validate the log at the end of the block.
 - If succeeds, atomically commit to shared memory.
 - If fails, restart the transaction.

State of the Art Circa 2003



See "[Optimizing Memory Transactions](#)" for more information.

New Implementation Techniques

■ Direct-update STM

- Allows transactions to make updates in place in the heap
- Avoids reads needing to search the log to see earlier writes that the transaction has made
- Makes successful commit operations faster at the cost of extra work on contention or when a transaction aborts

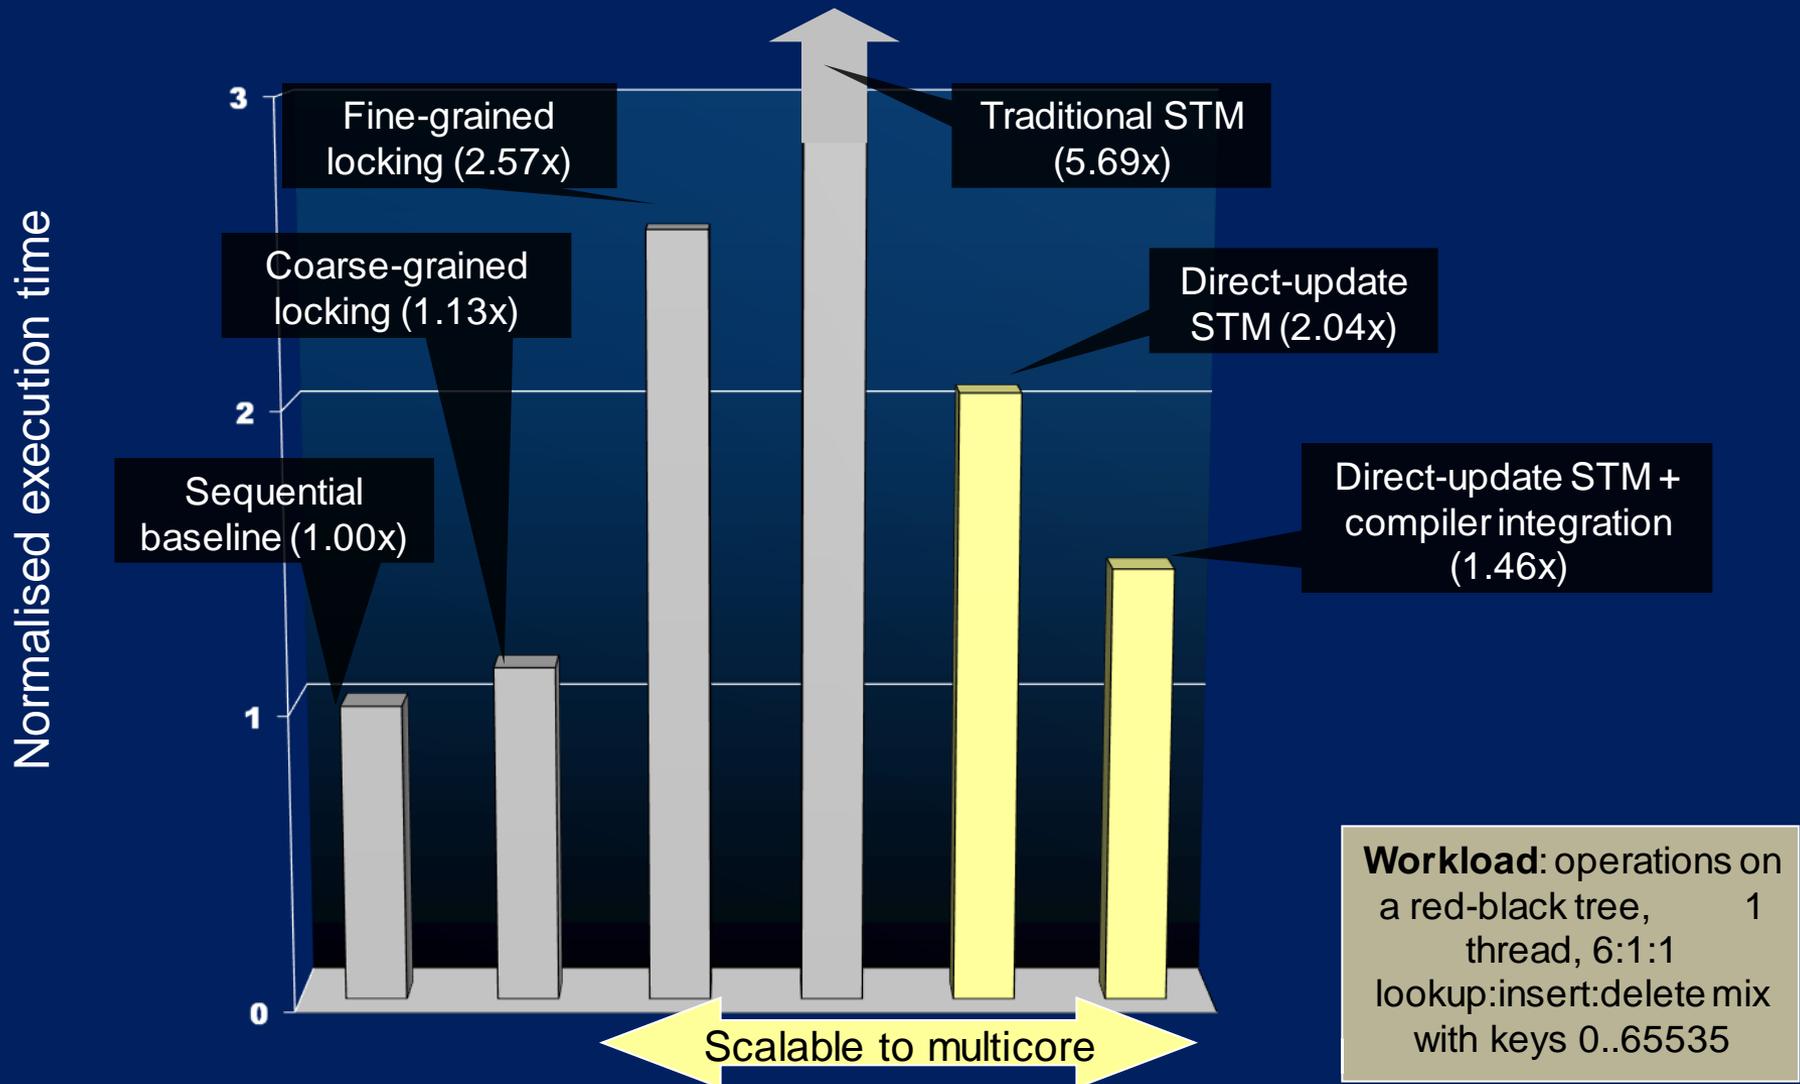
■ Compiler integration

- Decompose transactional memory operations into primitives
- Expose these primitives to compiler optimization (e.g. to hoist concurrency control operations out of a loop)

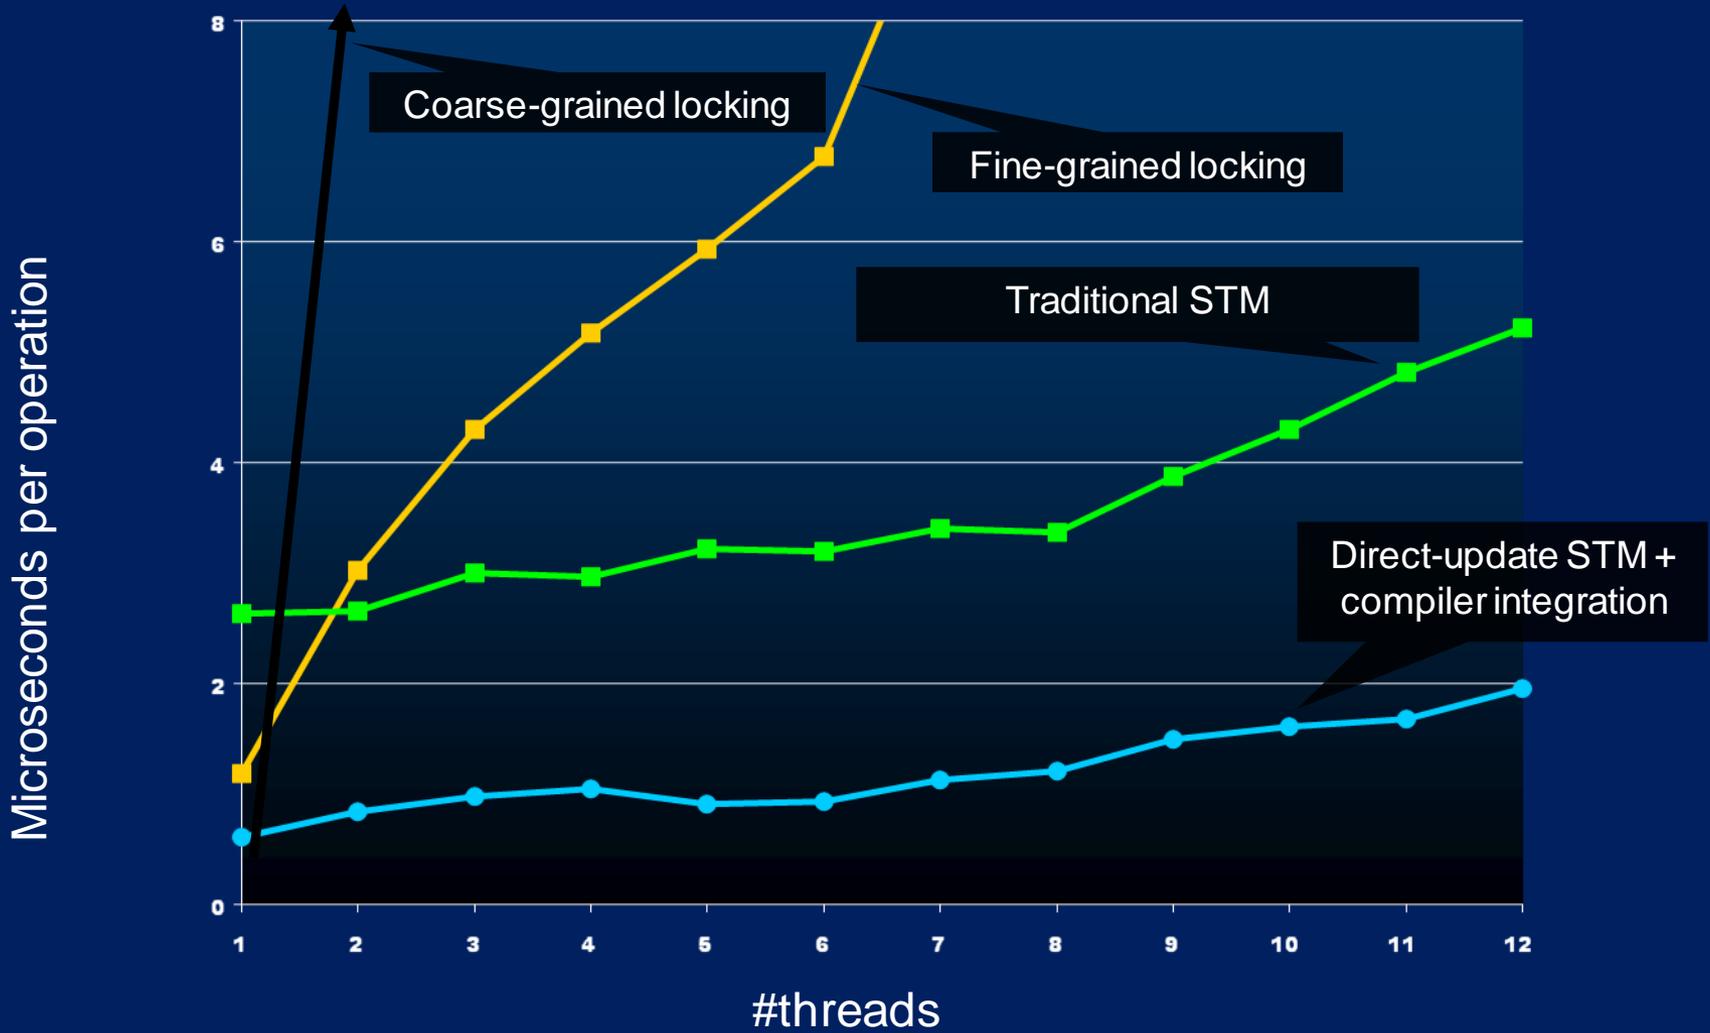
■ Runtime system integration

- Integrates transactions with the garbage collector to scale to atomic blocks containing 100M memory accesses

Results: Concurrency Control Overhead



Results: Scalability



Performance, Summary

- Naïve STM implementation is hopelessly inefficient.
- There is a lot of research going on in the compiler and architecture communities to optimize STM.
- This work typically assumes transactions are smallish and have low contention. If these assumptions are wrong, performance can degrade drastically.
- We need more experience with “real” workloads and various optimizations before we will be able to say for sure that we can implement STM sufficiently efficiently to be useful.

Still Not Easy, Example

- Consider the following program:

Initially, $x = y = 0$

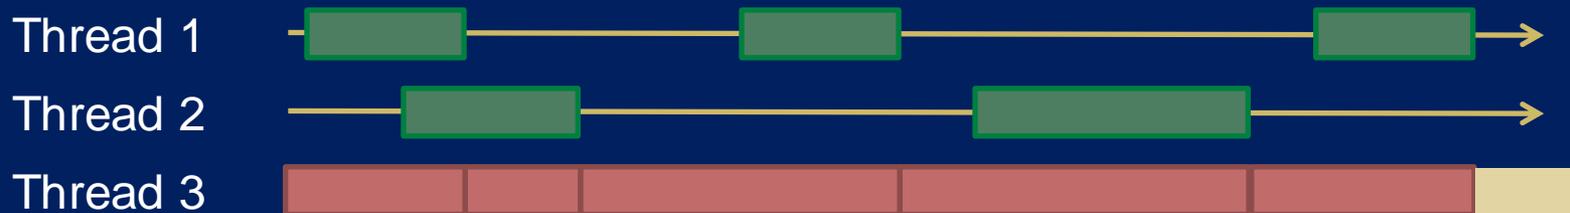
```
Thread 1
// atomic {                               //A0
    atomic { x = 1; }                       //A1
    atomic { if (y==0) abort; }           //A2
//}
```

```
Thread 2
atomic {                                     //A3
    if (x==0) abort;
    y = 1;
}
```

- Successful completion requires A3 to run after A1 but before A2.
- So adding a critical section (by uncommenting A0) changes the behavior of the program (from terminating to non-terminating).

Starvation

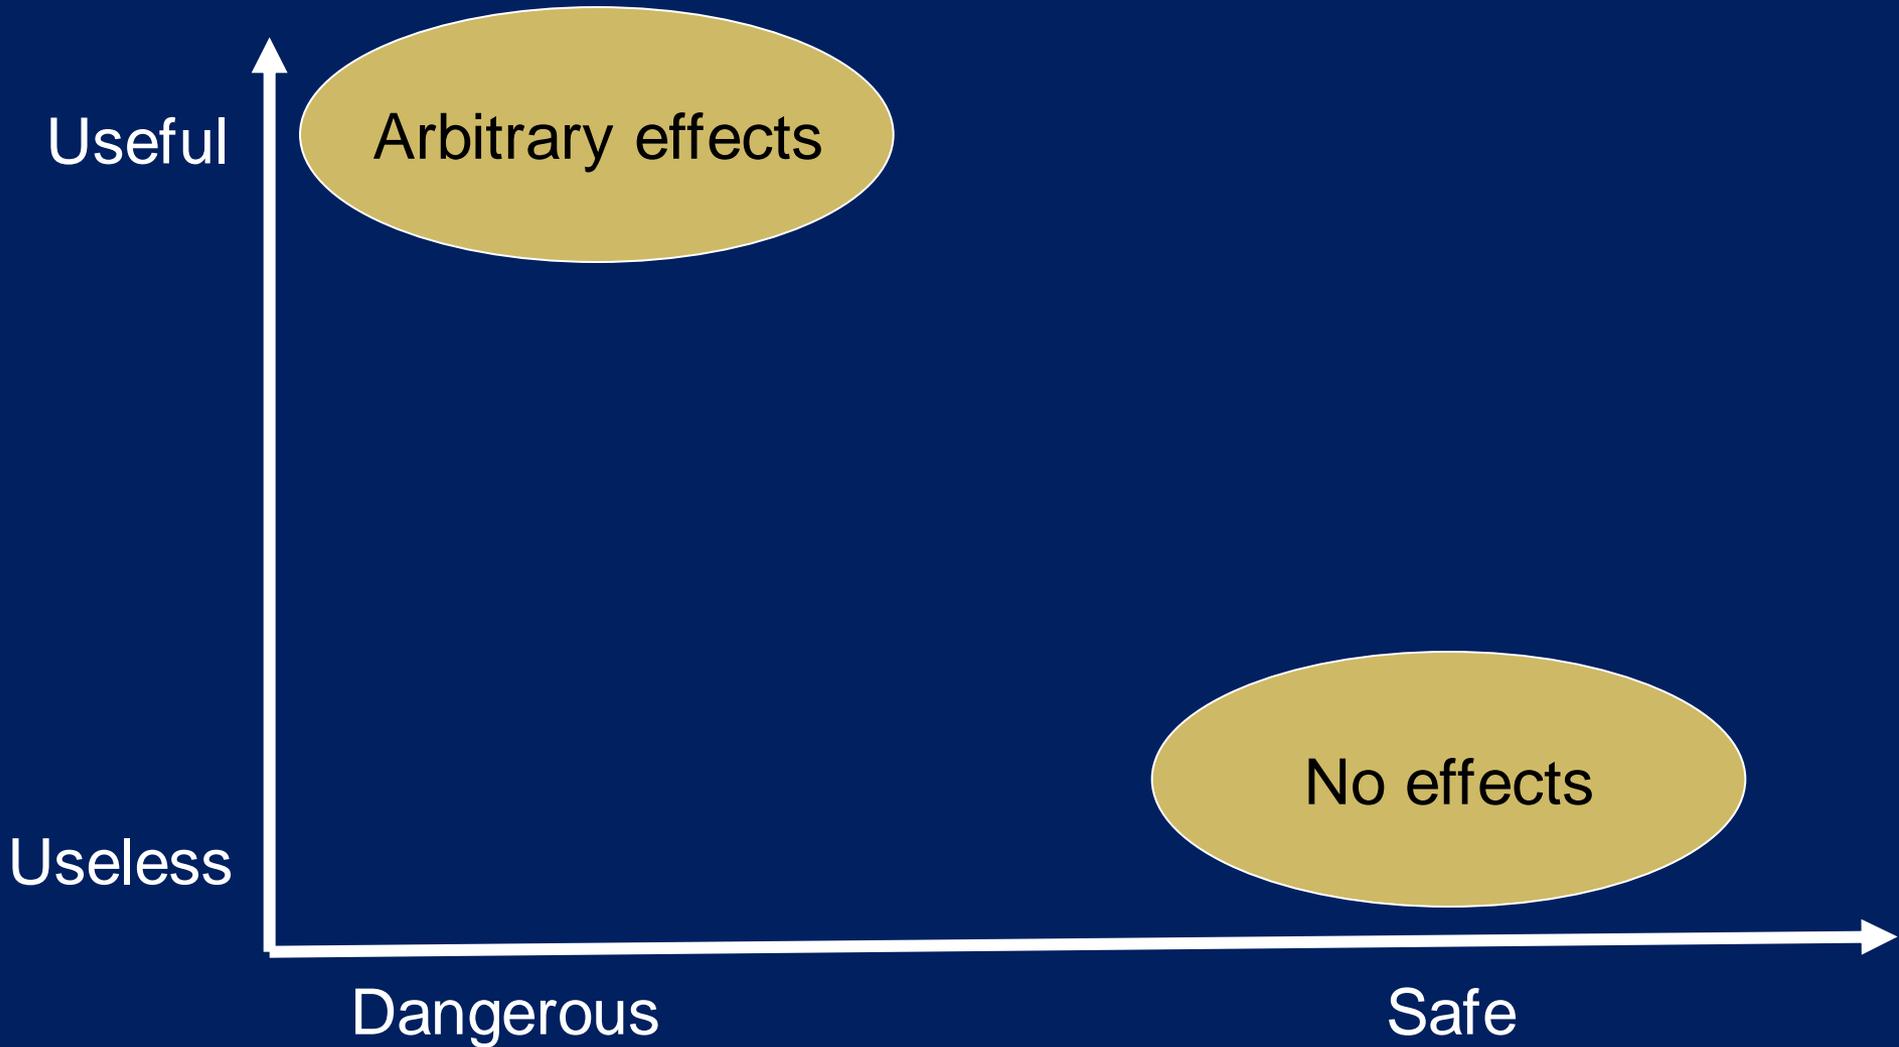
- **Worry:** Could the system “**thrash**” by continually colliding and re-executing?
- **No:** A transaction can be forced to re-execute only if another succeeds in committing. That gives a strong *progress guarantee*.
- **But:** A particular thread could **starve**:



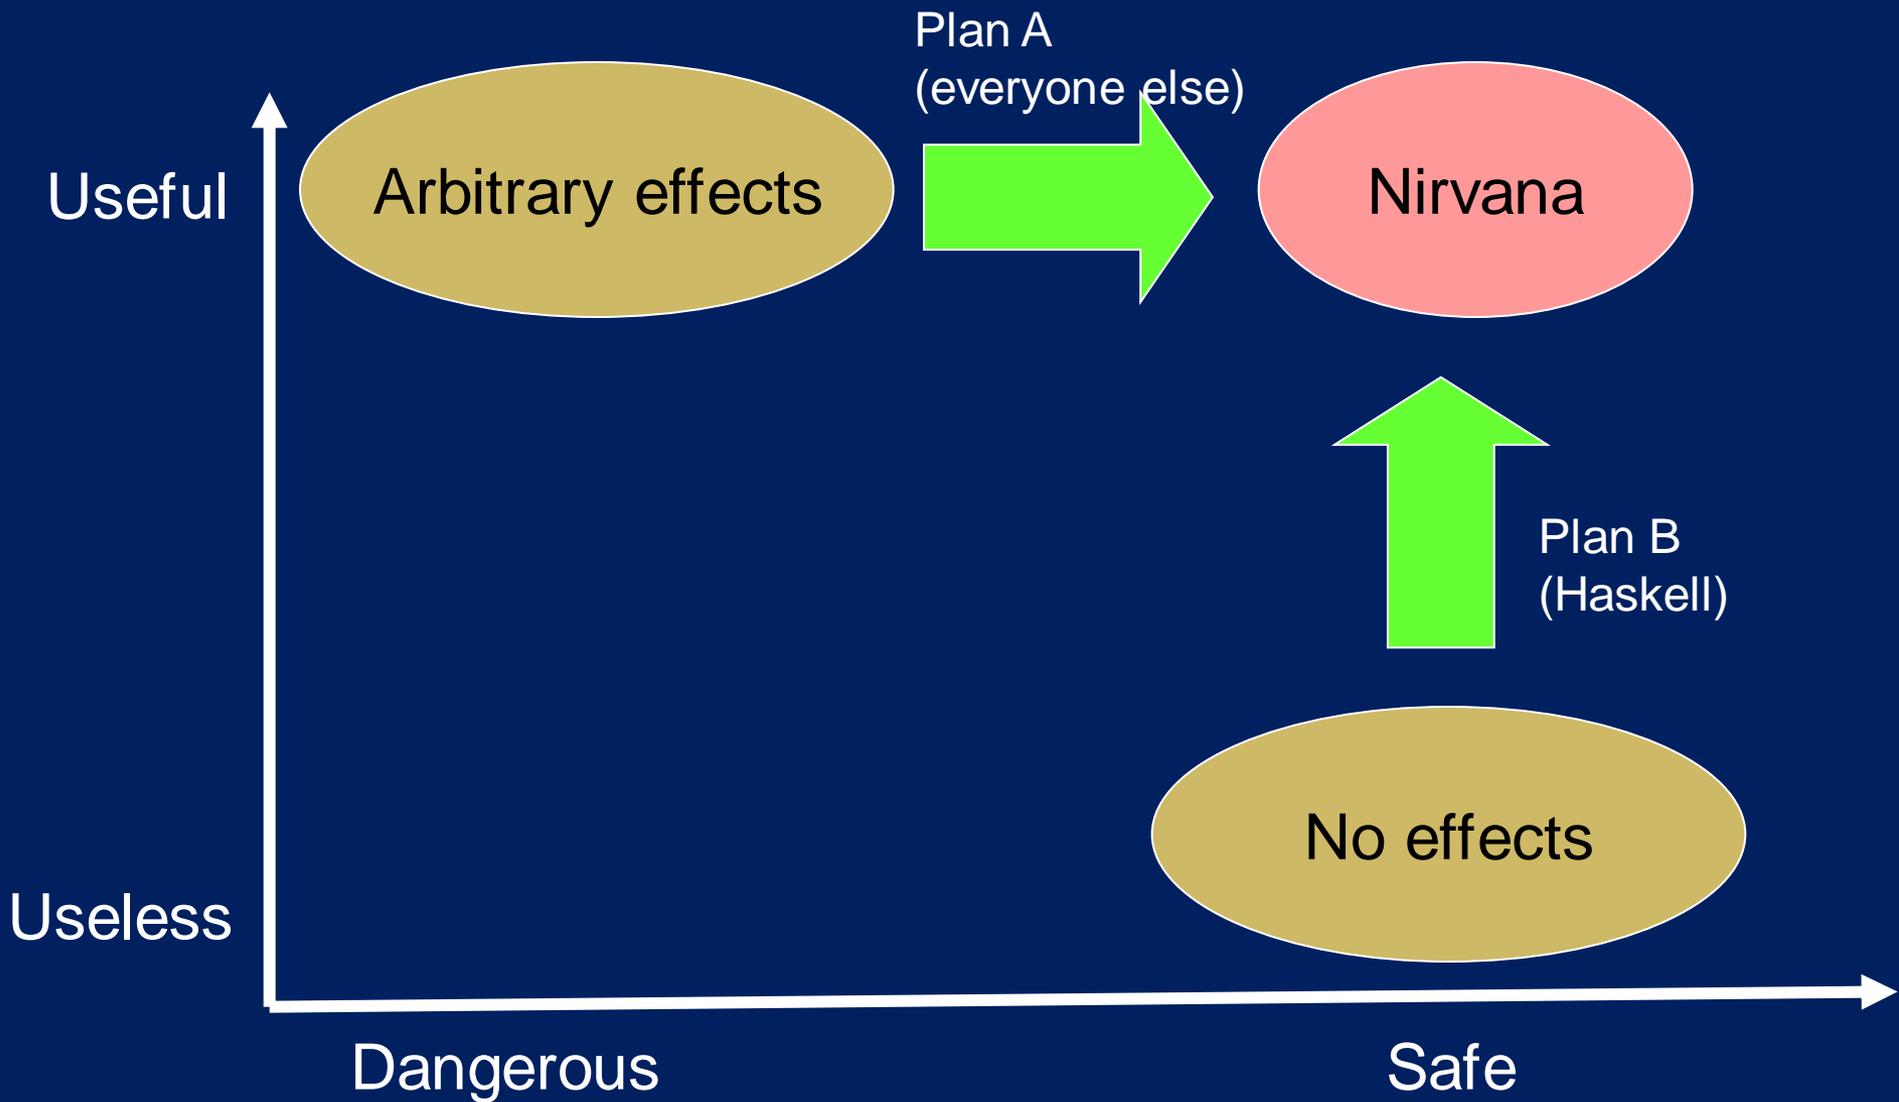
A Monadic Skin

- In languages like ML or Java, the fact that the language is in the IO monad is **baked in** to the language. There is no need to mark anything in the type system because IO is everywhere.
- In Haskell, the programmer can **choose** when to live in the IO monad and when to live in the realm of pure functional programming.
- **Interesting perspective**: It is not Haskell that lacks imperative features, but rather the other languages that lack the ability to have a statically distinguishable pure subset.
- This separation facilitates concurrent programming.

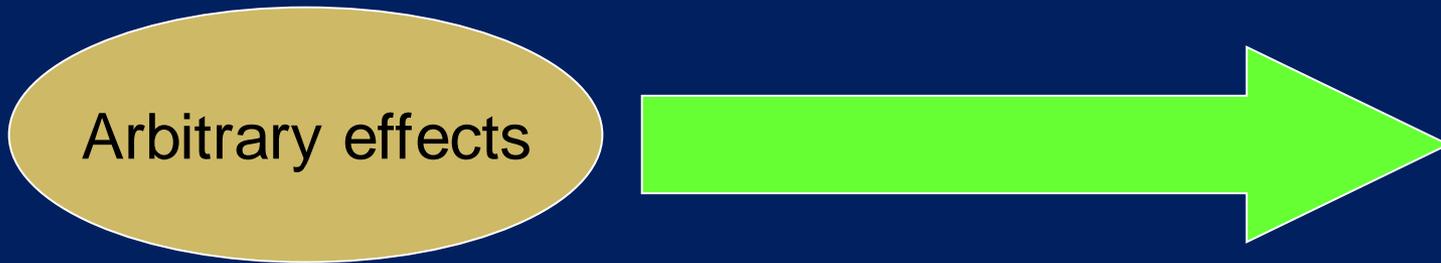
The Central Challenge



The Challenge of Effects



Two Basic Approaches: Plan A



Examples

- Regions
- Ownership types
- Vault, Spec#, Cyclone

Default = Any effect
Plan = Add restrictions

Two Basic Approaches: Plan B

Default = No effects

Plan = Selectively permit effects

Types play a major role

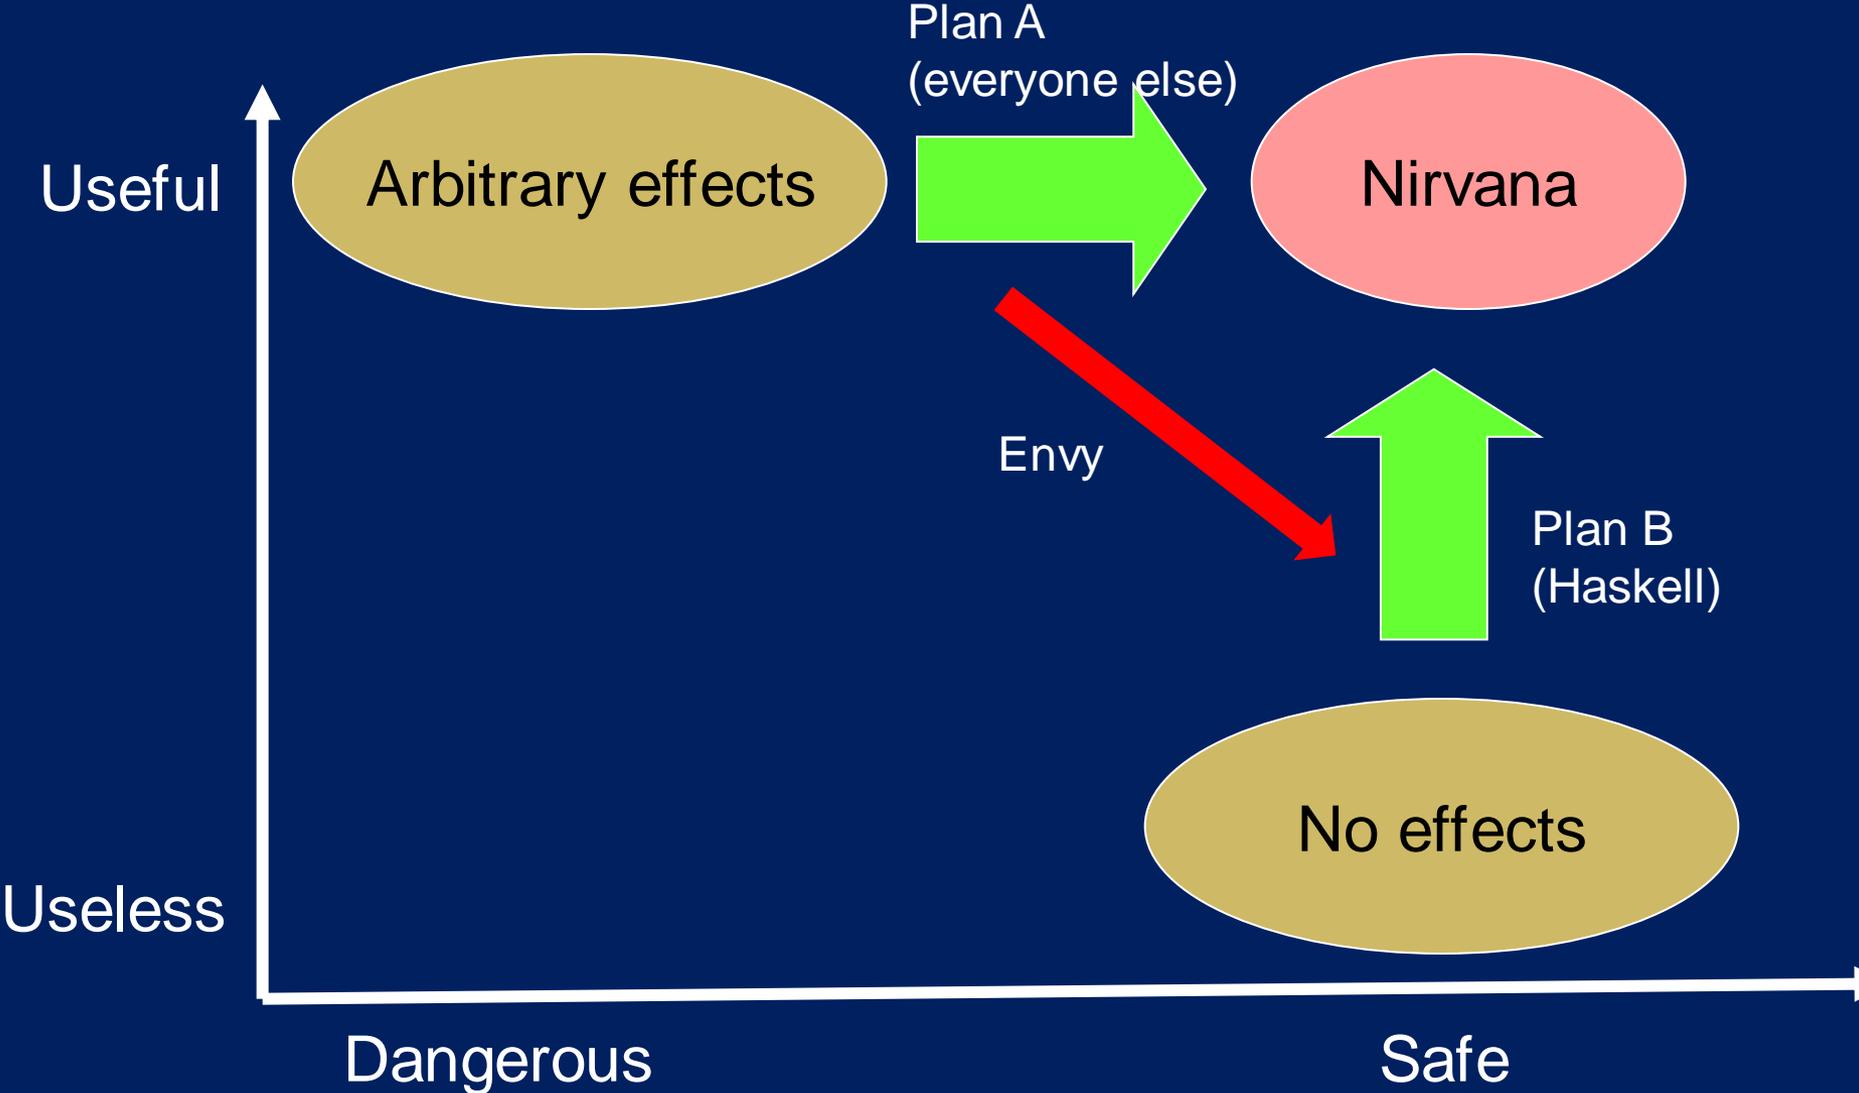
Two main approaches:

- Domain specific languages (SQL, Xquery, Google map/reduce)
- Wide-spectrum functional languages + controlled effects (e.g. Haskell)

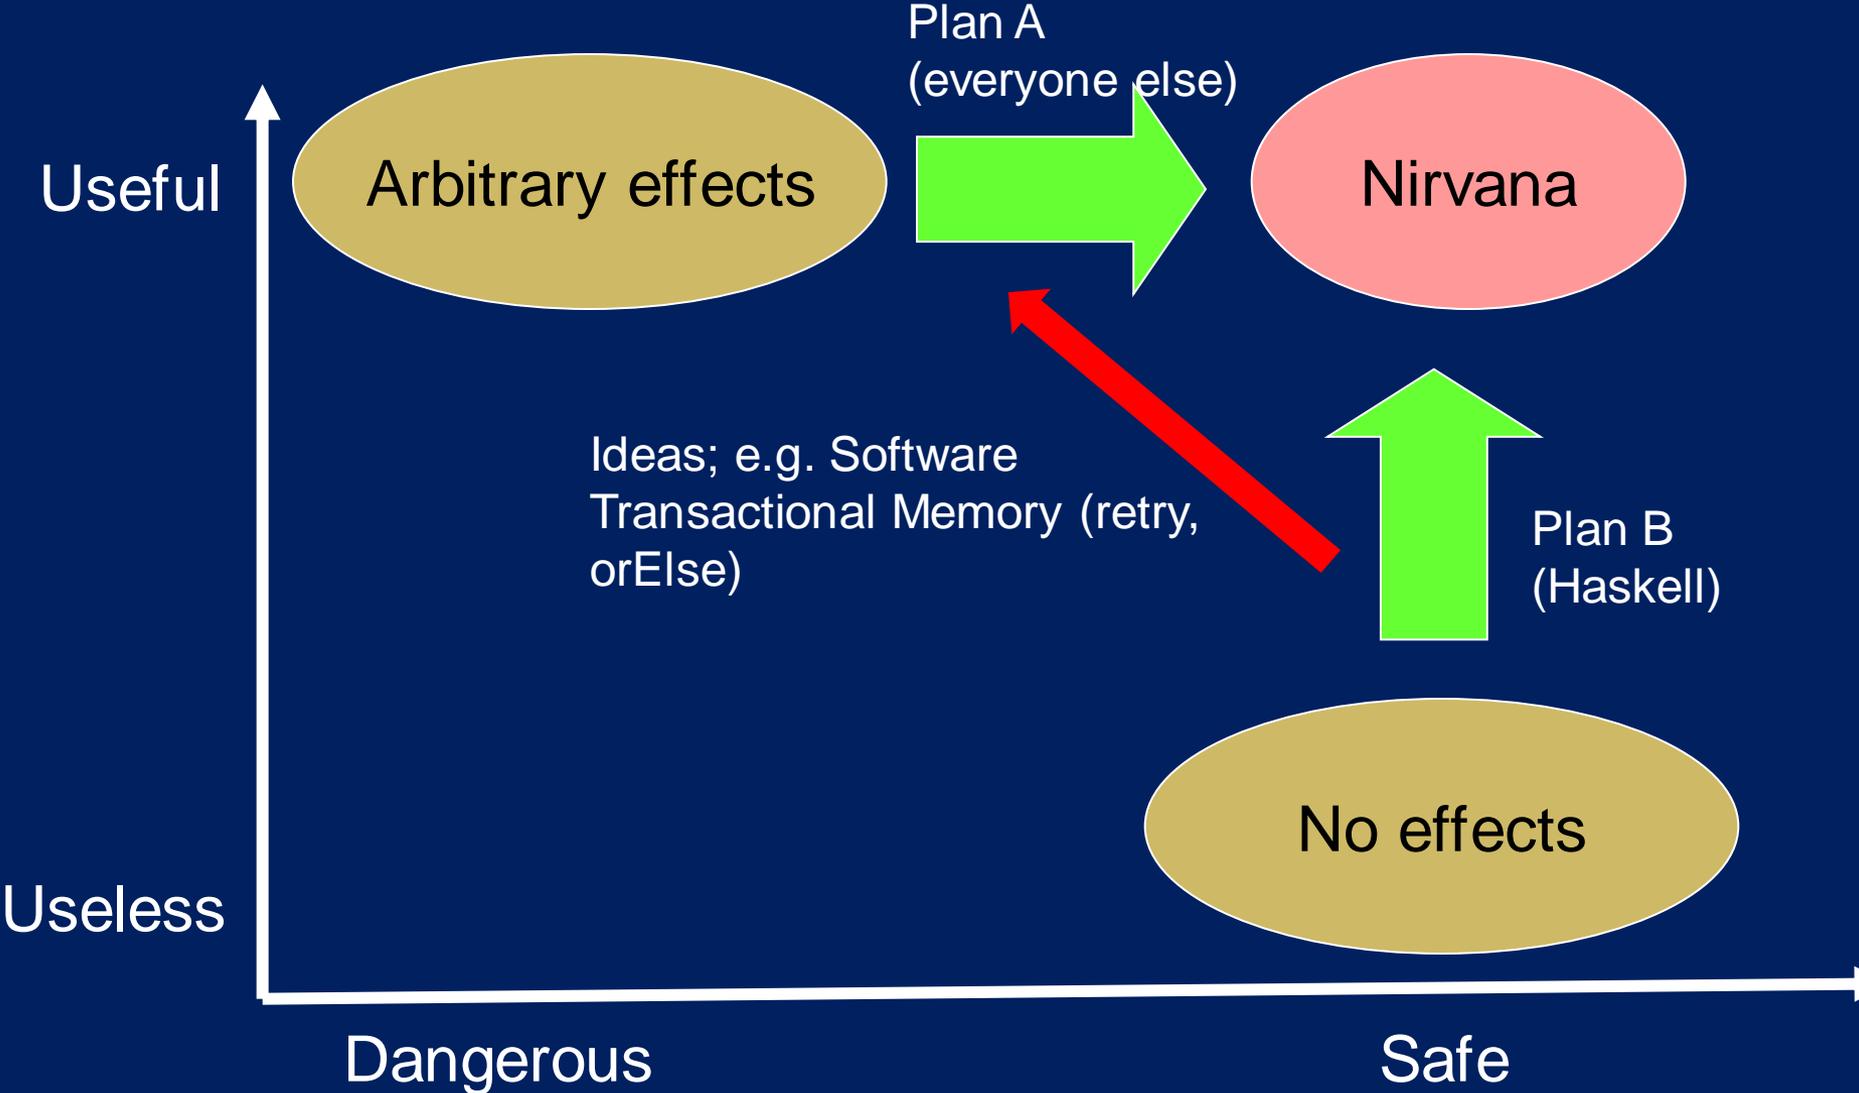


Value oriented programming

Lots of Cross Over



Lots of Cross Over



An Assessment and a Prediction

One of Haskell's most significant contributions is to take purity seriously, and relentlessly pursue Plan B.

Imperative languages will embody growing (and checkable) pure subsets.

-- Simon Peyton Jones