

IT days

#NewPerspectives

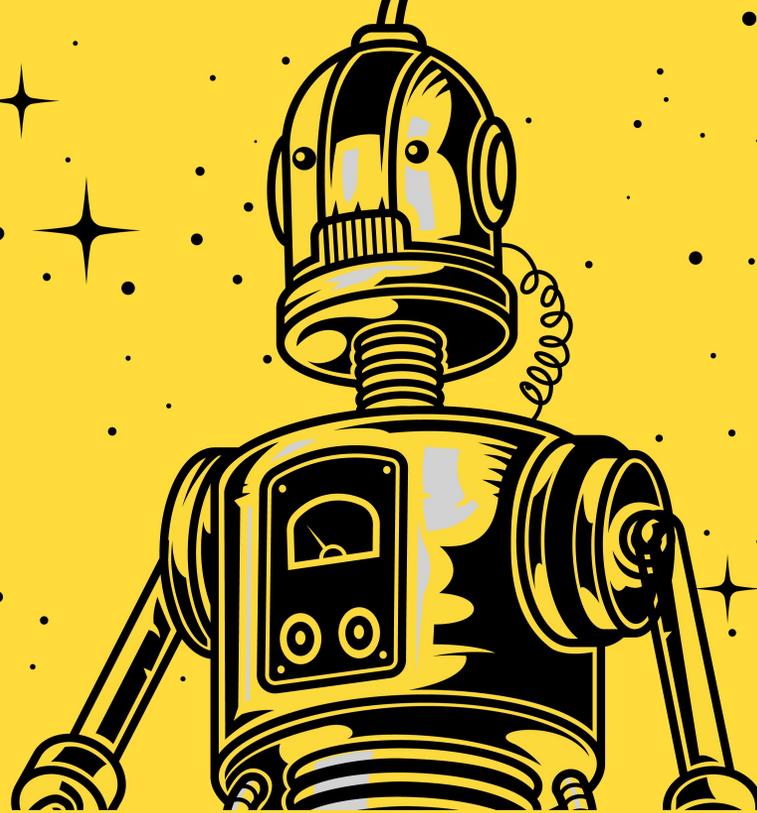


New Perspectives on Solving Concurrency

Lucian Radu Teodorescu

Staff Software Engineer @ Garmin

@LucT3o



New Perspectives on Solving Concurrency

lucteo.ro/pres/2021-itdays/



New Perspectives on Solving Concurrency

my perspectives

SOFTWARE ENGINEERING

Report on a conference sponsored by the
NATO SCIENCE COMMITTEE
Garmisch, Germany, 7th to 11th October 1968

Chairman: Professor Dr. F. L. Bauer
Co-chairmen: Professor L. Bollet, Dr. H. J. Helms

Editors: Peter Naur and Brian Randell

January 1969

software engineering

1968 NATO conference



concurrency in SE

Dijkstra, 1965

Solution of a Problem in Concurrent Programming Control

E. W. DIJKSTRA
Technological University, Eindhoven, The Netherlands

A number of mainly independent sequential-cyclic processes with restricted means of communication with each other can be made in such a way that at any moment one and only one of them is engaged in the "critical section" of its cycle.

Introduction

Given in this paper is a solution to a problem for which, to the knowledge of the author, has been an open question since at least 1962, irrespective of the solvability. The paper consists of three parts: the problem, the solution, and the proof. Although the setting of the problem might seem somewhat academic at first, the author trusts that anyone familiar with the logical problems that arise in computer coupling will appreciate the significance of the fact that this problem indeed can be solved.

The Problem

To begin, consider N computers, each engaged in a process which, for our aims, can be regarded as cyclic. In each of the cycles a so-called "critical section" occurs and the computers have to be programmed in such a way that at any moment only one of these N cyclic processes is in its critical section. In order to effectuate this mutual exclusion of critical-section execution the computers can communicate with each other via a common store. Writing a word into or nondestructively reading a word from this store are unidivisible operations; i.e., when two or more computers try to communicate (either for reading or for writing) simultaneously with the same common location, these communications will take place one after the other, but in an unknown order.

The solution must satisfy the following requirements.

(a) The solution must be symmetrical between the N computers; as a result we are not allowed to introduce a static priority.

(b) Nothing may be assumed about the relative speeds of the N computers; we may not even assume their speeds to be constant in time.

(c) If any of the computers is stopped well outside its critical section, this is not allowed to lead to potential blocking of the others.

(d) If more than one computer is about to enter its critical section, it must be impossible to devise for them such finite speeds, that the decision to determine which one of them will enter its critical section first is postponed until eternity. In other words, constructions in which "After you"-blocking is still possible, although improbable, are not to be regarded as valid solutions.

We beg the challenged reader to stop here for a while and have a try himself, for this seems the only way to get a feeling for the tricky consequences of the fact that each

computer can only request one one-way message at a time. And only this will make the reader realize to what extent this problem is far from trivial.

The Solution

The common store consists of:

"Boolean array b , $c[1:N]$; integer k "

The integer k will satisfy $1 \leq k \leq N$, $b[i]$ and $c[i]$ will only be set by the i th computer; they will be inspected by the others. It is assumed that all computers are started well outside their critical sections with all Boolean arrays mentioned set to **true**; the starting value of k is immaterial.

The program for the i th computer ($1 \leq i \leq N$) is:

```

"integer  $j$ ;
L10:  $b[i] := \text{false}$ ;
L11: if  $k \neq i$  then
L12: begin  $c[i] := \text{true}$ ;
L13: if  $b[k]$  then  $k := i$ ;
      go to L11
    end
  else
L14: begin  $c[i] := \text{false}$ ;
      for  $j := 1$  step 1 until  $N$  do
        if  $j \neq i$  and not  $c[j]$  then go to L11
      end;
    critical section;
     $c[i] := \text{true}$ ;  $b[i] := \text{true}$ ;
    remainder of the cycle in which stopping is allowed;
    go to L10'

```

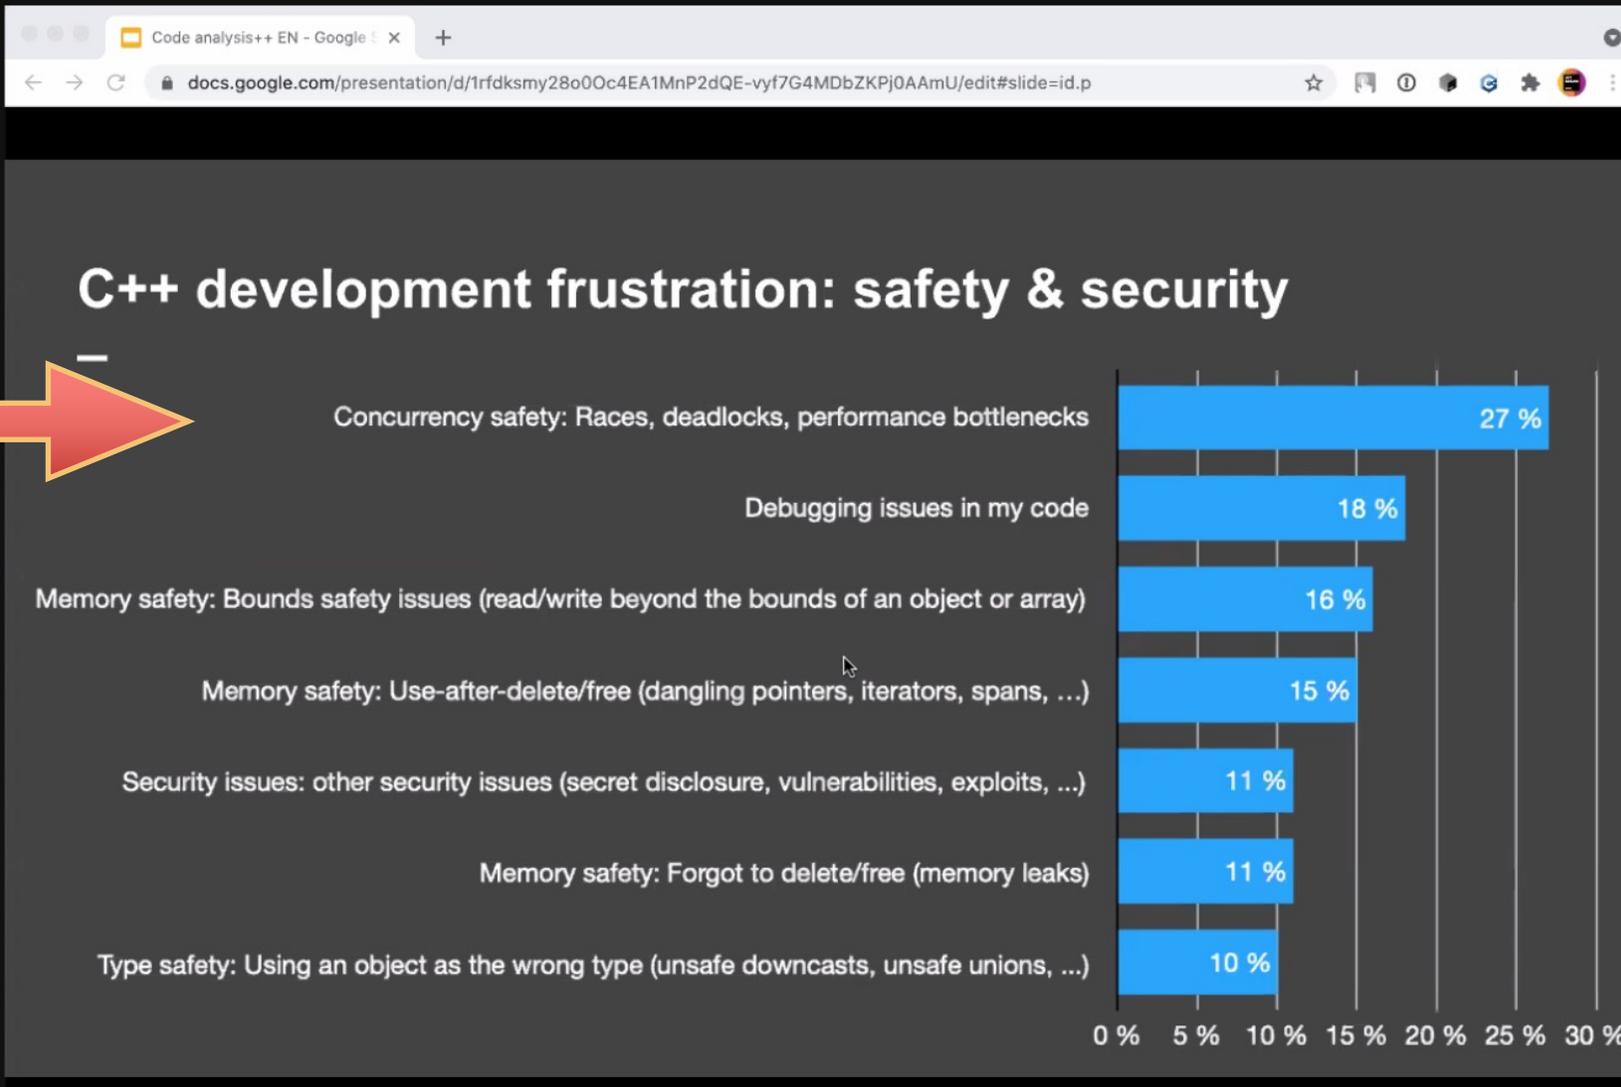
The Proof

We start by observing that the solution is safe in the sense that no two computers can be in their critical section simultaneously. For the only way to enter its critical section is the performance of the compound statement L14 without jumping back to L11, i.e., finding all other c 's **true** after having set its own c to **false**.

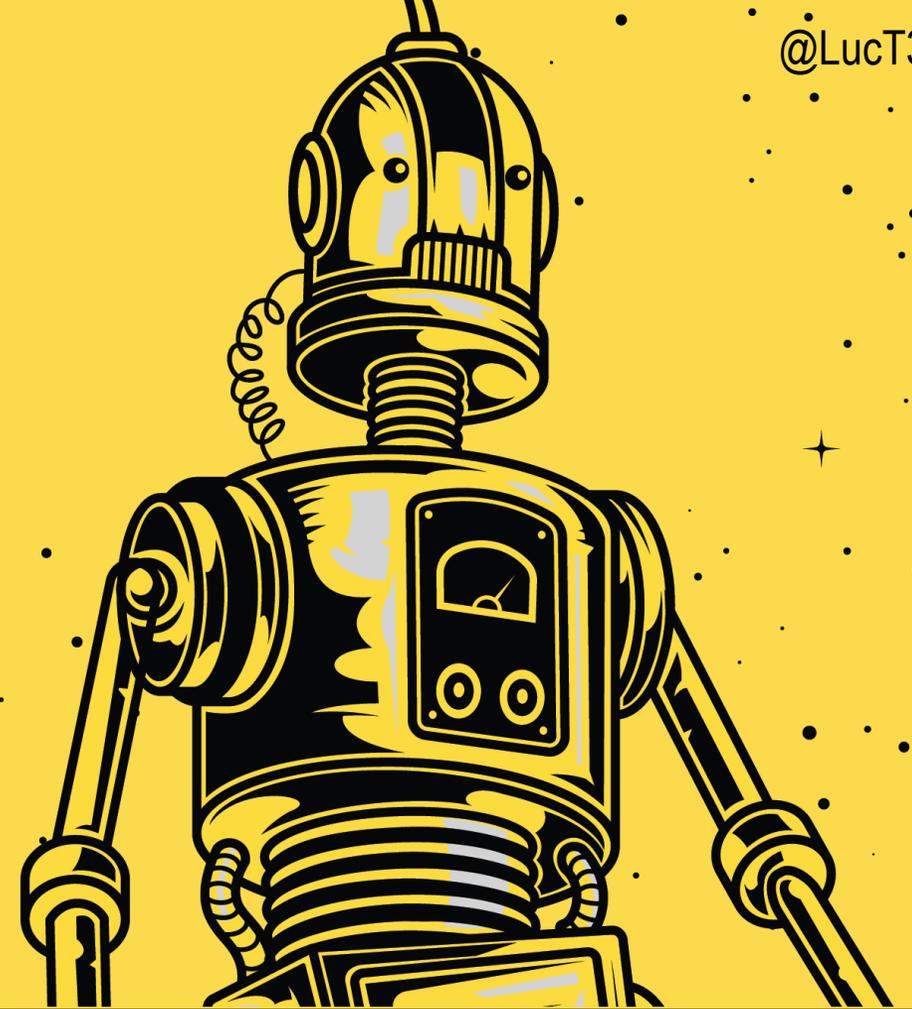
The second part of the proof must show that no infinite "After you"-blocking can occur; i.e., when none of the computers is in its critical section, of the computers looping (i.e., jumping back to L11) at least one—and therefore exactly one—will be allowed to enter its critical section in due time.

If the k th computer is not among the looping ones, $b[k]$ will be **true** and the looping ones will all find $k \neq i$. As a result one or more of them will find in L13 the Boolean $b[k]$ **true** and therefore one or more will decide to assign " $k := i$ ". After the first assignment " $k := i$ ", $b[k]$ becomes **false** and no new computers can decide again to assign a new value to k . When all decided assignments to k have been performed, k will point to one of the looping computers and will not change its value for the time being, i.e., until $b[k]$ becomes **true**, viz., until the k th computer has completed its critical section. As soon as the value of k does not change any more, the k th computer will wait (via the compound statement L14) until all other c 's are **true**, but this situation will certainly arise, if not already present, because all other looping ones are forced to set their c **true**, as they will find $k \neq i$. And this, the author believes, completes the proof.

is the problem **solved**?



Old perspective



building blocks

independent threads
synchronization primitives

roads analogy

thread \rightarrow road

sync primitive \rightarrow intersection

works well
for long roads
and few intersections





different reality

synchronization issues

deadlock

livelock

starvation

priority inversion

busy waiting

performance

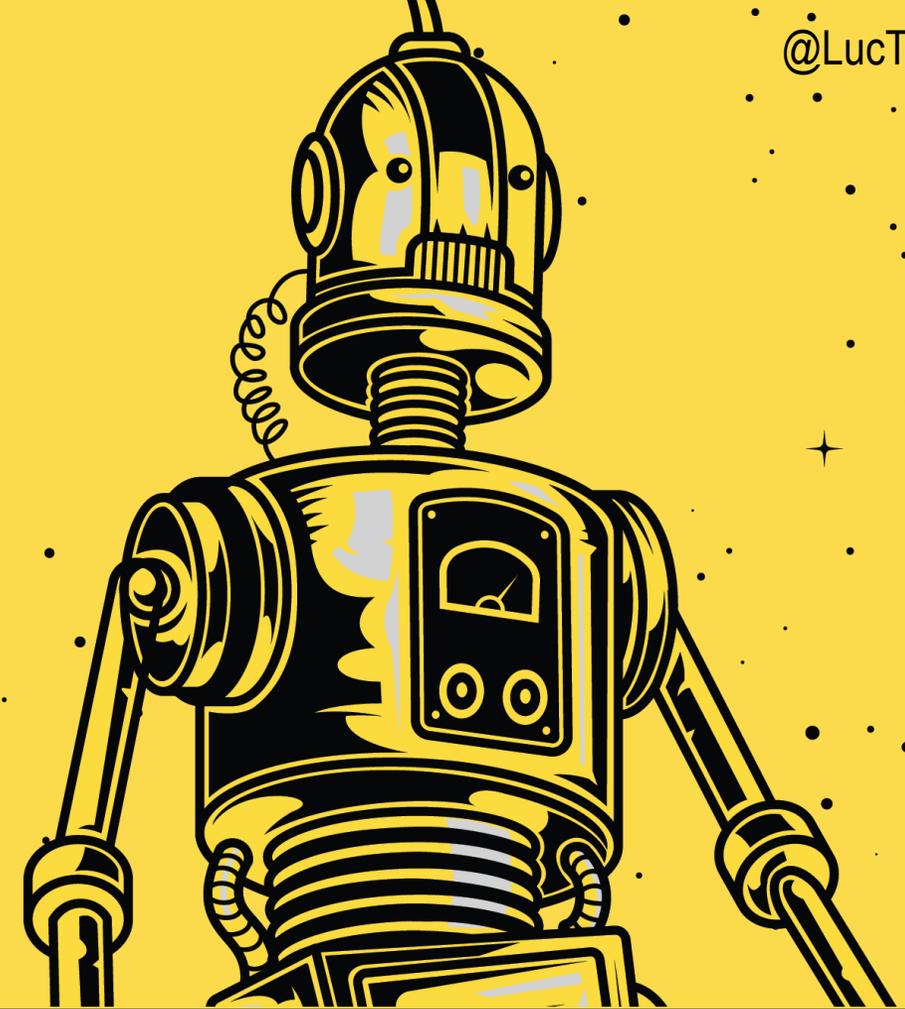
far from expected

threads

hard to think of them as independent

primitives are **not OK**

threads
synchronization



Concurrency with tasks

primitive

task = independent unit of work

primitive

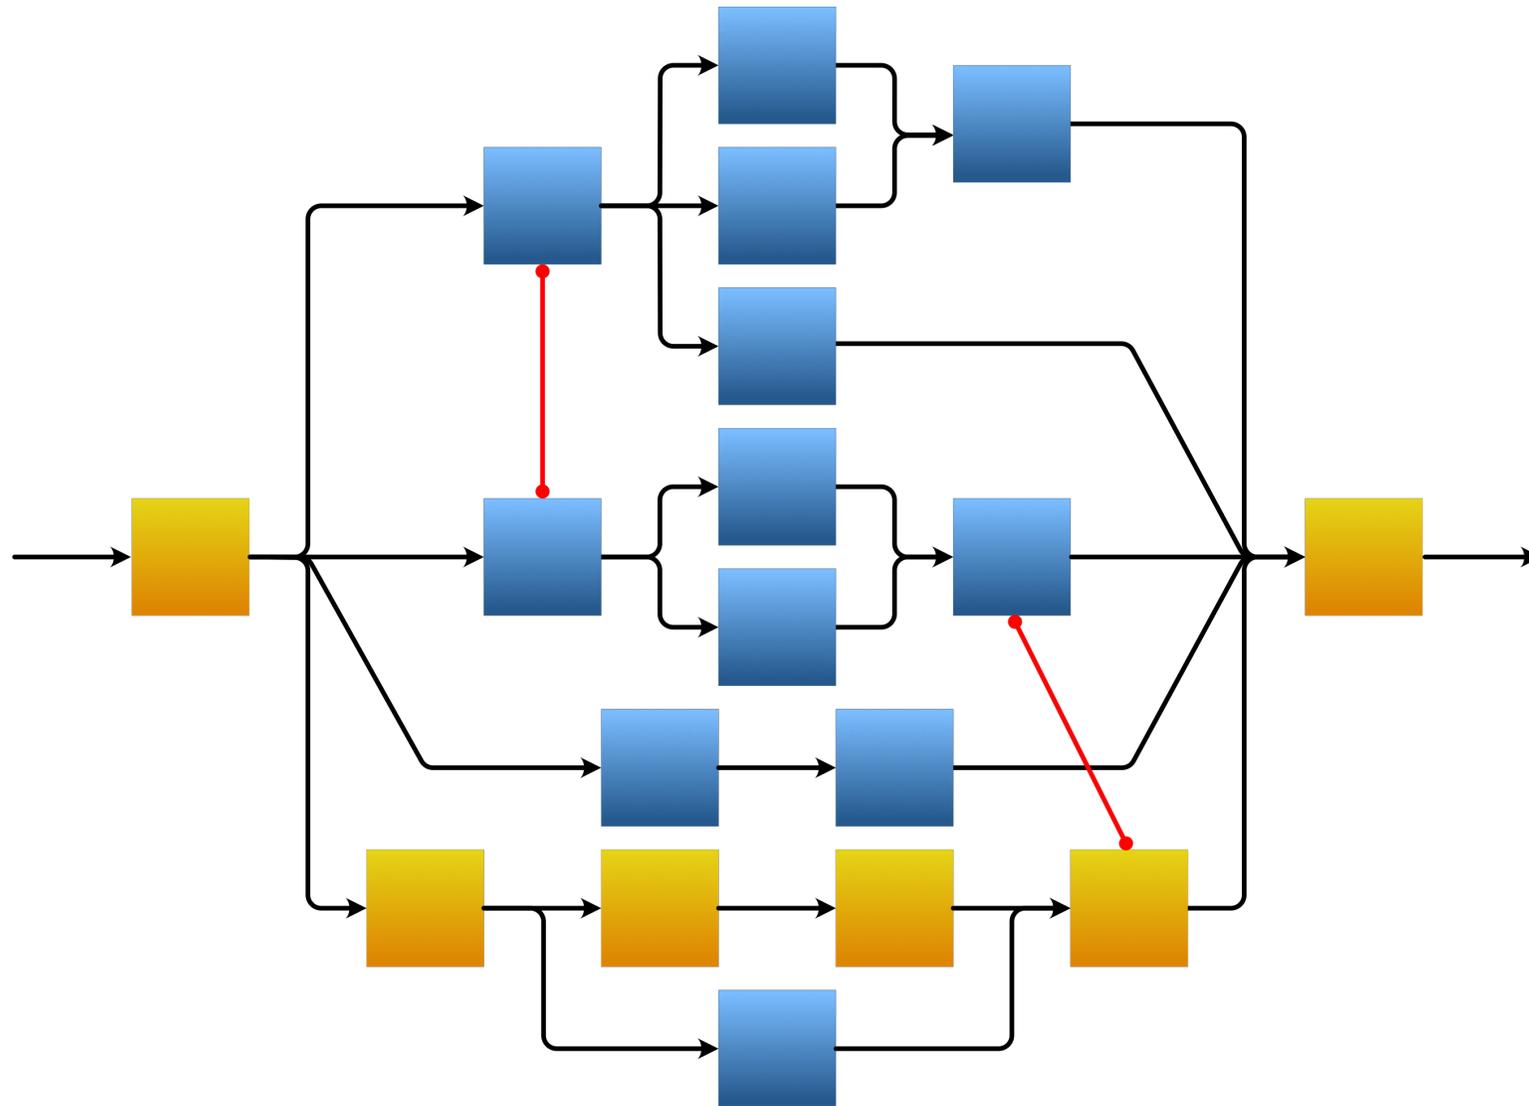
task = independent unit of work

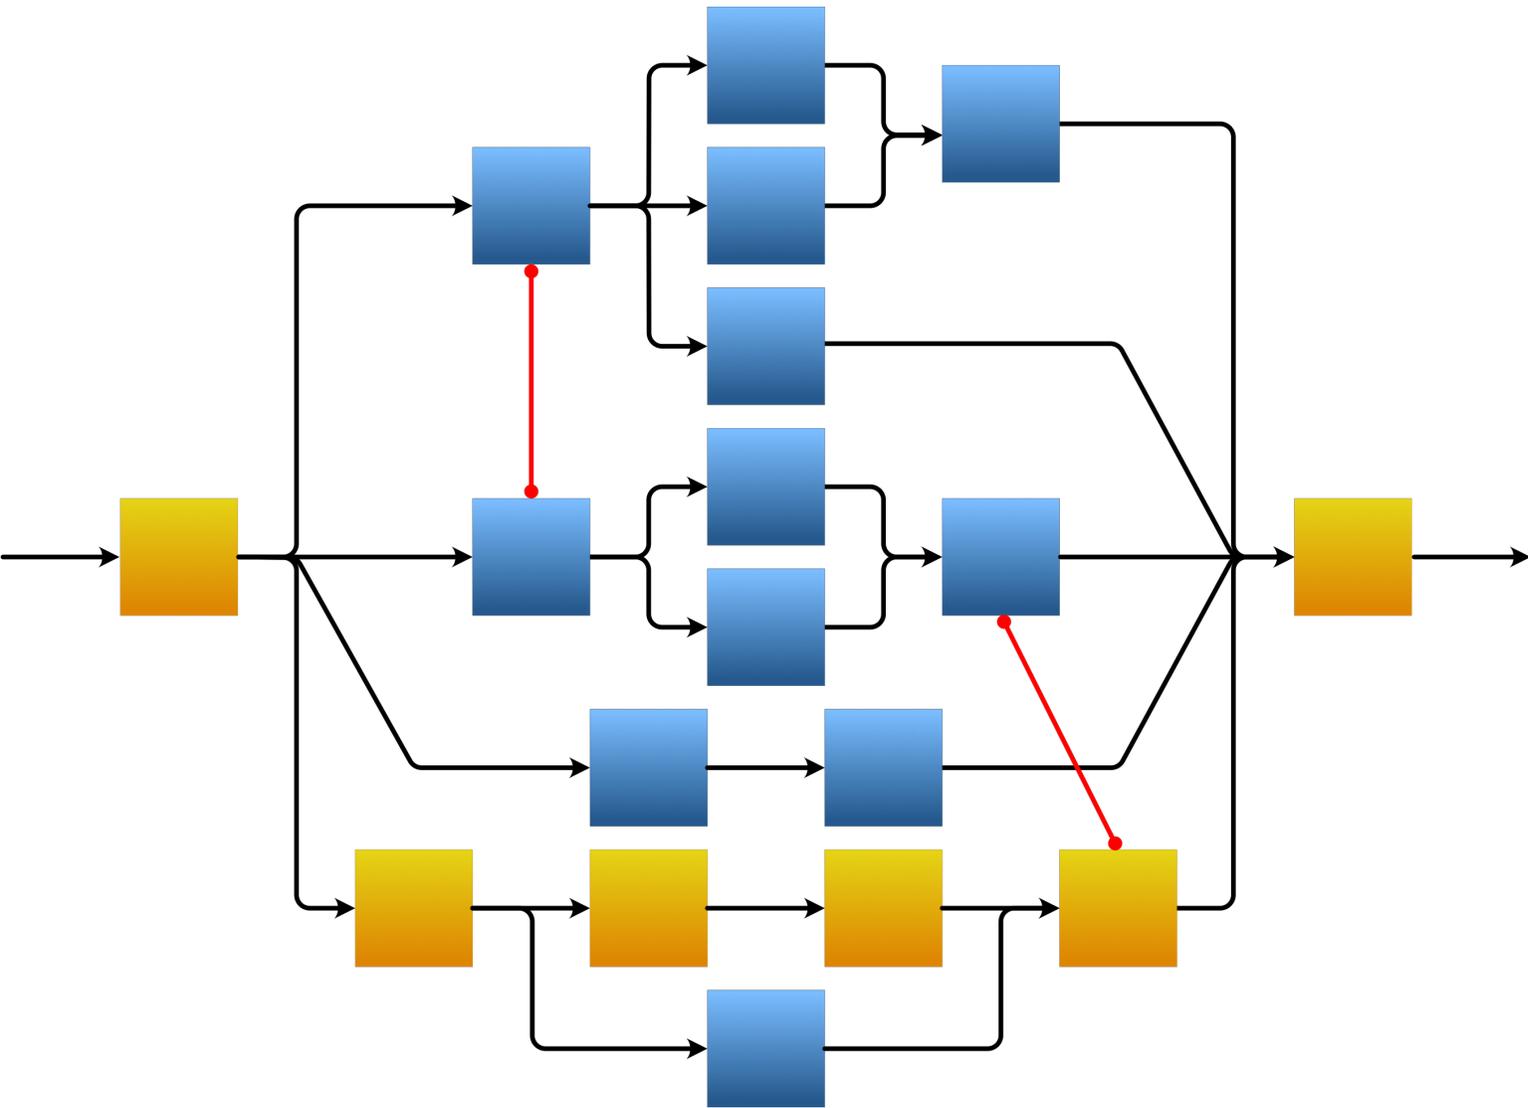
independent := does not depend on anything but its inputs

primitive

task = independent **unit of work**

unit := doesn't make sense to divide it





constraints
instead of locks

new problem

encoding concurrency with tasks

overload 157

JUNE 2020 £4.50

Refocusing Amdahl's Law

Parallelising code can make it faster. We explore how to get the most out of multi-threaded code.

Some Objects Are More Equal Than Others

We investigate different languages' approaches for object comparison.

Comment Only What the Code Cannot Say

A sensible approach to writing code comments

Afterword

Mantras are useful – but omitting vital information can lead to disaster



A magazine of ACCU

ISSN: 1354-3172

Refocusing Amdahl's Law

high efficiency for Greedy algo
high speedups

$$S_p \geq \frac{N}{K + \frac{N-K}{P}}$$

$$S_{1000} = 500.25$$

$$S_{10} = 9.91$$

$$N = 1000$$

$$K = 1$$

overload 158

AUGUST 2020 £4.50

The Global Lockdown of Locks

We demonstrate why you do not need mutexes in high-level code, since any concurrent algorithm can be implemented safely and efficiently with "tasks".

C++20: A Simple Math Module

An introduction to C++ 20 modules using a simple math library

A Thorough Introduction to Apache Kafka

An introduction to Kafka, which is the heart of many companies' architecture

An Example Confined User Shell

Employing snaps to provide bespoke confined Linux environments

The Global Lockdown of Locks

global solution
safety ensured
no need for locks

overload 159

OCTOBER 2020 £4.50

Concurrency Design Patterns

Orchestrating concurrent tasks using mutexes is seldom efficient. We investigate design patterns that help unlock concurrent performance.

C++ Modules: A Brief Tour

A tourist's guide to C++20's long awaited module system

Kafka Acks Explained

Visualizing Kafka's most misunderstood configuration setting

The Edge of C++

Every technology has a boundary; we look at the "outer limits" of C++

poly::vector - A Vector for Polymorphic Objects

An efficient C++ container of polymorphic objects, based on STL principles

A magazine of ACCU

ISSN: 1354-3172

Concurrency Design Patterns

building blocks for concurrent applications

**ACC4
2021**
VIRTUAL EVENT

Bloomberg
Engineering

undo

mosaic
CONSULTANTS TO FINANCIAL SERVICES

Threads Considered Harmful

Lucian Radu Teodorescu

https://youtu.be/_T1XjxXNSCs



Designing Concurrent C++ Applications

Lucian Radu Teodorescu

C++ now

20
MA
Aspen, Colo

<https://bit.ly/2YnVG5U>



tasks

a new **solution** for concurrency

overload 162

APRIL 2021 £4.50



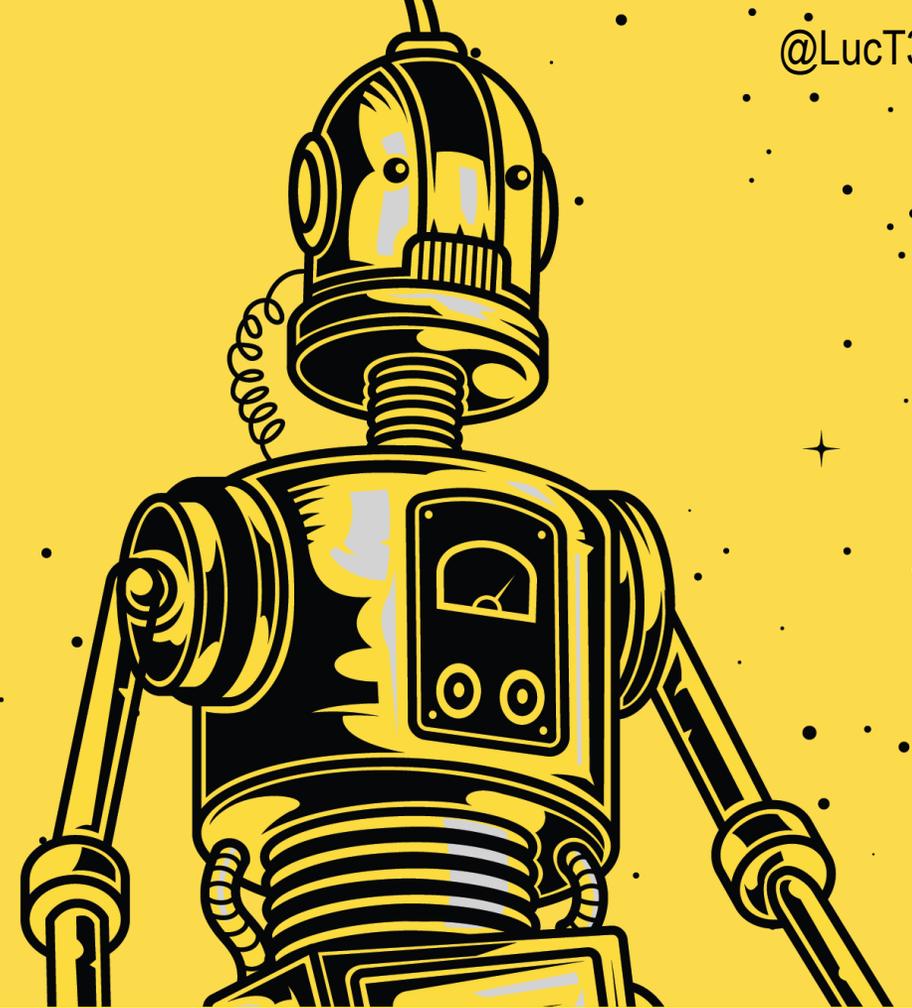
Composition and Decomposition of Task Systems

top-down and bottom-up design

composition of tasks

not the best solution

Async computations



C++ executors

P0443: A Unified Executors Proposal for C++

<https://wg21.link/p0443r14>

P2300: `std::execution`

<https://wg21.link/p2300r2>

overload 165

OCTOBER 2021 £4.50

Showing Variables Using the Windows Debugging API

Exploring the deep magic of debuggers.



Stufftar Revisited

Personal projects can provide valuable learning opportunities

Executors: a Change of Perspective

Exploring the new C++ proposal

Afterwood

Reflecting on reflection

A magazine of ACCU

ISSN: 1354-3172

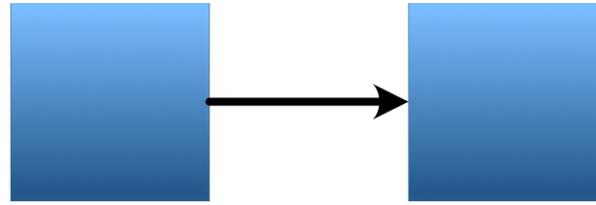
Executors: a Change of Perspective

senders/receivers

are a **better**

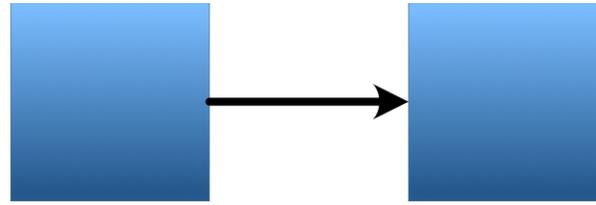
concurrency abstraction

tasks



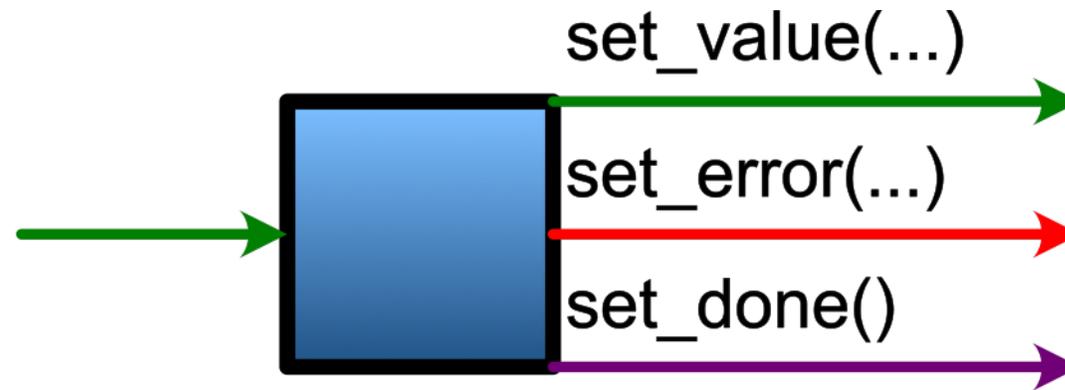
cannot directly pass values between tasks

tasks



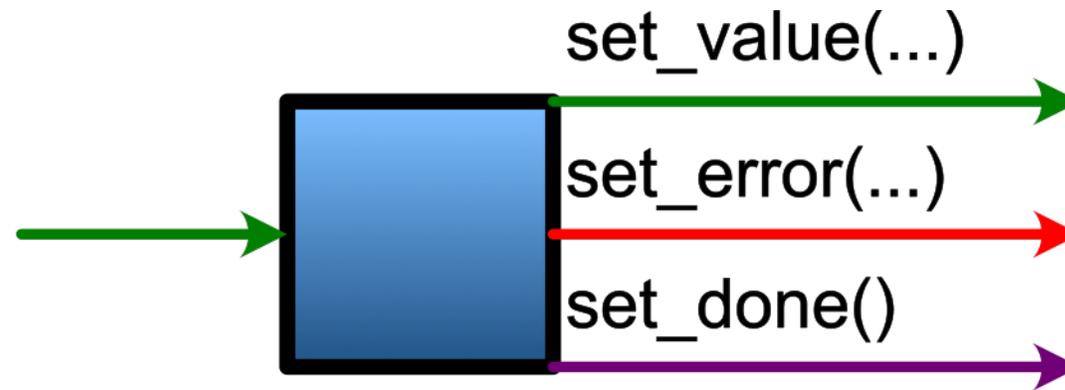
task body contains the call to the next task
what happens in case of error?

senders



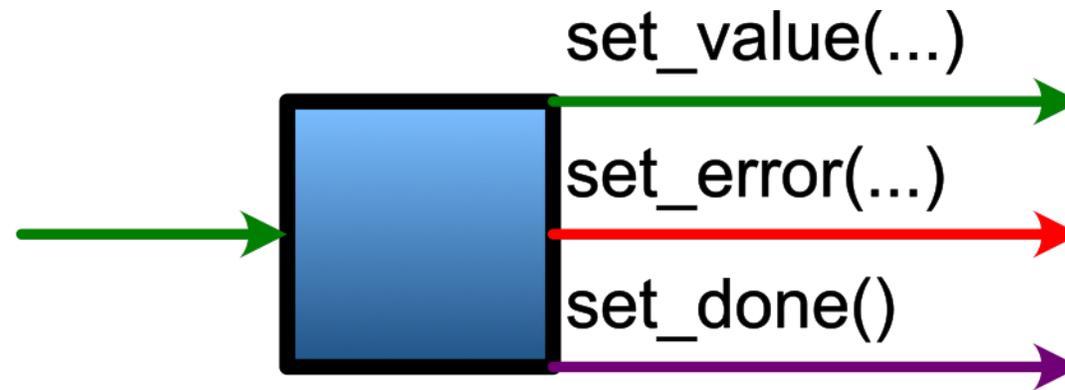
3 notification channels
wiring done by the framework

senders



no performance penalty

asynchronous computations



rename **senders** into “**async computations**”

what can be a **computation**?

a small chunk of work

a task

a group of tasks

a group of task groups

the entire application

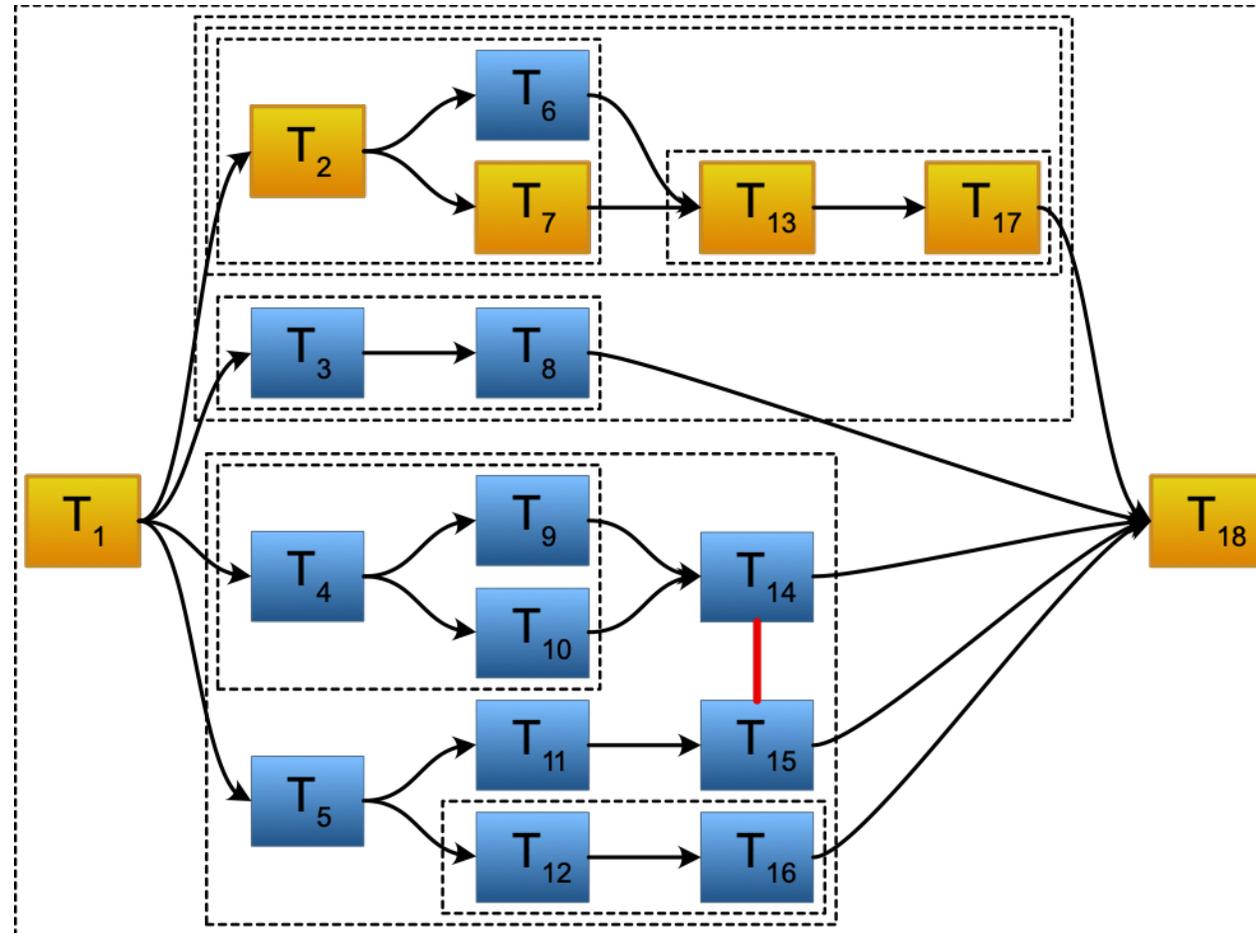
computations

general solution to concurrency

computations

compose better than tasks

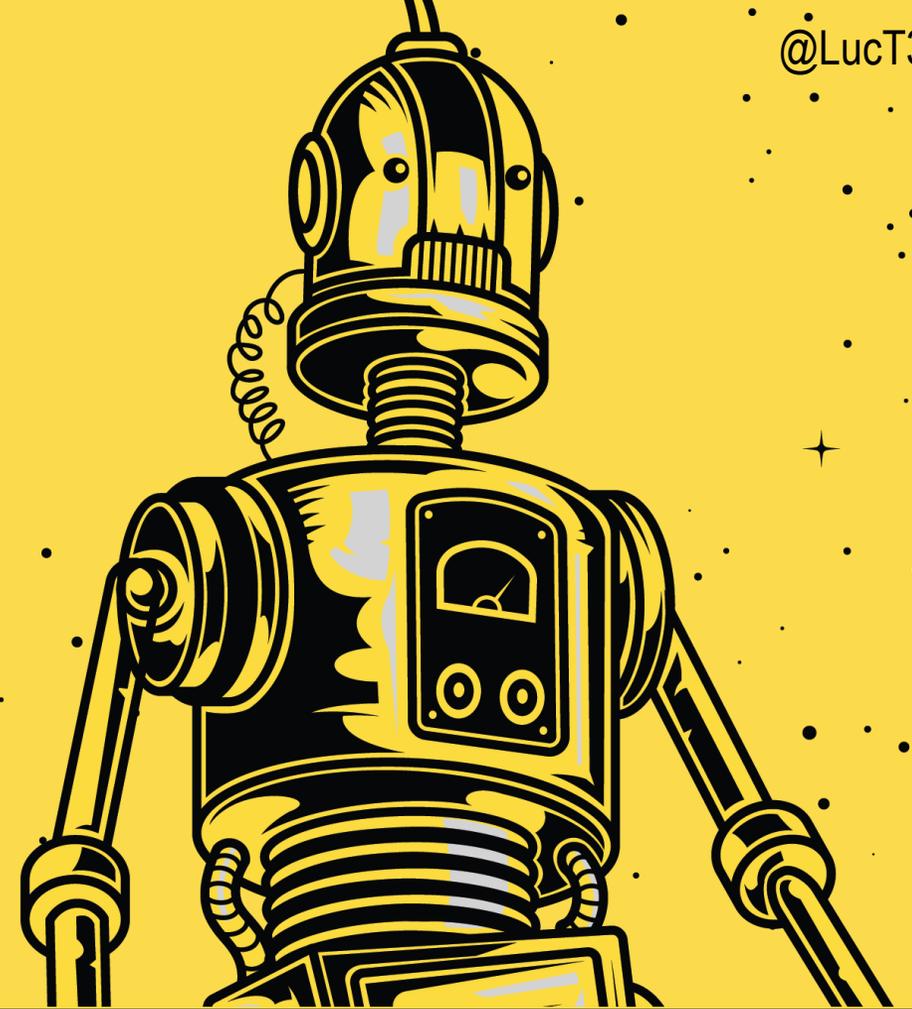
computations hierarchy



computation is an **abstraction**

allows us to incorporate concurrency in design

The future



goals

no more thread safety issues
clean design for concurrency

change of primitives

threads & locks



computations

change of approach

synchronization



constraints between computations

patterns & examples

make it teachable

widespread use

concurrency \neq frustration

IT days

#NewPerspectives



Q & A

