

Edsger W. Dijkstra

**Solution of a Problem  
in Concurrent Programming Control**

*Communications of the ACM,*  
*Vol. 8 (9), 1965*  
*pp. 569*

# 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 undividable 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"- "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  $j$ ;
Li0:  $b[i] := \text{false}$ ;
Li1: if  $k \neq i$  then
Li2: begin  $c[i] := \text{true}$ ;
Li3: if  $b[k]$  then  $k := i$ ;
      go to Li1
      end
      else

```

```

Li4: begin  $c[i] := \text{false}$ ;
      for  $j := 1$  step 1 until  $N$  do
        if  $j \neq i$  and not  $c[j]$  then go to Li1
      end;
      critical section;
       $c[i] := \text{true}$ ;  $b[i] := \text{true}$ ;
      remainder of the cycle in which stopping is allowed;
      go to Li0

```

### 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 *Li4* without jumping back to *Li1*, 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"-"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 *Li1*) 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 *Li3* 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 *Li4*) 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.

*Letters to the Editor*

Edsger W. Dijkstra

**Go To Statement Considered Harmful**

*Communications of the ACM,*  
*Vol. 11 (2), 1968*  
*pp. 147-148*

## Go To Statement Considered Harmful

**Key Words and Phrases:** go to statement, jump instruction, branch instruction, conditional clause, alternative clause, repetitive clause, program intelligibility, program sequencing

**CR Categories:** 4.22, 5.23, 5.24

### EDITOR:

For a number of years I have been familiar with the observation that the quality of programmers is a decreasing function of the density of **go to** statements in the programs they produce. More recently I discovered why the use of the **go to** statement has such disastrous effects, and I became convinced that the **go to** statement should be abolished from all “higher level” programming languages (i.e. everything except, perhaps, plain machine code). At that time I did not attach too much importance to this discovery; I now submit my considerations for publication because in very recent discussions in which the subject turned up, I have been urged to do so.

My first remark is that, although the programmer’s activity ends when he has constructed a correct program, the process taking place under control of his program is the true subject matter of his activity, for it is this process that has to accomplish the desired effect; it is this process that in its dynamic behavior has to satisfy the desired specifications. Yet, once the program has been made, the “making” of the corresponding process is delegated to the machine.

My second remark is that our intellectual powers are rather geared to master static relations and that our powers to visualize processes evolving in time are relatively poorly developed. For that reason we should do (as wise programmers aware of our limitations) our utmost to shorten the conceptual gap between the static program and the dynamic process, to make the correspondence between the program (spread out in text space) and the process (spread out in time) as trivial as possible.

Let us now consider how we can characterize the progress of a process. (You may think about this question in a very concrete manner: suppose that a process, considered as a time succession of actions, is stopped after an arbitrary action, what data do we have to fix in order that we can redo the process until the very same point?) If the program text is a pure concatenation of, say, assignment statements (for the purpose of this discussion regarded as the descriptions of single actions) it is sufficient to point in the

program text to a point between two successive action descriptions. (In the absence of **go to** statements I can permit myself the syntactic ambiguity in the last three words of the previous sentence: if we parse them as “successive (action descriptions)” we mean successive in text space; if we parse as “(successive action) descriptions” we mean successive in time.) Let us call such a pointer to a suitable place in the text a “textual index.”

When we include conditional clauses (**if  $B$  then  $A$** ), alternative clauses (**if  $B$  then  $A_1$  else  $A_2$** ), choice clauses as introduced by C. A. R. Hoare (case[i] of ( $A_1, A_2, \dots, A_n$ )), or conditional expressions as introduced by J. McCarthy ( $B_1 \rightarrow E_1, B_2 \rightarrow E_2, \dots, B_n \rightarrow E_n$ ), the fact remains that the progress of the process remains characterized by a single textual index.

As soon as we include in our language procedures we must admit that a single textual index is no longer sufficient. In the case that a textual index points to the interior of a procedure body the dynamic progress is only characterized when we also give to which call of the procedure we refer. With the inclusion of procedures we can characterize the progress of the process via a sequence of textual indices, the length of this sequence being equal to the dynamic depth of procedure calling.

Let us now consider repetition clauses (like, **while  $B$  repeat  $A$  or repeat  $A$  until  $B$** ). Logically speaking, such clauses are now superfluous, because we can express repetition with the aid of recursive procedures. For reasons of realism I don't wish to exclude them: on the one hand, repetition clauses can be implemented quite comfortably with present day finite equipment; on the other hand, the reasoning pattern known as “induction” makes us well equipped to retain our intellectual grasp on the processes generated by repetition clauses. With the inclusion of the repetition clauses textual indices are no longer sufficient to describe the dynamic progress of the process. With each entry into a repetition clause, however, we can associate a so-called “dynamic index,” inexorably counting the ordinal number of the corresponding current repetition. As repetition clauses (just as procedure calls) may be applied nestedly, we find that now the progress of the process can always be uniquely characterized by a (mixed) sequence of textual and/or dynamic indices.

The main point is that the values of these indices are outside programmer's control; they are generated (either by the write-up of his program or by the dynamic evolution of the process) whether he wishes or not. They provide independent coordinates in which to describe the progress of the process.

Why do we need such independent coordinates? The reason is—and this seems to be inherent to sequential processes—that we can interpret the value of a variable only with respect to the

progress of the process. If we wish to count the number,  $n$  say, of people in an initially empty room, we can achieve this by increasing  $n$  by one whenever we see someone entering the room. In the in-between moment that we have observed someone entering the room but have not yet performed the subsequent increase of  $n$ , its value equals the number of people in the room minus one!

The unbridled use of the **go to** statement has an immediate consequence that it becomes terribly hard to find a meaningful set of coordinates in which to describe the process progress. Usually, people take into account as well the values of some well chosen variables, but this is out of the question because it is relative to the progress that the meaning of these values is to be understood! With the **go to** statement one can, of course, still describe the progress uniquely by a counter counting the number of actions performed since program start (viz. a kind of normalized clock). The difficulty is that such a coordinate, although unique, is utterly unhelpful. In such a coordinate system it becomes an extremely complicated affair to define all those points of progress where, say,  $n$  equals the number of persons in the room minus one!

The **go to** statement as it stands is just too primitive; it is too much an invitation to make a mess of one's program. One can regard and appreciate the clauses considered as bridling its use. I do not claim that the clauses mentioned are exhaustive in the sense that they will satisfy all needs, but whatever clauses are suggested (e.g. abortion clauses) they should satisfy the requirement that a programmer independent coordinate system can be maintained to describe the process in a helpful and manageable way.

It is hard to end this with a fair acknowledgment. Am I to judge by whom my thinking has been influenced? It is fairly obvious that I am not uninfluenced by Peter Landin and Christopher Strachey. Finally I should like to record (as I remember it quite distinctly) how Heinz Zemanek at the pre-ALGOL meeting in early 1959 in Copenhagen quite explicitly expressed his doubts whether the **go to** statement should be treated on equal syntactic footing with the assignment statement. To a modest extent I blame myself for not having then drawn the consequences of his remark.

The remark about the undesirability of the **go to** statement is far from new. I remember having read the explicit recommendation to restrict the use of the **go to** statement to alarm exits, but I have not been able to trace it; presumably, it has been made by C. A. R. Hoare. In [1, Sec. 3.2.1.] Wirth and Hoare together make a remark in the same direction in motivating the case construction: "Like the conditional, it mirrors the dynamic structure of a program more clearly than **go to** statements and switches, and it eliminates the need for introducing a large number of labels in the program."



In [2] Guisepe Jacopini seems to have proved the (logical) superfluosness of the **go to** statement. The exercise to translate an arbitrary flow diagram more or less mechanically into a jumpless one, however, is not to be recommended. Then the resulting flow diagram cannot be expected to be more transparent than the original one.

REFERENCES:

1. WIRTH, NIKLAUS, AND HOARE, C. A. R. A contribution to the development of ALGOL. *Comm. ACM* 9 (June 1966), 413-432.
2. BÖHM, CORRADO, AND JACOPINI, GUISEPPE. Flow diagrams, Turing machines and languages with only two formation rules. *Comm. ACM* 9 (May 1966), 366-371.

EDSGER W. DIJKSTRA  
*Technological University  
Eindhoven, The Netherlands*