Fundamenta Informaticae 80 (2007) 1–9
IOS Press

The Decent Philosophers: An exercise in concurrent behaviour

Wolfgang Reisig*
Humboldt-Universität zu Berlin
Institute of Informatics
Unter den Linden 6, 10099 Berlin, Germany
reisig@informatik.hu-berlin.de

Abstract. Concurrent runs reveal more insight into distributed systems than interleaved runs. This is shown by help of Dijkstra's paradigm of five philosophers.

1. The problem

Antoni Mazurkiewicz' seminal work on traces boosted the acceptance of Carl Adam Petri's idea to conceive a single run (behaviour) of a distributed system not as a sequence of global states and steps, totally ordered in time, but as a set of local state and action occurrences, partially ordered by the cause - effect relation. This paper is an exercise in this kind of runs, called *concurrent runs* in the sequel. This exercise is based on Dijkstra's paradigm of *dining philosophers* as first introduced in [1]. Dijkstra never appreciated concurrent runs nor Petri Nets. But his paradigm nevertheless has interesting aspects to be nicely discussed in the framework of distributed runs and Petri Nets.

As the reader probably remembers, the dining philosophers paradigm is about processes that share resources, with "philosophers" and "forks" to represent processes and resources, respectively. Dijkstra describes the philosophers system in [1] as follows:

"Five philosophers, numbered from 0 to 4, are living in a house where the table is laid for them, each philosopher having his own place at the table. Their only problem – besides those of philosophy – is that the dish served is a very difficult kind of spaghetti, that has to be eaten with two forks. There are two forks next to each plate, so that presents no difficulty, as consequence, however, no neighbors may be eating simultaneously."

*Address for correspondence: Humboldt-Universität zu Berlin, Institute of Informatics, Unter den Linden 6, 10099 Berlin, Germany

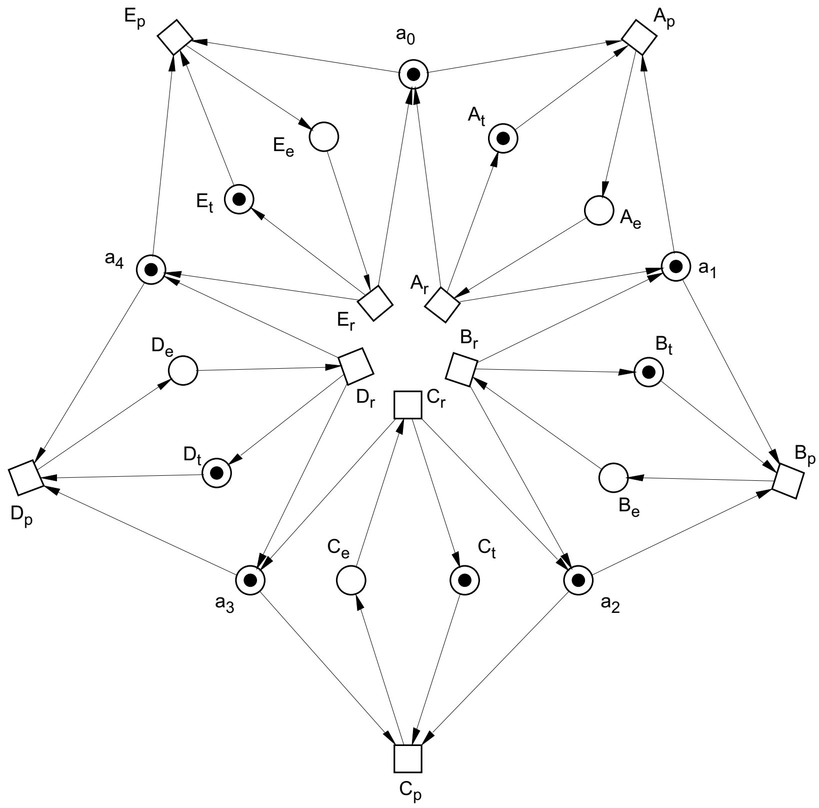

Figure 1. The five dining philosophers

Our first goal is a representation of the philosophers system as a Petri Net. The runs of the net are to describe the potential behaviors of the philosophers during their lifetime. Figure 1 shows this Petri Net. For convenience, the philosophers are denoted $A, \ldots, E$. Indices p, r, t, e stand for "picks up forks", "returns forks", "thinking" and "eating", respectively. For $i = 0, \ldots, 4$, condition a_i denotes that fork i is available for its users.

The initial marking M_0 enables each philosophers' "pick up forks" transition, in competition with his neighbor's "pick up forks" transitions. So, this Petri Net is nondeterministic, i.e. has many different (in fact, infinitely many different) runs.

A prefix of a typical, sequential, interleaved run of the philosophers system in Fig. 1 is

$$\begin{aligned} &M_0 \xrightarrow{A_p} M_1 \xrightarrow{C_p} M_2 \xrightarrow{C_r} M_3 \xrightarrow{A_r} M_4 \xrightarrow{B_p} M_5 \xrightarrow{D_p} M_6 \xrightarrow{B_r} \\ &M_7 \xrightarrow{B_p} M_8 \xrightarrow{D_r} M_9 \xrightarrow{E_p} M_{10} \xrightarrow{E_r} M_{11} \xrightarrow{B_r} M_{12} \end{aligned} \tag{1}$$

with markings $M_1, M_2 \ldots$ obvious from context. Philosopher B eats twice in this piece of a run and every other philosopher just once.

Turning now to concurrent runs, it is convenient to introduce a shorthand for pieces of runs. For philosopher A, call an occurrence of A_p, eventually followed by A_r, an *eating cycle* of A. As usual for concurrent runs, this eating cycle would be represented by an inscribed occurrence net, shown in Fig. 2.

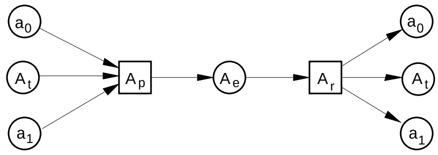

Figure 2. An eating cycle of A, represented as a concurrent run

As a shorthand we represent this eating cycle by

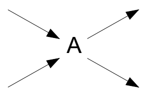

(2)

The upper ingoing and outgoing arrows represent the availability of fork 0. The lower arrows likewise represent fork 1. “Thinking” is not explicitly represented in (2).

Eating cycles of the other philosophers are analogously abbreviated. Composition of instances of eating cycles yield a concurrent run. Figure 3 shows an example. Its arrow sequences represent causal

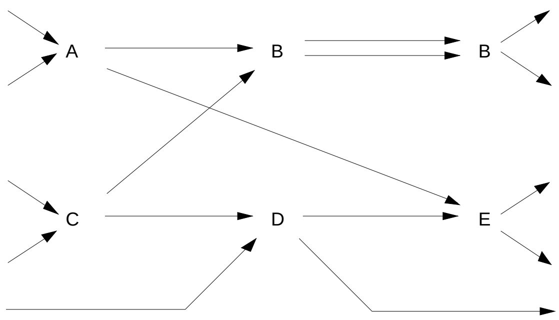

Figure 3. A concurrent run of Fig. 1

dependencies. For example, A occurs before B and E, and B occurs concurrently to D and E. B occurs twice, with no neighbors A or E in between. Any total extension of the partial order of Fig. 3 yields a sequential, interleaved run. One of them is (1). Obviously, the philosophers system of Fig. 1 has infinitely many, infinitely long concurrent runs.

A particularly *decent* kind of runs will be considered in the sequel:

A run of the Petri Net in Fig. 1 is *decent* iff
neighbored philosophers alternatingly use their shared fork. (3)

The sequential run of (1) as well as the concurrent run of Fig. 3 both are apparently not decent: B eats twice, without A and C eating in between. Figure 4 shows a decent run, R_1. For the sake of graphical clarity, each occurrence of B is represented twice. In the upper line as well as in the lower line. The reader must identify (“glue”) the corresponding occurrences as indicated.

We are now prepared to state the problem of this paper:

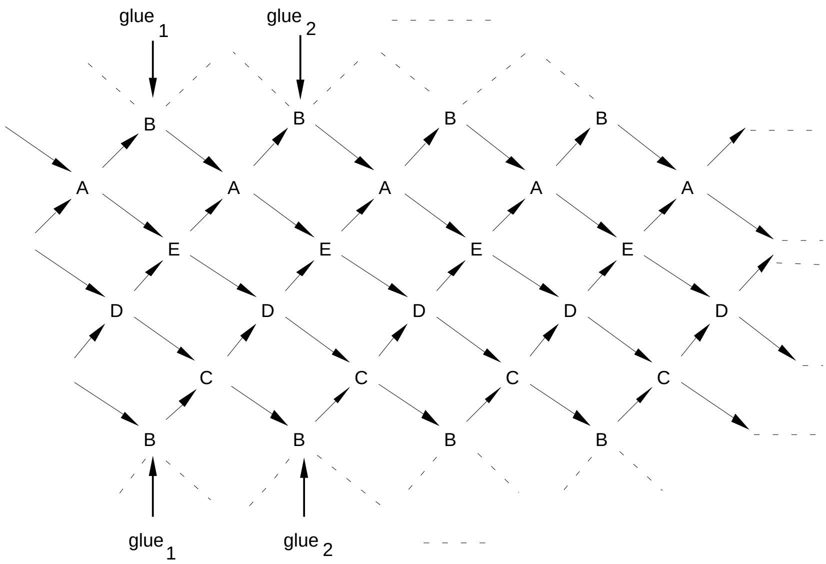

Figure 4. A decent concurrent run R_1 of Fig. 1

What can be said about the decent runs of the philosophers
system in Fig. 1? How many are there? (4)
How characterize the set of all decent runs?

2. The solution

To keep matters as simple as possible, we stick to *maximal* runs,i.e. runs that are not proper prefixes of other runs.

The reader may try to tackle the problem with interleaved, sequential runs. But he or she will almost certainly end up in a maze of technicalities, rarely revealing any insight into the problem. Much better is the approach based on concurrent runs. First of all, it is easy to realize that a decent concurrent run is fully determined by the forks' first use: After each fork has been used once, the rest of the run is absolutely fixed! Still, to each – infinite – concurrent run there exist infinitely many total extensions of its partial order. Hence, there are infinitely many decent interleaved runs, but apparently only finitely many decent concurrent runs. We try to characterize all of them.

Some decent runs can be derived form the run R_1 in Fig. 4, by skipping the first (i.e. the leftmost) occurrence of a subset M of eating cycles. In fact, M may be one of the sets $\{A\}$, $\{D\}$, $\{A, B\}$, $\{A, D\}$, $\{A, B, D\}$, $\{A, B, C, D\}$ or $\{A, B, D, E\}$.

Figure 5 shows the case of $M = \{A, B\}$. This way, R_1 of Fig. 4 yields 9 different decent runs. A further – and last – such run is gained by skipping the first *two* occurrences of A, as well as the first occurrences of B, D, and E (but no occurrence of C), as Fig. 6 shows.

There are more decent runs of Fig. 1. Though fairly similar to the decent run R_1 of Fig. 4, the run R_2 of Fig. 7 is substantially different, as the patterns of Fig. 8 show. The difference can be described by the relationship between occurrences of a philosopher's eating cycle and the concurrent occurrences of eating cycles of the two non-neighbored philosophers: They occur clockwise in R_1 and anticlockwise in

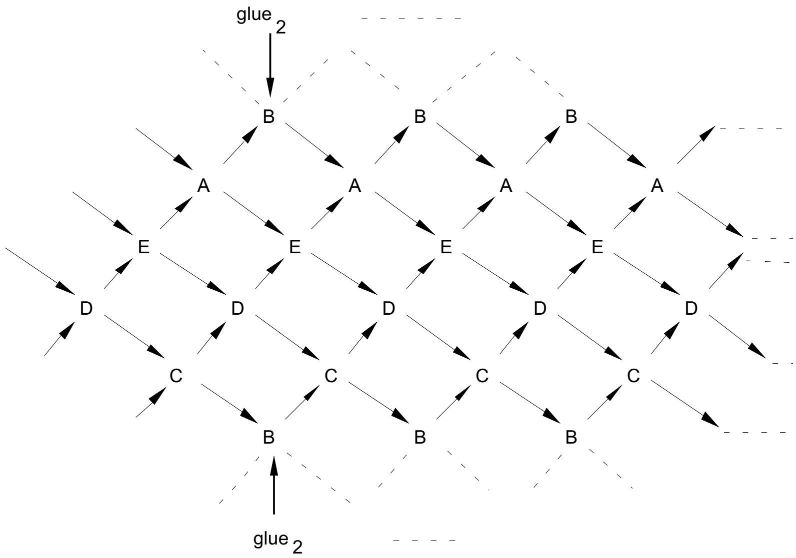

Figure 5. Concurrent run derived from Fig. 4 by skipping $\Sigma = \{A, B\}$.

R_2. As an example, in R_1, to each eating cycle of A there exist concurrent occurrences of cycles of C and D with C before D. Hence in R_1 the left pattern of Fig. 8 occurs and in R_2 the right pattern.

As in case of R_1, another 10 decent runs can be derived from R_2, by skipping a prefix of R_2.

Are there more than those 20 decent runs conceived so far? There are indeed. Due to an "unlucky" choice of the first users of the forks, there exist two further, less concurrent runs.

Figure 9 shows one of them, R_3. Another 4 decent runs can be derived from R_3, again by skipping subsets M of leftmost eating cycles, with $M = \{C\}$, $\{C, D\}$, $\{C, D, E\}$ and $\{C, D, E, A\}$. R_3 has a counterpart R_4 (the construction of which we leave to the reader) with another 4 derived runs. This sums up to $10 + 10 + 5 + 5 = 30$ decent runs.

As we have seen, the set of 30 decent runs has some internal structure: It is the union of four classes of "similar" decent runs. Intuitively formulated, two decent runs are similar if they differ only in their initial part, i.e. are identical at long sight. In technical terms: Two runs R and R' are *postfix equivalent* iff there are finite runs Q and Q' such that $R = QP$ and $R' = Q'P$ for some common postfix P. For example, the runs in Fig. 4, Fig. 5 and Fig. 6 are pairwise postfix equivalent.

It remains to be proven that the system in Fig. 1 in fact has 30 decent runs, and 4 postfix equivalence classes. We do so in the more general context of n philosophers, in the next section.

3. The case of n philosophers

The generalization of the system of Fig. 1 to the case of n dining philosophers is canonical and obvious. We will show two properties of those systems:

Proposition 3.1. The system of n dining philosophers has

a) $2^n - 2$ decent concurrent runs

b) $n - 1$ postfix equivalence classes

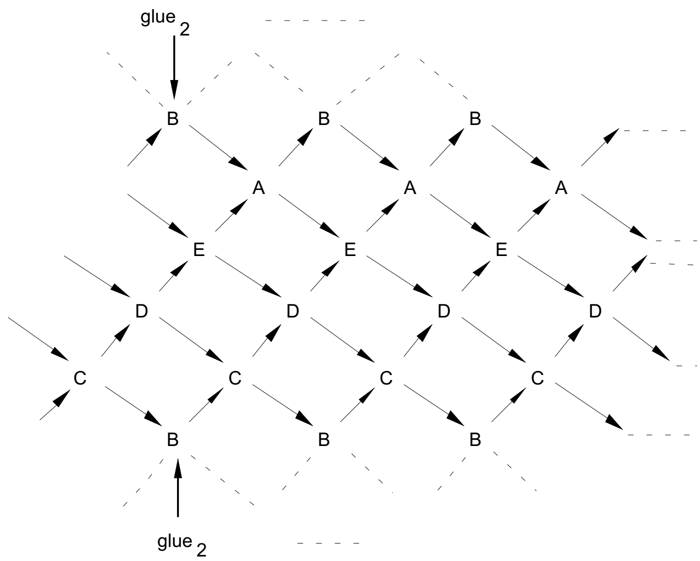

Figure 6. Concurrent run derived from Fig. 4 by skipping the first two occurrences of A, as well as the first occurrence of B, of D and of E (but no occurrence of C).

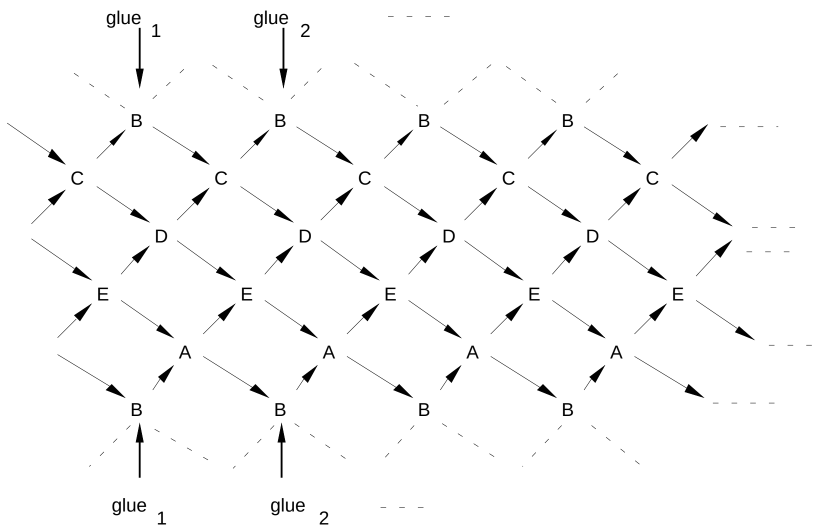

Figure 7. Concurrent run R_2 of Fig. 7

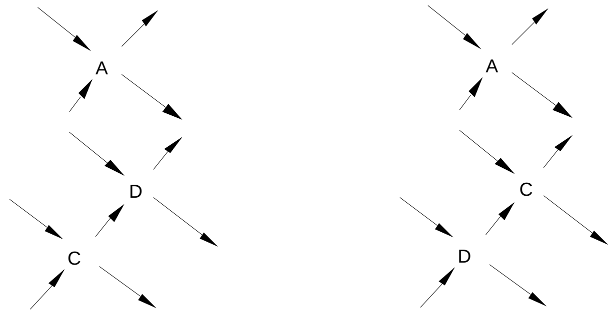

Figure 8. Different patterns of R_1 and R_2

Proof:
We show $a)$ in 6 steps

1. Each philosopher eats eventually (proof: assume a philosopher that never eats in a run. If none of its neighbors ever eats, the forks are available and he could start (maximality). If one or both neighbors eats, this philosopher eats only once and the respective fork will be and remains available afterwards (again maximality)).

2. Each fork is used eventually (immediate consequence).

3. Each fork behaves *left right left right* ... or *right left right left* ... (see above, and the definition of "decency").

4. Whenever each fork was used once, the total behavior is deterministic (choices are only possible initially).

5. It is not possible that each fork starts *left* or that each fork starts *right* (proof: at least one philosopher eats the first time before his neighbors eat the first time; one of his forks starts *left*, the other one *right*).

6. All other combinations of left- and right-starting forks are possible:

   We only have to show, that for each surjective mapping $init : forks \to \{left, right\}$ there is a run where each fork is used (at least once) starting its direction with *init (fork)*.

   Otherwise there is a decent run that can be extended but cannot be extended according to our rule that each fork starts with *init (fork)*. Clearly all forks are on the table after this run, because every run ending with an eating philosopher can be extended.

   This situation is only possible if a philosopher $p1$ and one of his neighbors $p2$ did not eat in this run, $p1$ could start, but the init mapping directs the fork towards $p2$.

   Assume w.l.o.g. that the problem is caused by the right fork, i.e., the right fork initially points to the right neighbor, and this is $p2$.

   Now consider the situation that the right neighbor $p3$ of $p2$ occurs in the run. Then the run could be extended by $p2$, because this extension would fulfil both decency and the *init* map. Since this

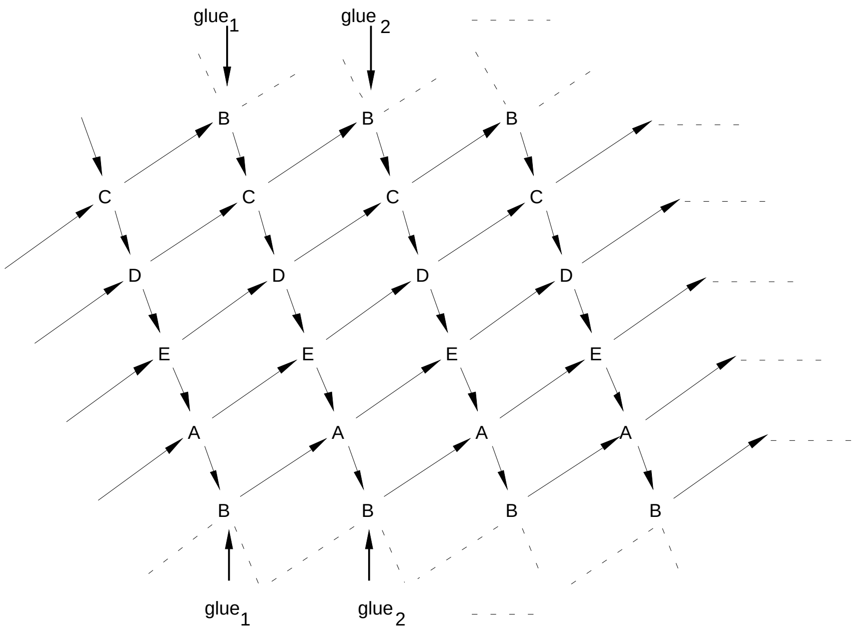

Figure 9. Concurrent run R_3

contradicts our assumption, $p3$ does not occur as well, and the fork between $p2$ and $p3$ initially points right, towards $p3$. But this means that $p2$ is in the very same situation as $p1$. By induction, we see that none of the philosophers started to eat yet and that all forks initially point in the same direction – a contradiction to our assumption.

There are $2^n - 2$ such combinations (trivial).

Property $b)$ is the consequence of an invariant property of decent runs; in fact, this invariant property *characterizes* the decent runs R of the n-philosophers system.

Each marking M of a decent run R defines an orientation function $orient_M : forks \rightarrow \{left, right\}$ on the forks, where $orient_M(f) = right$ if the next user of f in R is the right user of f. An example for such a function is the $init$ function as discussed in 6. of the above proof of a part $a) : init = orient_{M_0}$ for the initial marking M_0.

Let the ***right number*** of a function $f : forks \rightarrow \{left, right\}$ be a the number of forks g with $f(g) = right$.x

For all reachable markings M of R, the right number of $orient_M$ are identical. This proposition is obvious, because each occurrence of an eating cycle just swaps the orientation of two neighbored forks, as Fig. 10 outlines. So, the right number of the initial marking of a decent run R characterizes each reachable marking of R, it furthermore characterizes the postfix equivalence classes of R itself.

The right number of a surjective function $f : forks \rightarrow \{left, right\}$ with n forks varies between 1 and $n - 1$. Hence there are $n - 1$ postfix equivalence classes. □

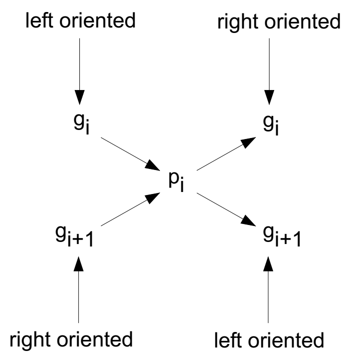

Figure 10. Orientation of forks in a concurrent run

Acknowledgements

A referee of a previous version of this paper pointed me at some inconsistencies. Jörg Desel contributed the idea and the proof to part $a)$ of Proposition 1. Thanks to both of them.

References

[1] E.W. Dijkstra *Hierarchical odering of sequential processes.* Acta Informatica 1, pp 115–138 (1971)