Problem Solving Using Process Algebra Considered Insightful

Jan Friso Groote$^{(\boxtimes)}$ ⓘ and Erik P. de Vink ⓘ

Department of Mathematics and Computer Science,
Eindhoven University of Technology, Eindhoven, The Netherlands
{J.F.Groote,E.P.d.Vink}@tue.nl

Abstract. Process algebras with data, such as LOTOS, PSF, FDR, and mCRL2, are very suitable to model and analyse combinatorial problems. Contrary to more traditional mathematics, many of these problems can very directly be formulated in process algebra. Using a wide range of techniques, such as behavioural reductions, model checking, and visualisation, the problems can subsequently be easily solved. With the advent of probabilistic process algebras this also extends to problems where probabilities play a role. In this paper we model and analyse a number of very well-known – yet tricky – problems and show the elegance of behavioural analysis.

1 Introduction

There is great joy in solving combinatorial puzzles. Numerous books have appeared describing those [13,33]. And although some of the puzzles are easy to solve once properly understood, they are real brain teasers for most people.

Many of these puzzles are about behaviour. Classical mathematics and logic hardly provides an effective context to solve such problems systematically. This is apparent if one considers classical analysis. But also fields like graph theory, combinatorics, combinatorial optimisation, probability theory, and even logic all require a translation of the problem to the mathematical domain that is generally not completely straightforward.

This is where process algebras come in. Process algebras are very suited to describe the behaviour often present in the puzzles mentioned. In the last decades numerous tools have been developed to provide insight in the behaviour denoted in a process algebra expression as it quickly became clear that the behaviour described in such an expression can be rather intricate. This gave rise to hiding of actions, behavioural reductions, various visualisation techniques, as well as modal logics to express and validate properties about behaviour.

The early 1970s can be seen as the period when process algebra was born. Both Milner and Bekič wrote a treatise expressing that actions were important to study behaviour [2,25,27]. It was the seminal work of Milner in 1981 that put process algebras on the map [28]. This had quite some effect. For instance Hoare presented CSP in 1978 as an advanced programming language [21], whereas

he presented it in 1985 as a process algebra [22]. The work on CSP has been developed into the impressive family of tools, FDR, that are based on failure divergence refinement [14,31].

The work on CCS also inspired the design of the language LOTOS [24] as a language to model communication services and protocols. A major role in its development was played by the Technische Hogeschool Twente (now Twente University) first in the completely formal standardisation of the language, with Brinksma as main editor, and later in activities to build tools around it. Notable are the extensive formal specifications of standard protocols, but also those of manufacturing systems, that were developed at the time [5,7,32]. The CADP toolset stems from this period [12]. It is the only major toolset still capable of analysing LOTOS specifications. Furthermore, it has become quite powerful throughout the years.

The Algebra of Communicating Processes (ACP) was developed in Amsterdam [3,4] around the same time. In order to model practical systems first PSF (Protocol Specification Formalism) was designed [26], which was followed by the simpler formalism μCRL [18], later renamed to mCRL2, which was also directed towards analysis of practical specifications [17]. All these LOTOS-like formalisms use data based on abstract equational datatypes. mCRL2 also supports time and these days also probabilities.

An important feature of mCRL2 is the support for a modal logic with time and data, which is very useful to investigate properties of the described behaviour. Temporal logic, with the operators $[F]$ and $[P]$, stems from [30]. Pnueli pointed to the applicability of formal logics to analyse behaviour [29]. For mCRL2 we are using the modal mu-calculus which is essentially Hennessy-Milner logic [20] with fixed points [23]. An alternative is the use of linear time logic (LTL [29]) or computational tree logic (CTL [8]), but these are far less expressive than the modal mu-calculus [15].

In this paper we show process algebraic models of a number of well-known mathematical puzzles. Most people find them hard to solve when they are confronted with them for the first time. We show that the puzzles can straightforwardly be modelled into process algebra and using the standard analysis tools, such as behavioural reduction, model checking, and visualisation, the solutions to these puzzles are easy to obtain.

The major observation is that process algebra is an industrious mathematical discipline in itself due to its capacity to understand worldly phenomena. Traditionally, there is a tendency to think that process algebras, or more generally formal methods, are intended to analyse software, protocols, and complex distributed algorithms. But the application to examples as in this paper shows that process algebra has an independent stand.

In this paper we use the language mCRL2, as we are acquainted with it, and it offers all we need, namely the capacity to express behaviour, data structures, probabilities, time (although we do not exploit time here), and modal formulas. mCRL2 has a very rich toolset offering a whole range of analysis methods, far more than we use for the examples in this article. In the following we do not

explain the tool nor the formalism. For this we refer to [17] or the webpage
www.mcrl2.org. The examples in this article are part of the mCRL2 distribution
downloadable from this website.

2 The Problem of the Wolf, Goat, and Cabbage

A problem that is well-known, at least to the people in Western Europe, is
the problem of the wolf, the goat, and the cabbage: A traveller walks through
stretched Russian woods together with a friendly wolf, a goat, and a cabbage.
Hungry and worn out, this companionship arrives at a river that they must cross.
There is a small boat only sufficient to carry our traveller and either the wolf,
the goat, or the cabbage. More than two do not fit. Crossing is complex as when
left unsupervised by the traveller, the wolf will eat the goat, while the goat will
eat the cabbage. The question to answer is whether it is possible to cross the
river without the goat or the cabbage being eaten.

This problem is quite old. It already appeared in a manuscript from the eighth
century A.D. [1]. Dijkstra wrote one of his well-known EWDs addressing this
problem [10]. The description in mCRL2 can be found in Table 1. The description
uses two shores, *left* and *right*, which are essentially sets of 'items', i.e. sets of
wolf, goat, and/or cabbage, resting at that shore. The opposite shore is given by
a function *opp*. An update function is used to remove items from one side and
add it to the other.

The behaviour of crossing the river is given by the process *WGC*. It has
two parameters, namely the shores s comprised of the sets of items at each side
of the river, and the current position p of the traveller. Observe that mCRL2
accommodates the use of data types such as sets which allows to neatly describe
the shores as a pair of sets containing items. The first two pairs of lines of the
WGC process express that if the wolf and the goat, or the goat and the cabbage
are at the side opposite of the traveller, something is eaten, expressed by the
action is_eaten. The symbol δ indicates that the process stops after this action.
Note that actions are typeset in a different font for easy recognition.

The third group of lines of the process expresses that the traveller can move
to the other shore alone, by performing the move action. To reduce the number of
transitions somewhat, we only allow this when no item can be eaten. The fourth
group of lines expresses that the traveller can transport one item from one shore
to the other. The last group of lines states that if the complete companionship
arrives at the right shore, the action done can take place. Initially, the traveller,
wolf, goat, and cabbage are at the left shore.

As the state space of this behaviour is small, it can nicely be visualised. See
Fig. 1. At the top we find the initial state, which is green. The goal state is
coloured blue at the bottom. All states where an action is_eaten can be done
are coloured red. They go to the white deadlocked state. All labels is_eaten
are removed for readability. States where nothing is eaten are green, yellow, or
blue. It is easy to see that there are paths from the green to the blue state
through yellow states by moving counter clock wise through the graph. One of
such paths is

Table 1. An mCRL2 description of the problem of the wolf, the goat, and the cabbage

---

sort *Item* = **struct** *wolf* | *goat* | *cabbage* ;
     Position = **struct** *left* | *right* ;
     Shores = **struct** *shores*(Set(*Item*), Set(*Item*));

map *opp* : *Position* → *Position* ;
     items : *Shores* × *Position* → Set(*Item*);
     update : *Shores* × *Position* × *Item* → *Shores* ;

var *s*, *t* : Set(*Item*);
     i : *Item*;
eqn *opp*(*left*) = *right*, *opp*(*right*) = *left* ;
     items(*shores*(*s*,*t*),*left*) = *s*, *items*(*shores*(*s*,*t*),*right*) = *t* ;
     update(*shores*(*s*,*t*),*right*,*i*) = *shores*(*s* − {*i*},*t* + {*i*});
     update(*shores*(*s*,*t*),*left*,*i*) = *shores*(*s* + {*i*},*t* − {*i*});

proc *WGC*(*s* : *Shores*, *p* : *Position*) =
        {*wolf*, *goat*} ⊆ *items*(*s*,*opp*(*p*)) →
           is_eaten(*goat*)·δ +
        {*goat*, *cabbage*} ⊆ *items*(*s*,*opp*(*p*)) →
           is_eaten(*cabbage*)·δ +
      ¬({*wolf*, *goat*} ⊆ *items*(*s*,*opp*(*p*))) ∧ ¬({*goat*, *cabbage*} ⊆ *items*(*s*,*opp*(*p*))) →
           move(*opp*(*p*))·*WGC*(*s*,*opp*(*p*)) +
        $\sum_{i:Item}$.(*i* ∈ *items*(*s*,*p*)) →
           move(*opp*(*p*),*i*)·*WGC*(*update*(*s*,*opp*(*p*),*i*),*opp*(*p*)) +
        items(*s*,*right*) ≈ {*wolf*, *goat*, *cabbage*} →
           done·δ;

init *WGC*(*shores*({*wolf*, *goat*, *cabbage*}, ∅),*left*);

$$\text{move}(\textit{right, goat}) \cdot \text{move}(\textit{left}) \cdot \text{move}(\textit{right, wolf}) \cdot \text{move}(\textit{left, goat}) \cdot$$
$$\text{move}(\textit{right, cabbage}) \cdot \text{move}(\textit{left}) \cdot \text{move}(\textit{right, goat}) \cdot \text{done}.$$

Inspection of the state space also reveals that there is one other essential solution to this problem, namely one where the places of the wolf and the cabbage are exchanged. This is no surprise as the wolf and the cabbage have symmetrical roles. Note that it is also clear why this puzzle is considered tricky. Each solution requires the counter intuitive step of moving the goat three times across the river, an insight that requires humans to overcome their default mental set.

Fig. 1. The state space of the problem with the wolf, the goat and the cabbage (Color figure online)

For this puzzle we are lucky that the number of states is sufficiently small to be depicted. In general this is not the case. Fortunately, modal formulas are a marvellous tool to investigate properties of behaviour. In this case we want to know whether there is a path from the initial state to a state where the action `done` is possible, while no action `is_eaten` is possible in any of the states on this path. In the modal mu-calculus as available in the mCRL2 toolset this is expressed by

$$\mu X.((\langle\text{true}\rangle X \vee \langle\texttt{done}\rangle\text{true}) \wedge \neg\langle\exists i: Item.\texttt{is\_eaten}(i)\rangle\text{true}).$$

The use of the minimal fixed point guarantees that the action `done` must be reached in a finite number of steps. The modality $\langle\text{true}\rangle$ says that an arbitrary action can be done. Checking this formula instantly yields true confirming that the traveller can safely reach the other shore with all the companions intact.

3   Crossing a Rope Bridge in the Dark

The second problem is similar in nature to the first but not as well-known. Four people of different age arrive at a rope bridge across a canyon in the night. They need to cross the bridge as quickly as possible. Each person has its own time to cross the bridge, namely, 1, 2, 5, and 10 min. Unfortunately, the bridge can only carry the weight of two persons simultaneously. To make matters worse, they

Table 2. The problem of crossing a rope bridge specified in mCRL2

```
sort  Position = struct this_side | far_side;
      Person = struct p₁ | p₂ | p₃ | p₄;

map   travel_time : Person → ℕ;
      initial_location : Person → Position;
      other_side : Position → Position;
      max_time : ℕ;

var   p : Person;
eqn   initial_location(p) = this_side;
      travel_time(p₁) = 1; travel_time(p₂) = 2;
      travel_time(p₃) = 5; travel_time(p₄) = 10;
      other_side(this_side) = far_side;
      other_side(far_side) = this_side;
      max_time = 20;

proc  X(light_position : Position, location : Person → Position, time : ℕ) =
          time ≤ max_time ∧ ∀p:Person.location(p) ≈ far_side →
             ready(time)·δ +
          Σ_{p:Person} ·
             time ≤ max_time ∧ location(p) ≈ light_position →
                move(p,other_side(location(p)))·
                    X(other_side(light_position),
                       location[p → other_side(location(p))],
                       time + travel_time(p)) +
          Σ_{p,p':Person} ·
             p ≉ p' ∧ time ≤ max_time ∧ location(p) ≈ light_position ∧
                       location(p') ≈ light_position →
                move(p,p',other_side(location(p)))·
                    X(other_side(light_position),
                       location[p → other_side(location(p))][p' → other_side(location(p'))],
                       time + max(travel_time(p),travel_time(p')));

init  X(this_side,initial_location,0);
```

only carry one flashlight. Crossing without the flashlight is impossible. So, the flashlight needs to be returned for others to cross. The question is to find the minimal time in which the group of people can cross the bridge.

The problem is modelled in mCRL2 in Table 2. The location of each person is now given by a function *location*: *Person → Position*. The function update construction is used to change a function. The expression *location*[$p → s$] represents a new function which is equal to *location* except that person p is now mapped

to position s. The parameter *time* records the total time to cross the bridge and *light_position* keeps track of the place of the flashlight.

The behaviour consists of three summands, and is a direct translation of the problem. The first summand expresses that if all people are at the far side, a `ready` action is done, reporting the time to cross. The second summand expresses that one person crosses the bridge, and the third summand indicates that two people move to the other side together.

Natural numbers in mCRL2 are specified using abstract data types and have no upper bound. This means that the state space of this problem is infinite as there are inefficient crossing strategies that can take arbitrarily large amounts of time. Although not strictly necessary, as mCRL2 is very suitable to investigate infinite state spaces, it is generally a wise strategy to keep state spaces finite and even as small as reasonably possible. Solving the problem naively, quickly leads to a crossing time of 19 min. We therefore limit the maximal crossing time to 20 min and focus on in the question whether crossing under 19 min is possible.

The generated state space is somewhat larger, namely 470 states and 1607 transitions, which disallows inspection as an explicit graph. Fortunately, we can use the tool `ltsview`, which can visualise the structure of large transition systems [16], in some case up to millions of states. Pictures made by `ltsview` appear to be rather pointless pieces of art at first glance, but when investigated, provide remarkable insight in the depicted behaviour.

The behaviour of crossing the rope bridge is depicted in Fig. 2 at the left. The initial state is at the top. The layering corresponds to the number of crossings of the bridge. The individually visible states and structures that grow to the side of the picture indicate deadlocks, i.e., states where the crossing time exceeds 20. For instance, the states at the end of the outward moving structure at the top right indicate that the bound of 20 min can be exceeded in three crossings. The red disk (the one but lowest) is the disk containing the action `ready`(17). There are no `ready` actions with a lower argument. This indicates that the bridge cannot be crossed in less than 17 min.

`ltsview` is not the most efficient way to inspect which `ready` actions are possible. By searching for actions while generating the state space it becomes immediately clear that the actions `ready`(17), `ready`(19) and `ready`(20) are possible. A trace to `ready`(17) is

$$\texttt{move}(p_2, p_1, \mathit{far\_side}) \cdot \texttt{move}(p_1, \mathit{this\_side}) \cdot \texttt{move}(p_4, p_3, \mathit{far\_side}) \cdot$$
$$\texttt{move}(p_2, \mathit{this\_side}) \cdot \texttt{move}(p_2, p_1, \mathit{far\_side}) \cdot \texttt{ready}(17).$$

This trace shows why this puzzle is hard to solve. The idea to save time to let the two slowest persons cross simultaneously does not easily come to mind for most people.

Using modal logics we can also easily check that 17 is the most optimal crossing time. The next formula, which says that there is a path to the action `ready`(17) and not to any action `ready`(n) for any $n < 17$, is readily proven to hold:

$$\langle \text{true}^\star \cdot \texttt{ready}(17) \rangle \, \text{true} \wedge \forall n{:}\mathbb{N}.(n < 17 \rightarrow [\text{true}^\star \cdot \texttt{ready}(n)] \, \text{false}).$$

Fig. 2. An `ltsview` visualisation of crossing a rope bridge and the game tic-tac-toe (Color figure online)

4   A Winning Strategy in Tic-tac-toe

Finding winning strategies in games can also be neatly expressed and studied in process theory. One of the simplest well-known games that can be analysed in this way is tic-tac-toe. Essentially, tic-tac-toe consists of a 3 by 3 board where two players alternatingly put a naught or cross at empty positions on the board. The first player that has three of naughts or crosses in a row, horizontally, vertically or diagonally, wins the game.

Table 3 contains a rather natural formalisation of this game. The playing board is given by a function from pairs of naturals to pieces. A less elegant formulation uses lists of lists of pieces, but for state space generation this is much faster. A player moves by putting its own piece at an empty position on the board using the action `put`. The action `win` is used to indicate that one of the players did win. The most complex function is $did\_win(p, b)$, checking whether player p, represented by a piece, did win the game. In the formalisation we use $c \rightarrow p \diamond q$ denoting 'if c then p else q'.

The total behaviour of this game has 5479 states and 17109 transitions, which is not very large. This behaviour is depicted in Fig. 2 at the right. The red dot at the right middle indicates where player 'naught' can win. There are more such states two disks lower, but they are hardly visible in the figure.

Although the transition system for this game is relatively small, it makes no sense to investigate it directly to determine whether the player that starts the game has a winning strategy. Fortunately, modal formulas come to the rescue.

Table 3. An mCRL2 formalisation of tic-tac-toe

<div style="border:1px solid">

sort *Piece* = **struct** *empty* | *naught* | *cross* ;
      Board = $\mathbb{N}^+ \to \mathbb{N}^+ \to Piece$;

map *empty_board* : *Board* ;
      did_win : *Piece* × *Board* → $\mathbb{B}$;
      other : *Piece* → *Piece* ;

var *b* : *Board* ;
      p : *Piece* ;
      i, *j* : $\mathbb{N}^+$;
eqn *empty_board*(*i*)(*j*) = *empty* ;
      other(*naught*) = *cross* ; *other*(*cross*) = *naught* ;
      did_win(*p*, *b*) =
          $(\exists i : \mathbb{N}^+ .(i \leqslant 3 \wedge b(i)(1) \approx p \wedge b(i)(2) \approx p \wedge b(i)(3) \approx p)) \vee$
          $(\exists j : \mathbb{N}^+ .(j \leqslant 3 \wedge b(1)(j) \approx p \wedge b(2)(j) \approx p \wedge b(3)(j) \approx p)) \vee$
          $(b(1)(1) \approx p \wedge b(2)(2) \approx p \wedge b(3)(3) \approx p) \vee$
          $(b(1)(3) \approx p \wedge b(2)(2) \wedge p \approx b(3)(1) \approx p)$;

proc *TicTacToe*(*board* : *Board*, *player* : *Piece*) =
          did_win(*other*(*player*), *board*) →
              win(*other*(*player*))·δ
          ◇ $\left(\sum_{i,j:Pos} .(i \leqslant 3 \wedge j \leqslant 3 \wedge board(i)(j) \approx empty) \right.$ →
              put(*player*, *i*, *j*)·
                  $TicTacToe(board[i \to board(i)[j \to player]], other(player)))$;

init *TicTacToe*(*empty_board*, *cross*);

</div>

The following formula expresses that player 'cross' has a winning strategy. It says that there is a way to put a cross on the board after which player 'cross' wins, or for every counter move by player 'naught', X must hold again, saying that also in that case player 'cross' has a winning strategy. This formula is invalid. There is no winning strategy for the player 'cross', and due to symmetry neither for player 'naught'.

$$\mu X.\langle \exists i, j : \mathbb{N}^+.\mathtt{put}(cross, i, j)\rangle \left(\langle \mathtt{win}(cross)\rangle \mathrm{true} \vee \right.$$
$$\left. ([\exists i, j : \mathbb{N}^+.\mathtt{put}(naught, i, j)]\, X \, \wedge \langle \mathrm{true}\rangle \mathrm{true}) \right)$$

Note that we use of a minimal fixed point operator expressing that winning must happen within a finite number of steps. As there are only a limited number of moves in tic-tac-toe this is always satisfied, hence a maximal fixed point operator could also have been used.

5   The Monty Hall Problem

Processes algebras have seen various extensions. One of these extensions is the addition of probabilities, which gives rise to the interesting combination of nondeterministic and probabilistic behaviour. This opens up the field of probabilistic puzzles to be modelled. The Monty Hall problem is a very nice example, because when understood is it very simple, yet most people fail to solve it properly.

Table 4. An mCRL2 specification of the Monty Hall quiz

sort *Doors* = **struct** $d_1 \mid d_2 \mid d_3$;
init dist *door_with_prize* : *Doors*[1/3].
     dist *initially_selected_door_by_player* : *Doors*[1/3].
       player_collects_prize(*initially_selected_door_by_player* $\not\approx$ *door_with_prize*)·δ;

The Monty Hall problem is a tv-quiz from the 1960s. A player can win a prize when he opens one of three doors with the prize behind it. Initially, the player selects a door with probability $\frac{1}{3}$. Subsequently, the quizmaster opens one of the remaining doors showing that it does not hide the prize. The question is whether the player should switch doors to optimise his winning probability.

The problem is expressed in the specification in Table 4. The process only consists of a single action player_collects_prize(b) where the boolean argument b is true if a prize is collected. The **dist** keyword is used to indicate a probability distribution. The process **dist** $x : S[D(x)].p$ indicates that variable x of sort S is selected with probability distribution $D(x)$. One of the doors hides the prize. This door is represented by the variable *door_with_prize* which can have values d_1, d_2, or d_3, each with a probability of $\frac{1}{3}$. Initially, the player selects a door. If the player decides to switch doors after the quizmaster opened a door, the player has a prize if and only if the initially chosen door did *not* carry the prize. This is expressed by the use of not equal sign ($\not\approx$) in the argument of the action. If the player decides to stick to the door that was initially selected, the not equal sign should be replaced by equality.

The resulting state space has 9 transitions each with a probability $\frac{1}{9}$. It is convenient to apply a probabilistic bisimulation reduction on the transition system. This leads to the reduced transition system in Fig. 3. It is clearly visible that the action player_collects_prize(true) can be done with probability $\frac{2}{3}$. Thus, when switching doors the probability of obtaining a prize is $\frac{2}{3}$, opposed to $\frac{1}{3}$ when not switching doors.

Fig. 3. The reduced probabilistic transition system for the Monty Hall problem

6   The Problem of the Lost Boarding Pass

More complex probabilistic problems can become rather hard even with the full strength of probability theory at ones disposal. Yet modelling the problem in mCRL2 is again pretty straightforward. The tools can subsequently help to obtain the required answer.

A particularly intriguing puzzle is that of the lost boarding pass as it has a remarkable answer, defying the intuition of most people trying to solve the problem: There is a plane with 100 seats. The first passenger boarding the plane lost his boarding ticket and selects a random seat. Each subsequent passenger will use his own seat unless it is already occupied. In that case he also selects a random seat. The question is what the probability is that the *last* passenger entering the plane will sit in his own seat.

The behaviour is modelled in Table 5. The number N is the number of seats, which is set to 100. The behaviour of entering the plane is characterised by two parameters. The parameter *number_of_empty_seats* indicates how many seats are still empty in the plane. The parameter *everybody_has_his_own_seat* indicates that all remaining seats correspond exactly with the places for all passengers that still have to board the plane. Except if the number of empty seats is 0. In that case it indicates whether the last passenger got its own seat.

Initially the first passenger selects his seat at random. With probability $\frac{1}{N}$ he will end up at his own seat. This corresponds with the situation where b is true. In the main process *Plane*, when all passengers have boarded the plane, the action last_passenger_has_his_own_seat indicates by its argument whether the last passenger got his own seat. If not all passengers boarded the plane yet, a next passenger enters (indicated by the action enter) and then it can either be that he finds his own seat free (b_0 is true) or occupied (b_0 is false). If everybody is sitting at is own seat this next passenger will for sure find his own seat free. Otherwise, he finds his own seat free with probability $1-1/number\_of\_empty\_seats$ as exactly one person is sitting on a wrong seat.

When this next passenger finds his own seat free he can sit down. This is done by the action select_seat with two parameters. But if his own seat is occupied, he must randomly select a seat for himself. If he selects the seat such

Table 5. An mCRL2 specification of the lost boarding pass

map $N : \mathbb{N}^+$;
eqn $N = 100$;

proc $Plane(everybody\_has\_his\_own\_seat : \mathbb{B}, number\_of\_empty\_seats : \mathbb{N}) =$
$\quad\quad (number\_of\_empty\_seats \approx 0) \rightarrow$
$\quad\quad\quad \texttt{last\_passenger\_has\_his\_own\_seat}(everybody\_has\_his\_own\_seat) \cdot \delta$
$\quad\quad \diamond (\texttt{enter} \cdot$
$\quad\quad \mathbf{dist}\ b_0 : \mathbb{B}[if(everybody\_has\_his\_own\_seat, if(b_0, 1, 0),$
$\quad\quad\quad\quad if(b_0, 1 - 1/number\_of\_empty\_seats, 1/number\_of\_empty\_seats))].$
$\quad\quad\quad b_0 \rightarrow \texttt{select\_seat} \cdot$
$\quad\quad\quad\quad Plane(everybody\_has\_his\_own\_seat, number\_of\_empty\_seats - 1)$
$\quad\quad\quad\quad \diamond \mathbf{dist}\ b_1 : \mathbb{B}[if(b_1, 1/number\_of\_empty\_seats, 1 - 1/number\_of\_empty\_seats)].$
$\quad\quad\quad\quad \texttt{select\_seat} \cdot$
$\quad\quad\quad\quad\quad Plane(if(number\_of\_empty\_seats \approx 1, everybody\_has\_his\_own\_seat, b_1),$
$\quad\quad\quad\quad\quad\quad number\_of\_empty\_seats - 1));$

init   **dist** $b : \mathbb{B}[if(b, 1/N, (N-1)/N)].Plane(b, N-1);$

that all passengers are sitting on their assigned seats (modulo a permutation) this is indicated in the variable b_1, where this passenger has probability

$$1/number\_of\_empty\_seats$$

of doing this.

The generated state space turns out to be linear in the size of the number of seats. It has 791 states and 790 transitions. Modulo strong probabilistic bisimulation there are 399 states and 398 transitions. It has the shape of a long sequence, as depicted in Fig. 4. Detailed exploration of this figure indicates that whence all the remaining passengers correspond to the remaining seats the last passenger will certainly get his own seat. Yet it is not obvious what the probability for the last passenger to get his own seat is. For this we use two – at present experimental – tools[1]. The first one applies a probabilistic weak trace reduction. The obtained state space, see Fig. 5, is rather non-exciting but indicates clearly that the probability of the last passenger to end up at its own seat is $\frac{1}{2}$. The remarkable property of this exercise is that this probability is independent of the number of seats.

---

[1] The tools are by Olav Bunte (evaluation of modal formulas on probabilistic transition systems) and Ferry Timmers (probabilistic trace reduction).
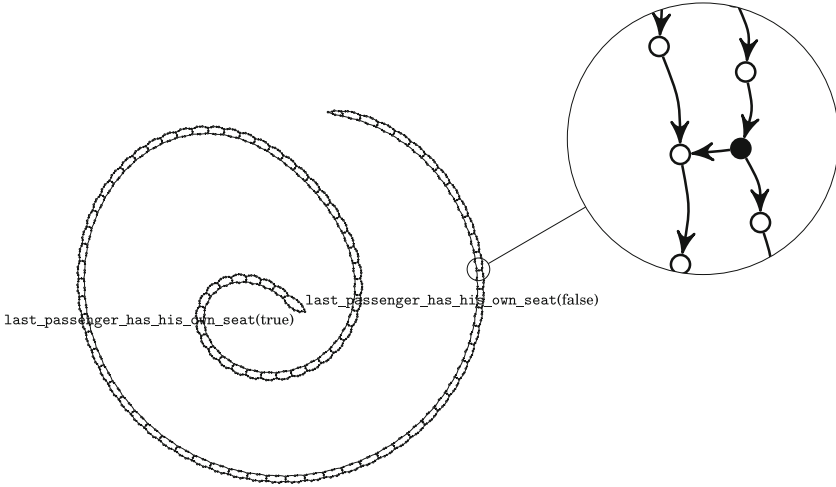
Fig. 4. The state space of the problem of the lost boarding pass with 100 passengers


Fig. 5. The state space of the lost boarding pass problem modulo weak trace equivalence

There is another way to obtain this probability by employing modal formulas over reals. These formulas are derived from the modal mu-calculus but deliver a real number, instead of a boolean. In this case the formula is just

$$\langle \text{true}^\star \cdot \texttt{last\_passenger\_has\_his\_own\_seat}(\text{true}) \rangle \, \text{true}$$

which is possible as the state space is deterministic. Needless to say that the verification of this formula yields $\frac{1}{2}$ as well.

7 Concluding Remarks

Process algebra is generally well-suited to solve many behaviour-oriented mathematical puzzles. In this paper we have used the process algebraic framework of mCRL2 to show how to model a number of such puzzles. Subsequently, the standard analysis tools available in mCRL2 (and occasionally an experimental one) were used for behavioural reduction, model checking and visualisation. From this it is clear that process algebra has a wider scope than the usual fields of software analysis and distributed computing in which it finds many applications.

Process algebra focuses on behavioural aspects of the subject of study. The underpinning algebraic and equational theory allows to relate to logics, in particular modal logics [6], as descriptions of properties or requirements over space, time, and probabilities. Logical characterisations and their assessment via model checking are a valuable replacement in situations where visual techniques, highlighted for the puzzles discussed here, become impractical.

Also other authors indicated that a notion of behaviour or state space is required for proper conditional reasoning, especially in the probabilistic setting. In [19] the distinction is made between 'naive' and 'sophisticated' space. For the Monty Hall puzzle this amounts to the three doors for the naive space, and to sequences of events for the sophisticated space. In the process algebraic modelling of the problem, it is exactly the latter that is determined by the specified behaviour, thus making the underlying protocol explicit.

Although we defend the use of process algebra as a qualitatively better approach to solving behavioural problems, this is a subjective opinion, influenced by our experience with process algebras. To substantiate this in a more objective manner one should measure how much time people need to solve particular problems with particular techniques, for instance by psychological tests.

If process algebraic techniques become commonplace, it might be that the nature of 'tricky' puzzles will shift where the proper behaviour is not directly obvious. Nice examples are for instance Freudenthal problems, containing knowledge, like the Muddy Children puzzle [11]. Translating knowledge into behaviour often requires a twist. In such cases dynamic epistemic logic might be more suitable [9].

References

1. Hadley, J., Singmaster, D.: Problems to sharpen the young. Math. Gaz. **76**(475), 102–126 (1992). doi:10.2307/3620384
2. Bekič, H.: Towards a mathematical theory of processes. Technical report TR25.125, IBM Laboratory, Vienna (1971). Also appeared in Jones, C.B. (ed.) Programming Languages and Their Definition, Lecture Notes in Computer Science, vol. 177. Springer (1984)
3. Bergstra, J.A., Klop, J.W.: Fixed point semantics in process algebras. Report IW 206, Mathematisch Centrum, Amsterdam (1982)
4. Bergstra, J.A., Klop, J.W.: Process algebra for synchronous communication. Inf. Comput. **60**(1/3), 109–137 (1984)
5. Biemans, F., Blonk, P.: On the formal specification and verification of CIM architectures using LOTOS. Comput. Ind. **7**(6), 491–504 (1986)
6. Blackburn, P., van Benthem, J., Wolter, F. (eds.): Handbook of Modal Logic. Studies in Logic and Practical Reasoning, vol. 3. Elsevier, Amsterdam (2007)
7. Brinksma, E., Karjoth, G.: A specification of the OSI transport service in LOTOS. In: Yemini, Y., Strom, R.E., Yemini, S. (eds.) Protocol Specification, Testing and Verification IV, pp. 227–251. North-Holland, Amsterdam (1984)

8. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Kozen, D. (ed.) Logic of Programs 1981. LNCS, vol. 131, pp. 52–71. Springer, Heidelberg (1982). doi:10.1007/BFb0025774

9. van Ditmarsch, H., van der Hoek, W., Kooij, B.: Dynamic Epistemic Logic. Studies in Epistemology, Logic, Methodology, and Philosophy of Science, vol. 337. Springer, Heidelberg (2008). doi:10.1007/978-1-4020-5839-4

10. Dijkstra, E.W.: Pruning the search tree. EWD1255. www.cs.utexas.edu/users/EWD/transcriptions/EWD12xx/EWD1255.html. Accessed June 2017

11. van Emde Boas, P., Groenendijk, J., Stokhof, M.: The conway paradox: its solution in an epistemic framework. In: Groenendijk, J., Janssen, T.M.V., Stokhof, M. (eds.) Truth, Interpretation, and Information: Selected Papers from the Third Amsterdam Colloquium, pp. 159–182. Foris Publications, New York (1984)

12. Garavel, H., Lang, F., Mateescu, R., Serwe, W.: CADP 2011: a toolbox for the construction and analysis of distributed processes. Int. J. Softw. Tools Technol. Transf. **15**(2), 89–107 (2013)

13. Gardner, M.: My Best Mathematical and Logic Puzzles. Dover, Downers Grove (1994)

14. Gibson-Robinson, T., Armstrong, P., Boulgakov, A., Roscoe, A.W.: FDR3 — a modern refinement checker for CSP. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 187–201. Springer, Heidelberg (2014). doi:10.1007/978-3-642-54862-8_13

15. Cranen, S., Groote, J.F., Reniers, M.A.: A linear translation from CTL* to the first-order modal μ-calculus. Theoret. Comput. Sci. **412**(28), 3129–3139 (2011)

16. Groote, J.F., van Ham, F.: Interactive visualization of large state spaces. Int. J. Softw. Tools Technol. Transf. **8**(1), 77–91 (2006)

17. Groote, J.F., Mousavi, M.R.: Modeling and Analysis of Communication Systems. The MIT Press, Cambridge (2014). (See for the toolset www.mcrl2.org)

18. Groote, J.F., Ponse, A.: The syntax and semantics of μCRL. Report CS-R9076, CWI, Amsterdam (1990)

19. Grünwald, P., Halpern, J.Y.: Updating probabilities. In: Darwiche, A., Friedman, N. (eds.) Uncertainty in Artificial Intelligence, pp. 187–196. Morgan Kaufman, Burlington (2002)

20. Hennessy, M., Milner, R.: On observing nondeterminism and concurrency. In: de Bakker, J., van Leeuwen, J. (eds.) ICALP 1980. LNCS, vol. 85, pp. 299–309. Springer, Heidelberg (1980). doi:10.1007/3-540-10003-2_79

21. Hoare, C.A.R.: Communicating sequential processes. Commun. ACM **21**(8), 666–677 (1978)

22. Hoare, C.A.R.: Communicating Sequential Processes. Prentice Hall International, Upper Saddle River (1985)

23. Kozen, D.: Results on the propositional μ-calculus. Theoret. Comput. Sci. **27**, 333–354 (1983)

24. ISO 8807:1989. Information processing systems - Open Systems Interconnection - LOTOS - A formal description technique based on the temporal ordering of observational behaviour. ISO/IECJTC1/SC7 (1989)

25. Milner, R.: An approach to the semantics of parallel programs. In: Proceedings Convegno di Informatica Teorica, pp. 283–302, Pisa (1973)

26. Mauw, S., Veltink, G.J.: A process specification formalism. Fundam. Inform. **XIII**, 85–139 (1990)

27. Milner, R.: Processes: a mathematical model of computing agents. In: Rose, H.E., Shepherdson, J.C. (eds.) Proceedings Logic Colloquium 1972, pp. 158–173. North-Holland, Amsterdam (1973)

28. Milner, R. (ed.): A Calculus of Communicating Systems. LNCS, vol. 92. Springer, Heidelberg (1980)
29. Pnueli, A.: The temporal logic of programs. In: Foundations of Computer Science, pp. 46–57. IEEE, Piscataway (1977)
30. Prior, A.N.: Time and Modality. Oxford University Press, Oxford (1957)
31. Roscoe, A.W.: Understanding Concurrent Systems. Springer, Heidelberg (2010). doi:10.1007/978-1-84882-258-0
32. van Sinderen, M., Ajubi, I., Caneschi, F.: The application of LOTOS for the formal description of the ISO session layer. In: Turner, K.J. (ed.) Formal Description Techniques, pp. 263–277. North-Holland, Amsterdam (1989)
33. Winkler, P.: Mathematical Puzzles. A Connaisseur's Collection. A.K. Peters, Natick (2004)