Conceptual Connections around Density Determination in Cellular Automata

Pedro P.B. de Oliveira

Universidade Presbiteriana Mackenzie, Faculdade de Computação e Informática,
Rua da Consolação 896, Consolação, 01302-907 São Paulo, SP – Brazil
pedrob@mackenzie.br

Abstract. A recurring and well studied benchmark problem in the context of computations with cellular automata is the attempt to determine which is the most frequent cell state in an arbitrary initial configuration. Although extremely simple in formulation, the problem has unveiled a rich web of conceptual connections which, at the same time, have enlarged and challenged our understanding about how to perform computations within cellular automata. Here, we outline such a conceptual web, and provide a personal assessment of some of its loose ends, with possibly fruitful paths to address them.

Keywords: Density classification task, DCT, global majority problem, emergent computation, cellular automata, number conserving rule.

1 Introduction

By looking up the cellular automata (CAs) literature in order to find out which are the most well studied computation performed by them, two tasks readily come out: the firing squad synchronisation problem and the density classification task, the latter outnumbered by the former. It is useful to compare the two.

In a simplified way, while the firing squad demands that a state synchronisation be achieved among the various cells of the entire lattice, density classification imposes that a decision be made, based upon information gathered over the entire lattice, namely, the more frequent cell state; in both cases, the effectiveness of the computation carried out is granted according to their reaching predefined global configurations, after sufficiently long time.

Crucial in these problems is the existence of trivial solutions by processes allowed global access to the entire lattice; this is naturally forbidden by a CA based solution, by definition, so as to enforce that the solution to a global problem be achieved by purely local means.

Notwithstanding such a necessary similarity, it is useful to identify differences between the two problems. While the firing squad can be solved by a single rule (with many examples already known, in fact), density classification cannot.

Also, while the CAs related to the former task have typically been obtained through stepwise and careful thought-out engineering, many times supplemented by formal methods, the CAs for the latter have been most effectively found by

J. Kari, M. Kutrib, and A. Malcher (Eds.): AUTOMATA 2013, LNCS 8155, pp. 1–14, 2013.

search methods, mostly evolutionary, in the space of possible rules. Additionally, density classification is structurally much simpler than the firing squad, in that it has a much simpler definition; it is more general in terms of the notion it is based upon – that of a majority – which also appears in various other contexts in science; it is more abstract, insofar as the notion of majority provides a direct link to the formal language notion of context sensitivity; and also more flexible in terms of the variations it admits.

So, such a diverse combination of features renders density classification a paradigmatic testbed for studying computations in the context of cellular automata ([Mitchell, 1998]). As such, much have already been learned in its respect, which ended up unveiling a web of conceptual connections, far beyond its original motivation of a mere tool to help addressing the notion of fault tolerance in cellular automata. Nevertheless, very challenging questions have been opened up, several of them still in need of a response. It is about those conceptual connections and their current loose ends that this paper is about, even if in condensed form.

In the following section we define the problem, and in the next we examine what is known and what remains unknown so far about it. The subsequent section concludes, by providing a personal assessment of some open issues still challenging us, and possibly fruitful paths to address them. The presentation has a survey flavour, systematic but general, with various details omitted.

2 Density Classification

Cellular automata (CAs) consist of a grid-like regular lattice of cells, together with a state transition rule, the cells having an identical pattern of local connections to others, the lattice being subjected to some boundary condition, usually periodic ([Wolfram, 2002]; [Kari, 2005]). Each cell can take on one of k discrete states, assumedly varying from 0 to k-1, and the neighbourhood of a cell is defined as the cell, together with those connected to it. The rule yields the next state for every cell, as a function of its neighbourhood; at each time step, all cells synchronously have their states updated. The size of the neighbourhood is usually written as $2r+1$, where r is the radius (or range) of the automaton. The particular case of binary one-dimensional CAs with $r = 1$ yields the so-called elementary CA rules and its corresponding elementary space.

In order to refer to rules of any given space, their numbers are defined herein according to Wolframs standard lexicographic ordering of the state transitions, in that the rule number is the decimal representation of the state transition outputs, with the left-most output referring to the neighbourhood consisting of all cells in the state $k-1$, with the remaining state transitions ordered in lexicographic order, down to the right-hand neighbourhood that consists of all cells in the 0-state.

In its standard and simplest formulation, the density classification task (DCT, for short) states that a binary, one-dimensional CA has to converge to a final configuration of all cells being in the 1-state, when an odd-sized initial configuration, in periodic boundary condition, has more 1s than 0s, and to a configuration