Phase Lengths in the Cyclic Cellular Automaton

Kiran Tomlinson

Department of Computer Science, Carleton College, Northfield, MN, USA

Abstract— *The cyclic cellular automaton (CCA) model of excitable media displays remarkable spiral wave behavior and evolves through distinct phases. Previous work has shown that the number of states k plays a pivotal role in the long-term behavior of the system, but there has been no systematic investigation into its phase lengths. We provide an empirical method for determining phase start and end points and describe the dependence of CCA phase lengths on k. We find that the length of each phase exhibits a distinctive power-law relation with k in a manner independent of lattice size. In addition, we perform bootstrapping and parameter sensitivity analysis to determine the uncertainty in the power law exponents and coefficients.*

Keywords: Cellular automata, dynamical systems.

1. Introduction

Excitable media are dynamical systems in which excited regions transmit their energetic state to their neighbors, but can only do so periodically. In two dimensions, excitable systems can form spiral waves [1]. Canonical examples of spiral wave excitable media include the oscillating Belousov-Zhabotinsky reaction [2], chemical signals in *Dictyostelium* amoebae [3], and electrical signals in the heart [4]. These systems are governed by complex physical and chemical processes—as such, their dynamics are difficult to study directly. Cellular automata are widely used to model dynamical systems due to their simplicity and complex emergent behavior, and several cellular automaton models of spiral wave excitable media have been proposed [5], [6], [7]. The cyclic cellular automaton (CCA) is one such model, originally described in [5].

For a wide range of parameters, the CCA model undergoes a sequence of distinct phases that eventually culminate in a stable spiral wave regime. These phases can be clearly seen in CCA visualizations (Fig. 1). Previous work on the CCA has described the range of parameters for which these phases develop [8], [9], [10], the shapes of spirals generated by different neighborhoods [11], and the self-organizing properties of the model [12]. However, no systematic investigation of phase lengths in the CCA model has yet been made.

In this paper, we present new computer simulation results describing the relationship between the number of states k and the lengths of CCA phases under two different neighborhoods. To do so, we describe a method of identifying phase transition points in the CCA. These empirical results are robust to grid size and provide new insight into the

behavior of the CCA model. In addition, the method we use to determine phase transition points from noisy data can be applied to other models that exhibit distinct phases.

1.1 CCA definition

Like other cellular automata, the CCA exists on a 2-dimensional integer lattice and consists of *cells* that occupy a finite number of *states* (also called *types* or *colors*). Formally, let $\zeta_t : \mathbb{Z}^2 \to \{0, \ldots, k-1\}$ denote the state of a cell at step t. We assign $\zeta_0(x)$ uniformly at random for all cells x and the system evolves according to the update rule:

$$\mathcal{N}_t^+(x) = \{y \in \mathcal{N}(x) \mid \zeta_t(x) + 1 \equiv \zeta_t(y) \ (\text{mod } k)\} \quad (1)$$

$$\zeta_{t+1}(x) = \begin{cases} (\zeta_t(x) + 1) \ \text{mod } k & \text{if } |\mathcal{N}_t^+(x)| \geq \theta \\ \zeta_t(x) & \text{otherwise} \end{cases} \quad (2)$$

We use $\mathcal{N}(x)$ to denote the set of neighbors of cell x, $\mathcal{N}_t^+(x)$ the neighbors of x in the state one higher at time t, and θ the threshold number of such neighbors required to increment x's state. We focus on the $\theta = 1$ case and on the two neighborhoods most widely used in cellular automata, the von Neumann and Moore neighborhoods. The former consists of the four cells that share an edge with x and the latter also includes the four diagonals. In theoretical discussions, the entire lattice $\mathbb{Z}^2$ is considered [5], [8]. However, practical implementations and experimental analyses restrict the system to a finite lattice, most often with periodic boundary conditions [10], [11], [12], a practice which we follow.

1.2 CCA phases

The cyclic cellular automaton progresses through a sequence of phases, named the *debris*, *droplet*, *defect*, and *demon* phases [5]. As long as k is not too small [10] (or too large, on a finite lattice), every CCA run follows the same pattern. See Fig. 1 for images of each phase.

Debris. In the first phase, the random grid goes through a short burst of activity during which neighboring cells with consecutively initialized states are promoted according to the update rule. After a brief flurry of updates, most of the grid is covered by inactive and disorganized debris.

Droplets. Some clusters of cells happen to border other cells in the correct sequence of ascending states, allowing the cluster to spread. These clusters, called droplets because of their circular and liquid appearance, are dominated by a small number of types that propagate through the droplet
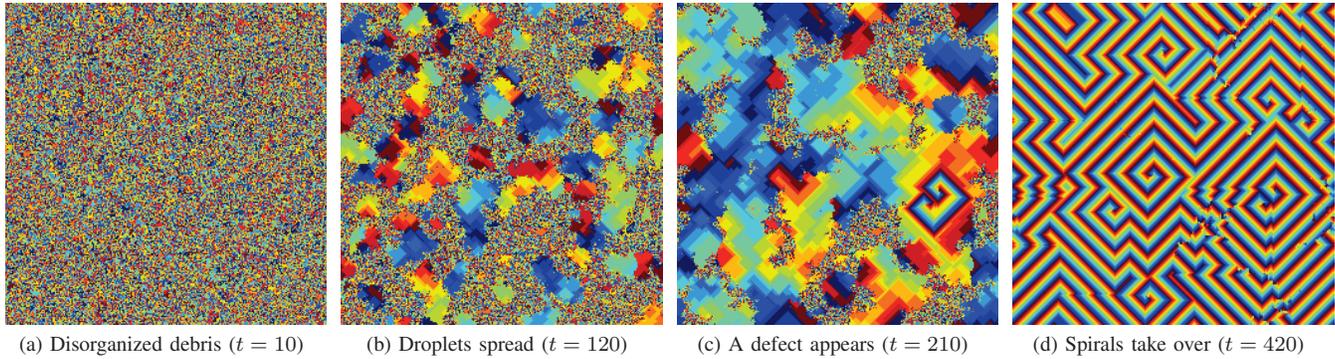
(a) Disorganized debris ($t = 10$)   (b) Droplets spread ($t = 120$)   (c) A defect appears ($t = 210$)   (d) Spirals take over ($t = 420$)

Fig. 1: Snapshots of the CCA (256×256 periodic grid, von Neumann $\mathcal{N}$, $k = 14$)

in waves. As time progresses, droplets increase in size by taking over the surrounding debris.

Defects. With high probability, a loop of updating cells called a defect forms at the intersection of droplets, or at the debris boundary of a droplet [5]. Defects then give rise to spirals, which spread across the grid, wiping out any remaining debris and droplets.

Demons. Some defects contain the correct number of cells to generate period k spirals whose constituent cells update at every step. These optimal spirals are known as demons. Demons absorb longer period spirals, eventually dominating the entire grid [5]. The system then stabilizes and the demons continue to spiral ad infinitum.

These phases characterize the behavior of the CCA for a wide range of neighborhoods and values of k [9]. Some basic experimentation reveals that the lengths of the phases change when when k is varied; this fact is apparent in figures from [5], [10]. However, the exact nature of these relationships is not immediately obvious because of the complexity of the model's emergent behavior.

In the following sections, we empirically identify the dependence of phase lengths in the CCA model on the number of states k under both the von Neumann and Moore neighborhoods. In particular, we analyze the lengths of the debris phase, the droplet phase, and defect phase.

2. Methods

We implemented the CCA in Python [13] using NumPy [14] and ran 1024 independent 500-step trials of the model for each parameter setting we considered. Specifically, we ran simulations on periodic lattices of size 256×256, 512×512, and 1024×1024 with both the Moore and von Neumann neighborhoods and with k ranging from 7 to 20. The vital part of our analysis is the extraction of phase length data from these simulations.

2.1 Identifying phase transitions

In order to determine the lengths of CCA phases, we need to be able to determine the steps at which phase transitions occur. This is accomplished by measuring the number of cells whose states change between consecutive lattice configurations ζ_{t-1} and ζ_t, which we denote $\Delta(t)$:

$$\Delta(t) = \sum_{x \in C} (\zeta_t(x) - \zeta_{t-1}(x) \bmod k) \qquad (3)$$

where C is the set of cells in the lattice. A similar quantity, the *open bond proportion*, is discussed in [5]. They call a bond between two neighboring cells x and y *open* at step t if $\zeta_t(x) - \zeta_t(y) \in \{-1, 1\} \bmod k$. By the update rule, the lesser of the two must then be promoted in step $t + 1$. The open bond proportion differs slightly from the cell update count $\Delta(t)$ since a cell update may be due to more than one open bond. We found it more straightforward to measure $\Delta(t)$, and the resulting plots are qualitatively identical (compare our Fig. 2 with Fig. 9 in [5]).

As seen in Fig. 2, the slope of $\Delta(t)$ is characteristic of the current phase. This key observation allows us to identify CCA phase transitions. In the debris phase, the initial flurry of updates rapidly dies down and most local regions become static. The global minimum of $\Delta(t)$ marks the end of the debris phase and the beginning of the droplet phase. Then, $\Delta(t)$ increases steadily as droplets expand and become more active. Once the first defect forms, the curve takes an upward turn as the spiral begins to spread. This marks the end of the droplet phase and the beginning of the defect phase. Once spirals have taken over the grid, the $\Delta(t)$ curve levels off, signaling the end of the defect phase. From this point, any period k spirals slowly take over the grid.

Based on these observations, the end of the debris phase can be identified as the step at which $\Delta(t)$ reaches its minimum. The ends of the droplet and defect phase are more difficult to identify. They occur in the bends between the roughly linear segments of the $\Delta(t)$ curve. It is then natural to identify the phase transition points as the steps when the curvature of the $\Delta(t)$ curve is highest—that is, at local extrema of the $\Delta''(t)$. However, the second derivative of the discrete data points is extremely noisy. Fortunately, differentiating noisy data is a well-studied problem [15].
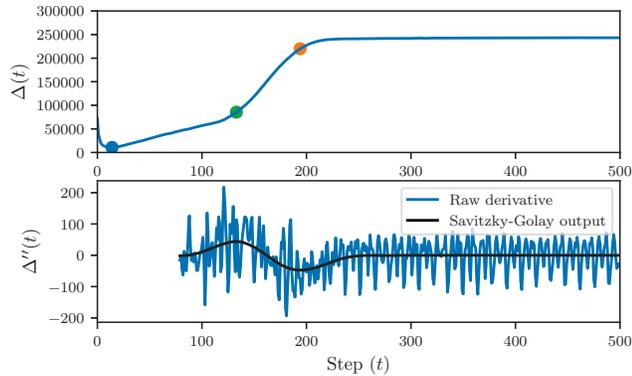
Fig. 2: $\Delta(t)$ (top) and $\Delta''(t)$ (bottom) curves from a single trial ($k = 13$, von Neumann $\mathcal{N}$, 512×512 grid). The $\Delta''(t)$ curve is shown with (black) and without (blue) the Savitzky-Golay filter. The three phase transition points are marked on the $\Delta(t)$ plot. The $\Delta''(t)$ curve is cut off on the left due to the large spike around the debris phase.

2.2 Smooth differentiation

We use a Savitzky-Golay filter to differentiate $\Delta(t)$ without compromising the signal-to-noise ratio [16], [17]. To process a data point $\Delta(t^*)$, the Savitzky-Golay filter fits a low-degree polynomial using linear least-squares to the points in a window around t^* and reports the desired derivative of this polynomial at t^*. For the CCA data, we use polynomials of degree 3, which allows us to compute the second derivative using the Savitzky-Golay filter. The widths of the windows were chosen manually such that the data was sufficiently smoothed without introducing unwanted distortion. The need to choose this parameter by hand is the primary drawback of the Savitzky-Golay filter [16]. To account for any systematic bias introduced by this manual choice, we perform a sensitivity analysis by varying all window widths randomly by up to $\pm 50\%$ and measuring the effect on our results.

From the smoothed $\Delta''(t)$ data, we can then easily identify the local maximum that signifies the upward turn of $\Delta(t)$ at the droplet-defect transition and the local minimum that signifies the downward turn at the defect-demon transition. Once these transition points are identified, the phase lengths can be found directly. By performing this on every one of 1024 trials, we can then find the mean phase lengths along with the associated variance.

3. Results

We found that all phase lengths in the cyclic cellular automaton exhibit a power law dependence on k. That is, we find that the length L of a phase is well described by $L = ak^b$, for constants a and b. Taking the logarithms of both sides allows us to find a and b by linear regression:

$$\log L = b \log k + \log a. \tag{4}$$

Table 1: Phase length power law exponents

(a) von Neumann $\mathcal{N}$

	256×256	512×512	1024×1024	**Mean**
Debris	2.55 ± 0.01	2.55 ± 0.01	2.57 ± 0.00	2.56 ± 0.01
Droplet	4.81 ± 0.08	4.89 ± 0.22	4.75 ± 0.23	4.81 ± 0.07
Defect	3.08 ± 0.10	3.16 ± 0.11	3.24 ± 0.11	3.15 ± 0.06

(b) Moore $\mathcal{N}$

	256×256	512×512	1024×1024	**Mean**
Debris	2.52 ± 0.01	2.57 ± 0.01	2.60 ± 0.01	2.56 ± 0.01
Droplet	4.34 ± 0.06	4.32 ± 0.07	4.37 ± 0.07	4.34 ± 0.03
Defect	2.88 ± 0.13	2.82 ± 0.09	2.77 ± 0.10	2.81 ± 0.06

Table 2: Phase length power law log coefficients

(a) von Neumann $\mathcal{N}$

	256×256	512×512	1024×1024	**Mean**
Debris	-3.99 ± 0.02	-4.00 ± 0.02	-4.05 ± 0.01	-4.03 ± 0.01
Droplet	-7.58 ± 0.20	-7.80 ± 0.57	-7.42 ± 0.59	-7.59 ± 0.18
Defect	-3.88 ± 0.23	-4.05 ± 0.29	-4.25 ± 0.29	-4.03 ± 0.15

(b) Moore $\mathcal{N}$

	256×256	512×512	1024×1024	**Mean**
Debris	-5.02 ± 0.03	-5.14 ± 0.02	-5.24 ± 0.01	-5.20 ± 0.01
Droplet	-8.31 ± 0.16	-8.25 ± 0.16	-8.38 ± 0.20	-8.30 ± 0.10
Defect	-4.76 ± 0.34	-4.59 ± 0.25	-4.45 ± 0.26	-4.58 ± 0.16

The exponent b captures the degree of the impact k has on each phase length. See Fig. 3 for the linear regressions and Tables 1 and 2 for the experimentally observed values of b and $\log a$. We estimate the uncertainty in b and $\log a$ through a combination of bootstrapping and perturbing the Savitzky-Golay parameters. First, we split the 1024 trials into 8 independent samples and separately perform the linear regression on each sample. We then calculate 95% confidence intervals from the standard error of the regression results. In addition, we redo the analysis 8 times with Savitzky-Golay window widths multiplied by random scalars drawn uniformly from $[0.5, 1.5]$. We again find the 95% confidence interval on these new regression results. This allows us to account for potential bias in the manual choice of these window widths. A conservative uncertainty estimate is then found by adding the sizes of the confidence intervals; these are the uncertainties reported in Tables 1 and 2. We use inverse-variance weighting to find the mean over grid sizes and the associated uncertainty.

Rarely, noise in the data may cause the local maximum and minimum of the $\Delta''(t)$ curve to swap places, yielding a spurious negative phase length. We filter out such trials. This occurred in only 48 of 28,000 CCA runs, so it has a negligible effect.

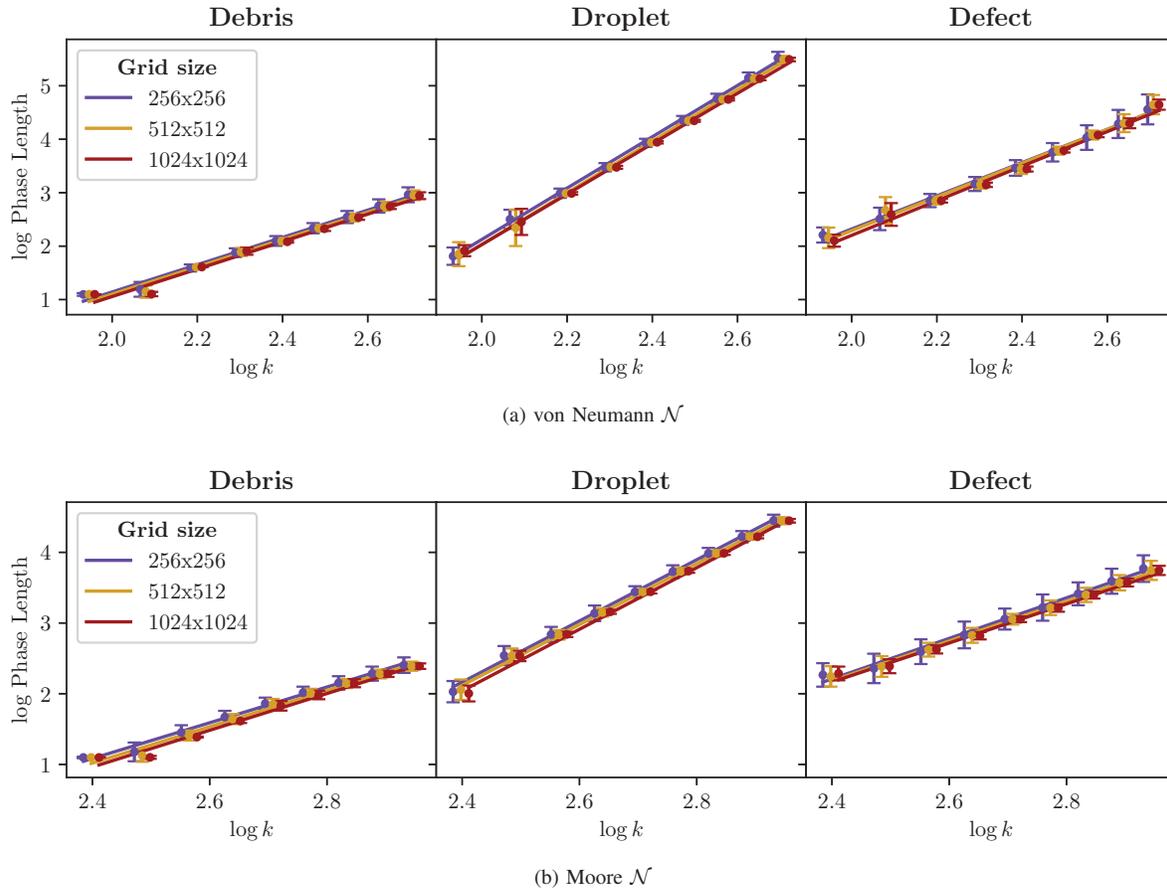We find that the droplet phase is most sensitive to changes

Fig. 3: Phase length dependences on k. Horizontal separation added to isolate data points from different grid sizes. Error bars show the sample standard deviations.

in k, followed by the defect and debris phases. This ordering is consistent between neighborhoods, and b is significantly higher for the droplet and defect phase under the von Neumann $\mathcal{N}$. These relationships are independent of grid size.

4. Discussion

We found that the lengths of the debris, droplet, and defect phases exhibit power-law relations with k with characteristic powers. These relationships have not been previously identified. We now briefly discuss intuitive interpretations of these results.

4.1 Phase length sensitivities

These phase length observations square well with our intuitive understanding of the model. As k increases, there are several compounding factors that elongate the droplet phase. First of all, it is known that defects can only form at the boundary of a droplet or at the intersection of droplets [5]. With a higher value of k, fewer droplets arise from the debris. With fewer droplets on the lattice, we then have fewer

droplet collisions and a smaller total droplet perimeter. Both of these factors therefore increase the expected time before a defect forms and brings an end to the droplet phase. In addition, in order for a droplet predominantly of type i to spread, it must encounter a cell of type $i + 1 \mod k$ on its border. Then, as the $i + 1$ wave spreads across the droplet, some part of it must encounter an $i + 2 \mod k$ cell on the border, and so the process continues. When k is higher, each wave must travel further before such a cell is encountered, slowing the expansion of the droplet.

The defect phase begins when the first spiral is seeded by a defect. The spiral then spreads out from this epicenter. Like a droplet, a spiral must encounter a cell of the correct type in order to absorb it. However, cells in spirals have very low periods compared to cells in droplets, so spirals spread much more quickly than droplets, with less extreme dependence on k.

4.2 Other methods for finding phase transitions

Methods that do not involve $\Delta(t)$ could also be used to identify phase transitions in the CCA. A more precise way to

find the droplet-defect phase transition would be to identify the exact step at which the first defect appears, since a defect can be directly detected as a loop of cells of adjacent types with a positive winding number [5]. To do this, we would need an efficient method for finding such loops. Perhaps a backtracking strategy could be used to eliminate the need to search the entire grid for such a loop at each step: we could play the history of the grid backward to narrow in on the defect that generated the first spiral. For the defect-demon phase transition, another approach would be to find the first step at which every cell belongs to a spiral by measuring all cell periods.

4.3 Conclusions

We found that the complex emergent behavior in the CCA model of excitable media depends on simple power law relationships with a key parameter, the number of cell types k. These relationships are conserved across grid sizes, showing that they are characteristic of the model itself. The relative lengths of phases are also maintained between neighborhoods. We presented a method for identifying phase transitions in the CCA by applying Savitzky-Golay differentiation to the $\Delta(t)$ curve. Future work could use alternate methods of identifying phase transitions or apply the same approach to other cellular automata. It is also possible that these power laws could be derived analytically from the probabilistic behavior of the CCA, although this appears to pose a considerable challenge.

Acknowledgment

References

[1] J. P. Keener, "Waves in excitable media," *SIAM J. Appl. Math.*, no. 3, pp. 528–548, 1980.

[2] J. P. Keener and J. J. Tyson, "Spiral waves in the Belousov-Zhabotinskii reaction," *Phys. D*, vol. 21, no. 2, pp. 307 – 324, 1986.

[3] E. Pálsson, K. J. Lee, R. E. Goldstein, J. Franke, R. H. Kessin, and E. C. Cox, "Selection for spiral waves in the social amoebae *Dictyostelium*," *Proc. Natl. Acad. Sci. USA*, vol. 94, no. 25, pp. 13 719–13 723, 1997.

[4] J. M. Davidenko, A. V. Pertsov, R. Salmonsz, W. Baxter, and J. Jalife, "Stationary and drifting spiral waves of excitation in isolated cardiac muscle," *Nature*, vol. 355, pp. 349–351, 1992.

[5] R. Fisch, J. Gravner, and D. Griffeath, "Cyclic cellular automata in two dimensions," in *Spatial Stochastic Processes*.   Springer, 1991, pp. 171–185.

[6] M. Gerhardt, H. Schuster, and J. J. Tyson, "A cellular automaton model of excitable media," *Phys. D*, vol. 46, no. 3, pp. 392–415, 1990.

[7] J. M. Greenberg and S. P. Hastings, "Spatial patterns for discrete models of diffusion in excitable media," *SIAM J. Appl. Math.*, vol. 34, no. 3, pp. 515–523, 1978.

[8] R. Durrett and D. Griffeath, "Asymptotic behavior of excitable cellular automata," *Exp. Math.*, vol. 2, no. 3, pp. 183–208, 1993.

[9] R. Fisch, J. Gravner, and D. Griffeath, "Threshold-range scaling of excitable cellular automata," *Stat. Comput.*, vol. 1, pp. 23–39, 1991.

[10] K. A. Hawick, "Neighbourhood and number of states dependence of the transient period and cluster patterns in cyclic cellular automata," in *Proc. 10th Int. Conf. on Sci. Comput. (CSC'13)*, 2013. [Online]. Available: http://worldcomp-proceedings.com/proc/p2013/CSC7339.pdf

[11] C. A. Reiter, "Medley of spirals from cyclic cellular automata," *Comput. & Graphics*, vol. 34, pp. 72–76, 2010.

[12] C. R. Shalizi and K. L. Shalizi, "Quantifying self-organization in cyclic cellular automata," in *Proc. SPIE 5114, Noise in Complex Systems and Stochastic Dynamics*, L. Schimansky-Geier, D. Abbott, A. Neiman, and C. V. den Broeck, Eds., 2003.

[13] Python Software Foundation, "Python," 2018, version 3.7.1. [Online]. Available: https://www.python.org

[14] NumPy Developers, "Numpy," 2018, version 1.15.4. [Online]. Available: https://www.numpy.org

[15] I. Knowles and R. J. Renkak, "Methods for numerical differentiation of noisy data," *Electron. J. Diff. Eqns.,* Conference 21, pp. 235–246, 2014.

[16] W. Gander and J. Hřebíček, *Solving Problems in Scientific Computing Using Maple and MATLAB*, 3rd ed.   Berlin: Springer, 1997, ch. 9, pp. 135–139.

[17] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Anal. Chem.*, vol. 36, no. 8, pp. 1627–1639, 1964.