18.204: CHIP FIRING GAMES

ANNE KELLEY

ABSTRACT. Chip firing is a one-player game where piles start with an initial number of chips and any pile with at least two chips can send one chip to the piles on either side of it. When all of the piles have no more than a single chip, the game ends. In this paper we review fundamental theorems related to this game on a two dimensional number line, including the fact that termination and final configuration are independent of the sequence of moves made and prove the number of moves required for termination is bounded. We then extend the game to consider distinct chips also on a two dimensional number line, where chips are represented by integers and firings result in a comparison of two chips in a pile such that the smaller is sent left and larger is sent right. We prove that for odd numbers of chips some final configurations are sorted while others are unsorted and conjecture that for even numbers of chips the final configuration is necessarily sorted.

1. INTRODUCTION

Chip-firing as a field is relatively young compared to most areas of mathematical research. The first chip-firing game was described by Spencer in 1986 when he wrote a paper about a so-called "balancing game". Interestingly, he did not set out to study the properties of a chip-firing process, but instead was studying a vector balancing problem where he tried to assign positive or negative labels to a set of vectors such that when multiplied by their label, the sum of all the vectors was within a certain bound. His answer, however, came in terms of what is now considered to be a chip-firing game. [3]

In Spencer's original game a pile of chips begin at the origin of an infinite two dimensional number line of piles. On each move, every pile of chips sends half of its contents one pile to the right and the other half of its contents one pile to the left. When a pile sends its chips to its neighboring piles it is said to have been $fired$. If it contains an odd number of chips, a single chip remains in its original pile. This process continues until each pile contains no more than a single chip and it is no longer possible for any chips to move as a result. We will formally describe this process in the next section.

Naturally, once other mathematicians were introduced to the idea, they sought to generalize it so it could be applied to more processes. In 1989 R.J. Anderson, L. Lovász, P. Shor, J. Spencer, E. Tardos, and S. Winograd generalized the process to allowing the player to only move a single pair of chips at a time in their paper *Disks, balls, and walls: analysis of a combinatorial game.* Their paper is the focus

Date: April 24, 2016.
Key words and phrases. Chip Firing, Abelian Sandpile Model.

of this summary paper, and we will explore and restate their results in the third section. In particular, we ask ourselves natural questions such as: Which piles are able to ever hold chips? Will the firing process ever end? If so where will the chips be when it does? And how many firings will be needed until it does?

Since the work of Anderson et al. the idea has been generalized beyond a two dimension number line onto arbitrary graphs, both directed and undirected. Furthermore, shortly after Spencer's work, theoretical physicists studying the toppling of sandpiles and pink noise introduced a similar idea they called the Abelian Sandpile Model [4]. In general, the two describe the same process, and so the ideas explored in this paper are often referred to as both chip-firing and the Abelian Sandpile Model. While these more general ideas are not explored within this paper, the reader is encouraged to explore them at their leisure.

Since this field of research is still young, there are still many open questions and interesting variations that have not been fully explored. In the last major section of this paper we present one such variation where the chips are considered to be distinct. We explore some very basic results, such as showing final configurations of odd numbers of chips can either be sorted or unsorted, but the problem in general for even numbers of chips is still open.

We'll begin in section two by restating Spencer's original game. Then, in section three we recap classical results on the more general game as described by Anderson et al. Finally in section four we extend to the case where chips are considered to be distinct.

2. Preliminaries: Spencer's Original Game

Formally, Spencer's original chip-firing game can be defined as follows:

Let $A_{i,t}$ be the number of chips in the i^{th} pile on the t^{th} timestep. Then we can write the following recurrence relationship for this process:

$$A_{i,t} = \lfloor \frac{A_{i-1,t-1}}{2} \rfloor + \lfloor \frac{A_{i+1,t-1}}{2} \rfloor + a_{i,t-1} \bmod 2$$

.

In this recurrence the first term refers to the chips that come from the pile to the left, the second term refers to the chips that come from the pile to the right, and the third term refers to the chip that may remain stationary in the pile on a given move.

We're specifically interested in the initial conditions

$$A_{i,0} = \begin{cases} 2n+1 & i = 0 \\ 0 & i \neq 0 \end{cases}$$

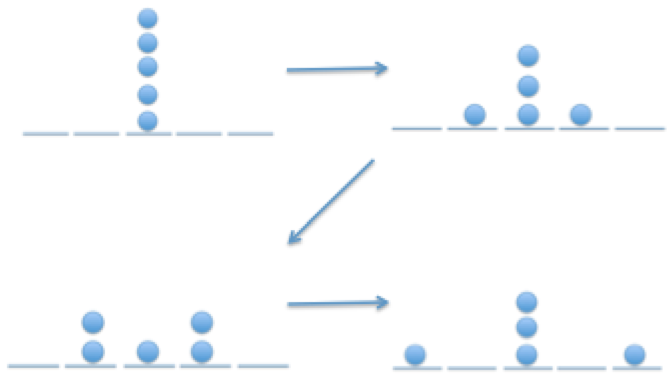

FIGURE 1. Chip-Firing Game. To transition from the first to second configuration the pile at the origin is fired. Then to transition to the third it is fired again. The transition from the third to fourth configuration takes a single move in Spencer's original game where all piles are fired simultaneously, but two moves in the generalized game where the pile at -1 and pile at 1 must be fired individually

Throughout many of the proofs, we are interested in which piles chips are present in and how many chips are in each pile. To this end, we make the following definition that describes the placement of chips.

Definition 2.1. Configuration
Let the assignment of chips to piles be called the *configuration*. We say that two configurations are the same if corresponding piles with the same position relative to the origin contain the same number of chips in both of the configurations.

However, as is sometimes the case in mathematics, in order to answer the questions we have about Spencer's original game we seek to describe the results of a more general chip-firing game.

Our more general version of the game is largely set up the same, except on each turn we fire only a single pile and from this pile we move only a single pair of chips, again sending one to the right and one to the left. The example in Figure 1 illustrates the process. The termination condition remains unchanged. This game is more general in the sense that there is a series of moves in this game that mimics the behavior of Spencer's original game and therefore anything we can show for any sequence of moves in the more general game must also apply to Spencer's original game.

3. Classical Results: Anderson Et Al's Generalization

We consider the more general game as presented above, with $2n+1$ chips beginning at the origin. The following lemmas, theorems, and proofs in this section are

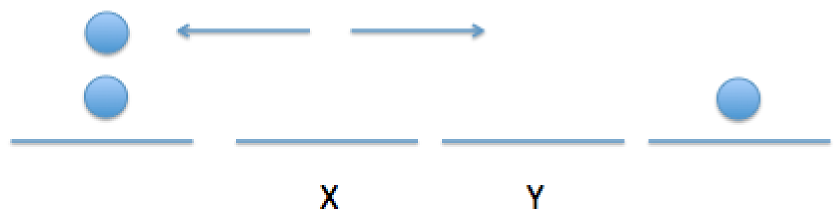

Figure 2. Lemma 3.1: x and y are two consecutive unoccupied piles. If x was fired last, one chip would have been moved both left and right as indicated by the arrows. However, this would have put a chip into pile y resulting in a contradiction.

adapted from the paper of Ansderson et al. who originally explored this variation.

A natural question is to consider whether or not such a game would terminate, and if so what sort of final configuration the chips might be left in when it does. To this end, we begin with a series of lemmas to build up to a pair of theorems that show that there is only a single final configuration and describe the final configuration.

Lemma 3.1. *It is impossible to have two unoccupied piles in a row that fall between occupied piles on either side.* [1]

Proof. We proceed by contradiction. Consider the two piles labeled x and y in Figure 2. Since there are occupied piles on either side of x and y and all the chips began at the origin and can move only a single pile in each turn, at some point both x and y must have been occupied. Suppose, without loss of generality, that pile x was the last occupied of the two. Then, in order for x to now be empty, x must have been fired. However, firing x means that a chip must have been moved from x to y and y must have been occupied after x was occupied. This is a contradiction, and therefore it is impossible to ever have two unoccupied piles between occupied piles. □

Next we set to bound how far away from the origin the chips are able to move.

Lemma 3.2. *Chips never move outside the range of piles [-4n,4n].* [1]

Proof. We have $2n+1$ chips and as a result of the previous lemma, there can be at most $2n$ interior empty piles if there is one empty pile between each of the occupied piles. As a result, we know that the range of the chips must always be less than or equal to $4n+1$. An example of this is shown in the Figure 3. The final thing to notice is that after the first move where all the chips begin at the origin, there must always be chips on both the positive and negative sides of the origin. As a

FIGURE 3. Lemma 3.2: Consider this example of $n = 2$. There are 5 chips, and at most 4 interior empty piles, since there can never be two empty piles in a row. The total possible maximum range between chips is therefore 9. As a result none of the chips can ever move outside of $[-9, 9]$ since the chips can't all be on the same side of the origin.

result, we know that the chips can never move more than $4n$ from the origin in either direction. □

Note that this bound is not tight, as we will show later. Also of interest is how many moves it takes before we reach the final configuration.

Lemma 3.3. *The number of moves, M, until completion is bounded by* $16n^3 + 8n^2$. [1]

Proof. Let a_i be the number of chips in the i^{th} pile in the final configuration. We consider the sum of the squared distances from the center, which we can write as $D = \sum_i a_i * i^2$. In the very worst case, all of the chips are at either $4n$ or $-4n$ at termination, since we determined in the last lemma that no chip can ever move outside these bounds. Therefore, we know that

$$D = \sum_i a_i * i^2 \leq (2n + 1)(\pm 4n)^2 = 32n^3 + 16n^2$$

Then we consider what happens to the sum of the squared distances for each move that is made. If pile i is fired, then we add a chip to pile $i - 1$, add a chip to pile $i + 1$, and take two away from pile i so

$$\Delta D = (i + 1)^2 + (i - 1)^2 - 2i^2) = 2$$

This means that for every move that is made, the sum of the squared distances increases by two. Since we've bounded $D \leq 32n^3 + 16n^2$ this means that there can be at most $M \leq 16n^3 + 8n^2$ moves. □

Asymptotically this means that the number of moves needed until termination is $O(n^3)$. A natural question, of course, is how tight this bound is. The actual number of moves to completion is actually the same regardless of the sequence of moves made, and equals [2]

$$M = \frac{k(k-1)(2k-1)}{6}$$

where

$$k = \lfloor \frac{n+2}{2} \rfloor.$$

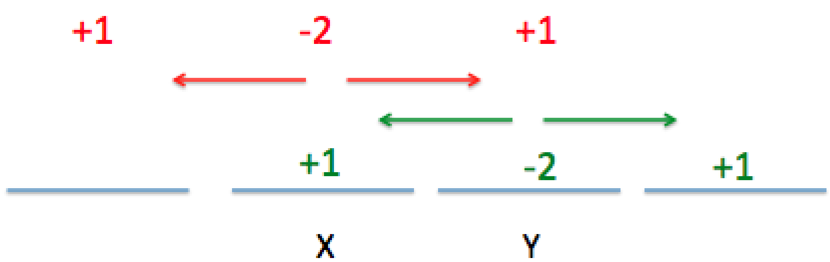

FIGURE 4. Lemma 3.5: The red arrows and numbers show the changes in number of chips in each pile when x is fired. The green arrows and numbers show the changes in number of chips in each pile when y is fired. Notice that because addition commutes, the total number of chips in each pile is the same regardless of whether x is fired and then y or y is fired and then x.

Notice that this is also $O(n^3)$ so our rough bound above is fairly tight.

Corollary 3.4. *The chip-firing game must terminate*

Proof. In the previous lemma we proved that number of moves is bounded. Since one move occurs on each time step, the number of time step must also be finite and therefore the game must terminate. □

We need one final supporting lemma, before we can prove that there is only a single final configuration that is always reached.

Lemma 3.5. *Moves Commute* [1]

Proof. Consider two subsequent firings of piles x and y, where x and y are distinct. There are two cases, either x and y have more than a single pile between them or x and y have at most one pile between them.

If the two piles have more than a single pile between them, there is no overlap between the piles affected when they are fired. As a result, it doesn't matter in which order the two are fired, the result must still be the same.

In the other case, we first note that firing a pile only adds chips to other piles so it cannot make a pile that previously was able to fire unable to subsequently fire. Then, we note that the affects of firing the two piles amount to additions and subtractions from the relevant piles. Since both addition and subtraction commute, the moves also commute.

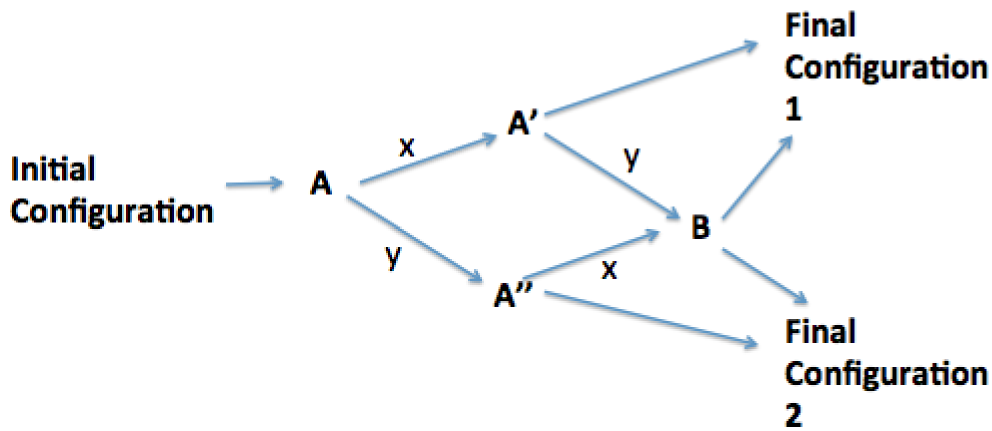

FIGURE 5. Theorem 3.8: We claim that A is the last configuration with two final configurations as descendants. We then consider making moves x and y such that the two resulting configurations result in different final configurations. However, since moves commute, we then apply moves y and x such that both end up in configuration B. Then, B must have both final configurations as descendants, resulting in a contradiction

Therefore, starting from configuration A and making move x followed by move y must result in the same configuration as starting from configuration A and making move y followed by move x. □

We are now ready to prove our first major result, that there is only a single final configuration but first we need the following two definitions.

Definition 3.6. Descendant
We say that configuration B is a *descendant* of configuration A if there is a sequence of moves that transforms configuration A into configuration B

Definition 3.7. Final
We say that a configuration is $final$ if there are no more piles that can be fired.

Theorem 3.8. *There is only a single final configuration for the chips* [1]

Proof. We assume for the purposes of contradiction that there are at least two final configurations. We choose the last configuration A such that it has two final configurations but each of its descendants has only a single final configuration. We consider, as shown in Figure 5, the configurations A' and A'' that come from firing piles x and y respectively. However, as we proved in the previous lemma, moves commute, so making move y and x from configurations A'' and A' respectively both result in configuration B. However, this means that configuration B must be able to reach both final states. This is a contradiction because we said that A was the last state that was able to reach both final states. Therefore, there can only be a single final configuration. □

The natural next question, once we've shown that there's only a single final configuration is to show what the final configuration that must be reached is. That is our next major result. We take advantage of the fact that since there is only a single final configuration, we can choose the sequence of moves which makes it easiest to determine the resulting final configuration and know that this final configuration also applies to any other sequence of moves that we could have made.

Theorem 3.9. *A pile of $2n+1$ chips at the origin results in piles of size 1 on the interval $[-n, n]$.* [1]

Proof. We claim that there is a sequence of moves that can transform m consecutive piles of size 1 followed by a pile of n chips followed by another m consecutive piles of size 1 into m+1 consecutive piles of size 1 followed by a pile of n-2 chips followed by another m+1 consecutive piles of chips of size 1. Such a sequence of moves makes it clear that the resulting final configuration must be piles of size 1 on the interval $[-n, n]$ since the center pile is slowly reduced to one as the piles expand outwards. This is illustrated in the example below.

9
1 7 1
1 1 5 1 1
1 1 1 3 1 1 1
1 1 1 1 1 1 1 1 1

We now describe the sequence of moves that makes this possible. At each time step, we simply fire all of the piles that are able to be fired at the same time. Consider the patterns apparent in the following example we borrow from Anderson et al. [1]:

1 1 1 1 3 1 1 1 1
1 1 1 2 1 2 1 1 1
1 1 2 0 3 0 2 1 1
1 2 0 2 1 2 0 2 1
2 0 2 0 3 0 2 0 1
1 0 2 0 2 1 2 0 2 0 1
1 1 0 2 0 3 0 2 0 1 1
1 1 1 0 2 1 2 0 1 1 1
1 1 1 1 0 3 0 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1

An alternating sequence of 2s and 0s ripples outwards from the center until it reaches the most extreme piles. Then a single chip is pushed one spot further out in both directions, and another alternating sequence of 0s and 2s retreats towards the center. Once it reaches the center, we realize that we have reduced the center pile by 2 and added another pile of size 1 on both sides of it. This is exactly the sequence of moves that we sought to describe. Then, this sequence of moves can be applied over and over again until the configuration is fully reduced, each time firing the center pile, rippling the effect outwards, and then back inwards until only a single chip remains in the center. □

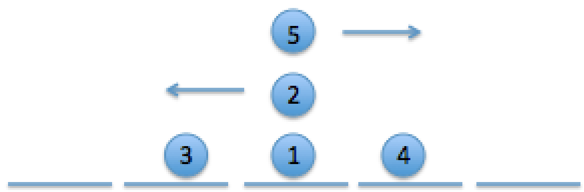

FIGURE 6. Distinct Chips: In the above example, first chips 3 and 4 were compared, and 3 is smaller so it was sent left while 4 is larger so it was sent right. Now, on the next turn, we are comparing 2 and 5. Since 2 is less than 5, 2 will be sent to the right and 5 will be sent to the left.

To extend this result to an even number of chips, note that there is a single chip that remains in the center for all time and never moves. If we instead began with only $2n$ chips at the origin, by the same argument we'd end up with chips on $[-n, -1] \cup [1, n]$ with a single empty pile at the origin.

Now that we've proved these results about this more general game, we return to Spencer's original game. As we stated early in the paper, there is a sequence of moves such that the more general game models Spencer's game. Since the more general game has the same final configuration regardless of the sequence of moves made, Spencer's game must also have this same unique final configuration.

4. Distinct Chips

In the previous chip firing games all of the chips were considered to be identical. As a result, when we considered configurations we were only concerned with the number of chips in each pile. Now we'd again like to consider the case where we again have chips that begin at the origin, but the chips themselves are distinct. We represent the chips by the integers 1 through $2n+1$. Now, on each firing, we choose two of the chips from the pile to compare and move the smaller of the two one pile to the left and the larger of the two one pile to the right. An example of this is shown in Figure 6. Other than this, the game is identical to the general version presented above, and as such all of the previous lemmas and theorems still apply to this case. Now, however, we are interested not only in the number of chips in each pile when we talk about configurations, but also what the chips are that exist in each pile.

Recall that the final configuration must be piles of size 1 on $[-n, n]$ for odd numbers of chips and piles of size 1 on $[-n, -1] \cup [1, n]$ for even numbers of chips. Of interest, however, is the final position of the chips relative to each other in the final configuration. We begin with a definition of one possible final configuration of the chips relative to each other.

Definition 4.1. Sorted
We say that a final configuration is sorted if for every pile the chips in piles to its left are less than the chip it contains and the chips in piles to its right are greater than the chip it contains.

We claim that for every odd number of chips there is both a sequence of moves such that the final configuration is sorted and a sequence of moves such the final configuration is not sorted. To show that there is a final configuration in which the chips are not sorted is slightly easier, so we begin with this result.

Lemma 4.2. *There exists a sequence of moves for every odd number of chips such that the final configuration is not sorted*

Proof. Consider the sequence of moves that we used above to prove the final configuration for an odd number of chips. Notice that there is a single chip, which ends up in the middle, that never moves from the center position. Since we choose which chips to compare in a given pile, if we never choose a given chip for a comparison, it will remain stationary for all time and therefore occupy the center location in the final configuration. If we choose this chip to be anything other than the middle element, the resulting configuration will not be sorted. □

While it is possible to adversarily make comparisons such that the resulting configuration is not sorted, the next lemma shows that it is also possible to advantageously make comparisons such that the resulting configuration does end up sorted.

Lemma 4.3. *There exists a sequence of moves for every odd number of chips such that the final configuration is sorted*

Proof. Each time you choose a pile to be fired, choose the smallest and largest elements that it contains to be the elements used in the comparison. This allows us to maintain the invariant that a chip never ends up to the left of something that is smaller than it. To see this, suppose that the invariant was broken and consider the firing that broke it, such that the invariant was true before the firing. For the invariant to be broken by the firing, the chip that is moved left must be smaller than something that is now to its right. However, the only chips to its right that were not previously to its right are the other chips in the pile that was fired. For the invariant to be broken, one of these elements must be smaller than the element that was moved left, but this is a contradiction because we chose the smallest element in the pile that was fired. If the invariant is always true, then in the final configuration, no element is to the left of something smaller than it, so the configuration is necessarily sorted. □

Conjecture 4.4. *Conjecture: For every sequences of moves for an even number of chips the final configuration is sorted.*

At this time, to the best of my knowledge this conjecture is still an open question. While tricky to prove, it is easy to explore through simulation by writing a short script.

Running repeated simulations re-enforces our conjecture and fails to provide an obvious counterexample. It is easy to see that the smallest chip must end up in the position farthest to the left and the largest chip must end up in the position farthest to the right since nothing can ever be fired to the extreme side of the smallest and largest chips. Interestingly, simulation shows that other than these two extreme chips, all of the interior chips appear to end up being compared to the other interior chips that ultimately occupy neighboring positions in the final configuration, regardless of the sequence of moves made. While this doesn't immediately imply that the final configuration must be sorted, as an observation it may provide inspiration as to why this is necessarily the case.

5. Final Remarks

Since chip firing has only been studied for about the past thirty years, it is still a very exciting area of research. The results originally proven on the infinite number line, as repeated above, have since been extended to both directed and undirected graphs. Chip-firing as a process is also interesting because of the multi-disciplinary implications it has. The process is important to theoretical physicists who are interested in its applications to frequency in noise and sandpile shapes and to theoretical computer science, where sorting, such as the process explored in the last section of this paper, is often of interest. Furthermore, this may have applications to latencies in networks and load balancing. It is likely in the coming years as researchers have even more time to explore this area even more exciting and novel results are produced.

6. Acknowledgments

I'd like to thank Professor Postnikov and my classmates in 18.204 (Undergraduate Seminar in Discrete Mathematics) for introducing me to chip-firing and the Abelian Sandpile Model, and for helping me to refine my presentation and writing skills. I hope to have more time in the future to further explore this area of research.

References

1. Anderson, Richard et al.. *Disks, Balls, and Walls: Analysis of a Combinatorial Game.* The American Mathematical Monthly 96.6 (1989): 481493. Web...
2. Anders Björner, László Lovász, Peter W. Shor, *Chip-firing Games on Graphs, European Journal of Combinatorics*, Volume 12, Issue 4, 1991, Pages 283-291, ISSN 0195-6698, http://dx.doi.org/10.1016/S0195-6698(13)80111-4. (http://www.sciencedirect.com/science/article/pii/S0195669813801114)
3. J. Spencer *Balancing vectors in the max norm* Combinatorica, 6 (1986), pp. 5566
4. Bak, Per and Tang, Chao and Wiesenfeld, Kurt *Self-organized criticality: An explanation of the 1/f noise, Phys. Rev. Lett.*, Volume 59, Issue 4, 1987, Pages 381-384, PhysRevLett.59.381. Americal Physical Society. 10.1103/PhysRevLett.59.381

Department of Mathematics, MIT, Cambridge , MA 02139