Course 18.312: Algebraic Combinatorics

Lecture Notes #29-31 Addendum by Gregg Musiker

April 24th - 29th, 2009

The following material can be found in several sources including Sections 14.9-14.13 of "Algebraic Graph Theory" by Chris Godsil and Gordon Royle, as well as the papers "Chip-firing Games on Graphs" by Anders Björner , Lásló Lovász, and Peter Shor (1991), and "Chip-firing and the Critical Group of a Graph" by Norman Biggs (1999).

1 Chip-Firing on an undirected graph

Throughout these notes, let $G = (V, E)$ be a finite connected undirected graph without loops. Assign a nonnegative integer $C(v)$ to each vertex v of G. We say that vertex v has $C(V)$ chips on it and call this assignment a **chip configuration**. Let N denote the total number of chips, i.e.

$$\sum_{v \in V(G)} C(v) = N.$$

Definition. If a vertex v has at least as many chips on it as its degree, i.e. $C(v) \geq deg(v)$, we say that v is **ready to fire**.

Definition. A vertex v that is ready to fire may send chips to its neighbors by sending one chip along each of its incident edges.

Definition. Given graph G, we say that a chip configuration C is **stable** if no vertex is ready to fire, i.e. if $C(v) < deg(v)$ for all vertices $v \in V(G)$.

Question. Given an initial chip configuration C on graph G, does C reduce to a stable chip-configuration, and if so, can we describe the possible stable states that are reachable? In particular, can the chip-firing process, i.e. the game, go on indefinitely or will it terminate in finite time?

Theorem. **[Björner , Lovász, and Shor (1991)]**

a) If $N > 2|E| - |V|$ then the game is infinite.

b) If $N < |E|$, then the configuration reduces to a unique stable configuration (that does not depend on the choice of which vertices are fired).

c) If $|E| \le N \le 2|E| - |V|$, then there exists some initial chip configuration which leads to an infinite process, and some other initial configuration which terminates in finite time (and terminates to a stable configuration which does not depend on the choice of the order vertices are fired).

Proof of Theorem. We start by showing that if $N \le 2|E| - |V|$, then there exists an initial configuration that is stable. Namely, if $N = 2|E| - |V|$, we can let $C(v) = deg(v) - 1$ for each vertex $v \in V(G)$. (Recall that $\sum_{v \in V(G)} deg(v) = 2|E|$.) If $N < 2|E| - |V|$, we simply take away chips from some of the vertices. Note that this shows the second half of (c).

On the other hand, if $N > 2|E| - |V|$ then even if we try to flatten the chip configuration as much as possible, there must be at least one vertex v' such that $C(v') \ge deg(v')$. Thus C is not stable and v' can fire. However, firing does not change the total number of chips, and hence this process will go on forever in this case and we have proven (a).

To see (b), we consider an acyclic orientation $\mathcal{O}$ on graph G, that is we orient the edges of G in such a way so that there are no cycles. This will always be possible. Let

$$T = \sum_{v \in V(G)} \max\Big(0, \quad C(v) - \text{outdeg}_{\mathcal{O}}(v)\Big)$$

where $\text{outdeg}_{\mathcal{O}}(v)$ is the number of outgoing edges from vertex v with respect to the chosen acyclic orientation.

Note that $\sum_{v \in V(G)} \text{outdeg}_{\mathcal{O}}(v) = |E|$ and since we are assuming that $N < |E|$, there exits at least one vertex $v \in V(G)$ such that $C(v) < \text{outdeg}_{\mathcal{O}}(v)$. We use this quantity T to show that the firing process must terminate: assume that C is an initial configuration with chip total $N < |E|$ and that vertex v_1 is ready to fire. (If there is no such v_1 to fire, then C is already stable and finite termination has been achieved.)

Fire v_1 and reverse the orientation of all edges leaving v_1. Call this new orientation $\mathcal{O}'$ and new chip configuration C'. This turns v_1 into a sink and cannot create a cycle. Moreover, the quantity T will never be increased by this change and will in fact decrease if there is a new vertex u with $C'(u) < \text{outdeg}_{\mathcal{O}'}(u)$. Since T is non-negative, this quantity can only decrease a finite number of times. However, it only takes a finite number of firings to create a new vertex u with a small enough $C'(u)$, so it follows that for such an initial C, a stable configuration must be reachable. We omit the proof that the stable configuration reached is *unique*. This shows (b) (except for uniqueness).

We use a similar technique to exhibit a chip configuration leading to an infinite game whenever $N \geq |E|$. (In fact we assume $N = |E|$ since adding chips can never enhance a configuration's stability.) In particular, let $\mathcal{O}$ be an acyclic orientation, and put $\text{outdeg}_{\mathcal{O}}(v)$ chips on each vertex. Call this configuration $C_{\mathcal{O}}$. Before proceeding further we need the following Lemma about such orientations.

Lemma. In an acyclic orientation on a finite graph, there is at least one vertex that is a sink and one that is a source.

Proof of Lemma. Suppose otherwise. Then each vertex v has at least one outgoing edge. Pick v_1 arbitrarily and follow one of the outgoing edges of v_1 and let v_2 denote the adjacent vertex. Then v_2 will also have at least one outgoing edge and so we continue this process. However, since there are a finite number of vertices, we must eventually revisit a vertex which would imply the orientation was not acyclic. Similar logic shows that there is a source.

Let v be the source of $\mathcal{O}$, and note that v can fire since $C(v) = deg(v)$. We reverse all edges incident to this source and then get a new acyclic orientation with a new source v'. We then can fire v' and continue indefinitely in this way. This proves the first part of (c) and finishes the proof.

Remark: While we did not give the details in this proof, part of the theorem is the fact that whether or not a configuration stabilizes or not and what configuration it stabilizes to does not depend on the choice of ordering of the firing vertices but only on the initial configuration.

2 The Dollar game: Chip-firing with a sink vertex

We now designate one of the vertices v_0 to be special, i.e. a **sink vertex**.

New Firing Rule for v_0: Vertex v_0 can fire if and only if the rest of the configuration C is stable.

Given a graph G (finite, undirected, loopless, connected) and $v_0 \in V(G)$ we say that a **chip configuration** C is an assignment of a nonnegative integer to each $v \in V(G) \setminus \{v_0\}$. We say that such a configuration is **stable** (v_0-stable if the choice of v_0 ambiguous) if $C(v) < deg(v)$ for each $v \in V(G) \setminus \{v_0\}$.

In particular, we do not think of v_0 having chips on it but instead it is a black hole that can absorb chips from the rest of the system.

Claim: If G is finite and connected, then for any initial chip configuration, there is a sequence of legal firings that results in a stable configuration.

Sketch of Proof: We prove this inductively on the radius of G with respect to v_0, i.e. the size of the longest minimal path between v_0 and any other vertex $v \in V(G)$. When neighbors of v_0 fire, the total number of chips in the system decreases and eventually is small enough so that the configuration stabilizes. When neighbors of v_0's neighbors fire, eventually one of v_0's neighbors has enough chips on it so that it is ready to fire.

Definition. A configuration C is **recurrent** if there is a legal firing sequence (including v_0 firing) such that C recurs.

Definition. A configuration C is **critical** if it is both stable and recurrent.

Theorem. **(Uniqueness)** Given an initial configuration C, there exists a **unique** critical configuration that is reachable via a legal firing sequence.

Proof. The proof of this Theorem is deferred to the next section.

We use this theorem to give a group structure to the set of critical configurations.

Definition. We let $K(G, v_0)$ denote the **critical group** of graph G with respect to the sink vertex v_0. This group is given as the set of critical configurations of G (with respect to v_0) with the group law given by

$$C_1 \oplus C_2 := \overline{C_1 + C_2}.$$

Here $C_1 + C_2$ is pointwise addition of chip assignments to vertices and $\overline{C}$ denotes the unique critical configuration reachable from C.

Theorem. We have the isomorphism of groups

$$K(G, v_0) \cong \mathbb{Z}^{|V(G)|-1} \Big/ \text{Im } L_0(G, v_0)$$

where $L_0(G, v_0)$ denotes the reduced Laplacian matrix of graph G with the row and column corresponding to vertex v_0 deleted. Thus $K(G, v_0)$ is an abelian group.

Proof. Sketched in class, follows from the fact that firing a vertex corresponds to subtracting the corresponding column of the reduced Laplacian matrix.

Corollary 1. The number of critical configurations of undirected connected finite graph G with respect to v_0 is independent of v_0, and equals the number of spanning trees of G.

Corollary 2. We can use Smith Normal Form of $L_0(G, v_0)$ to compute the decomposition of $K(G, v_0)$ into a product of cyclic subgroups.

Definition. We say that two matrices M and M' have the same **Smith normal form** if we can change M to M' by the following three transformations:

(1) Add an integer multiple of one row or column to another.

(2) Multiply a row or column by (-1).

(3) Interchange two rows or two columns.

Theorem. (1) If M and M' are integer matrices with k (resp. k' columns), the groups $\mathbb{Z}^k \big/ Im\ M$ and $\mathbb{Z}^{k'} \big/ Im\ M'$ are isomorphic if and only if M and M' have the same Smith Normal Form.

(2) Any integer matrix M has the same Smith normal form as a "diagonal" rectangular matrix D with d_i's on the diagonal such that $d_i | d_j$ for any $i \leq j$. In particular,

$$\mathbb{Z}^k / Im\ D \cong \mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z}.$$

3 Proof of the Uniqueness Theorem

Before proving the theorem showing uniqueness of a critical configuration, we introduce some notation. Given an undirected graph G, we let $D = D(G)$ denote a digraph on the same set of vertices and edges as G, only with an orientation arbitrarily chosen for each edge. Note that the choice of orientation will not be relevant and that we are *not* replacing each undirected edge with a pair of directed edges.

Definition. Let G, and hence D, have n vertices and m edges. We let $\partial = \partial(D)$ denote the **incidence matrix** of D, which is the $n \times m$ matrix with (i,j)th entry given as

$$\partial_{ij} = \begin{cases} -1 & \text{if vertex } v_i \text{ is the tail of edge } e_j \\ 1 & \text{if vertex } v_i \text{ is the head of edge } e_j \\ 0 & \text{otherwise} \end{cases}.$$

Lemma. The Laplacian matrix is given as the product

$$L(G) = L(D) = \partial(D) \cdot (\partial(D))^T.$$

Proof. For any $n \times m$ matrix A, it follows that

$$(AA^T)_{ij} = \sum_{e=1}^{m} a_{ie} a^T_{ej} = \sum_{e=1}^{m} a_{ie} a_{je}.$$

Since edge e is of the form $v_i \longrightarrow v_j$ or $v_i \longleftarrow v_j$, and both contribute a (-1) to this sum and any other edge contributes a zero, it follows that $\partial\partial^T_{ij} = -m_{ij}(G)$ whenever $i \neq j$. (We remind the reader that m_{ij} denotes the number of undirected edges between v_i and v_j in G.) By similar logic, we see that $\partial\partial^T_{ii} = \sum_{e=1}^{m} \partial^2_{ie} = \deg(v_i)$, thus completing the proof.

We can think of matrix $L(G)$ as a linear transformation from $\mathbb{Z}^n \to \mathbb{Z}^n$.

Lemma. If G is a finite connected graph, the kernel of $L(G)$ is spanned by the all ones vector $[1, 1, \ldots, 1]^T$.

Proof. Let $w \in \mathbb{Z}^n$. We use the fact that $L(G) = \partial\partial^T$ to obtain

$$L(G)w = 0 \longleftrightarrow \partial\partial^T w = 0.$$

Furthermore, ∂^T is a linear transformation from $\mathbb{Z}^n \to \mathbb{Z}^m$ which takes in a vector $[w_1, w_2, \ldots, w_n]^T \in \mathbb{Z}^{|V|}$ and assigns a value $w_i - w_j$ to each edge e from v_j to v_i. Hence $\partial^T w = 0$ if and only if all of the potential differences are zero. Since G is connected, this means the *potential*, i.e. the value C_i at each vertex must be constant throughout the graph.

Furthermore, we let Δ_e denote the potential difference along edge e and then obtain

$$w^T L(G) w = w^T \partial\partial^T w = (\partial^T w)^T (\partial^T w) = \sum_{e=1}^{m} \Delta_e^2.$$

As a consequence, $L(G)w = 0$ implies that $\partial^T w = 0$, and the converse is more easily seen to be true. This completes the proof that the kernel of $L(G)$ is spanned by the all ones vector.

Given an integer vector $w \in \mathbb{Z}^{|V|}$ and a choice of sink vertex v_0, we can subtract the (v_0)th column of the Laplacian matrix a sufficient number of times until the integer assigned to each other $v \in V(G) \setminus \{v_0\}$ is nonnegative. In this way, we can turn any such integer vector into a chip configuration C on $V(G) \setminus \{v_0\}$. We thus wish to show that there exists a unique critical configuration in the same coset (i.e. same orbit under action of adding columns of Laplacian matrix) as C. We now introduce some useful notation.

Starting with an initial chip configuration C, we let $C^{(v_{i_1}, v_{i_2}, \dots, v_{i_\ell})}$ denote the configuration obtained after the vertices v_{i_1} through v_{i_ℓ} fire. We say that $\mathcal{X} = (v_{i_1}, v_{i_2}, \dots, v_{i_\ell}) \in V(G)^\ell$ is a **legal firing sequence** if $\mathcal{X}$ does not contain v_0 and no vertex goes into deficit along the way, that is for every j, the number of chips on v_{i_j} in configuration $C^{(v_{i_1}, v_{i_2}, \dots, v_{i_{j-1}})}$ is greater than or equal to the degree of v_{i_j}.

We define the **score** of $\mathcal{X}$ to be $\mathbf{x} = (x_1, \dots, x_{|V|-1})$ if vertex v_i appears x_i times in $\mathcal{X}$.

If $\mathcal{Y}$ is another legal firing sequence (also not containing v_0), with score given by vector $\mathbf{y}$, then we let $\mathcal{X} \setminus \mathcal{Y}$ denote the sequence obtained from $\mathcal{X}$ by deleting the first y_i occurrences of vertex v_i. (If $x_i \leq y_i$ then we delete all occurrences.) Lastly, we denote the concatenation of sequence $\mathcal{X}$ followed by $\mathcal{Y}$ as $(\mathcal{X}, \mathcal{Y})$.

Proposition. Let C be an initial unstable configuration, and $\mathcal{X}$, $\mathcal{Y}$ both be legal firing sequences not containing v_0. Then $(\mathcal{Y}, \mathcal{X} \setminus \mathcal{Y})$ is also a legal firing sequence and has score equal to z where $z_i = \max\{x_i, y_i\}$ for all $i \in \{1, 2, \dots, |V| - 1\}$.

We prove this Proposition momentarily. As an application we show that if $\mathcal{X}$ and $\mathcal{Y}$ are both legal firing sequences for an initial configuration C, then there is a configuration, $C^{\mathcal{Z}}$, which is reachable from both $C^{\mathcal{X}}$ and $C^{\mathcal{Y}}$.

Namely, the concatenations $(\mathcal{Y}, \mathcal{X} \setminus \mathcal{Y})$ and $(\mathcal{X}, \mathcal{Y} \setminus \mathcal{X})$ are both legal by the Lemma, and have the same score. Thus letting $\mathcal{Z} = (\mathcal{Y}, \mathcal{X} \setminus \mathcal{Y})$, we find $C^{\mathcal{Z}}$ satisfies the desired hypothesis.

In particular, if $C^{\mathcal{X}}$ and $C^{\mathcal{Y}}$ are both stable configurations, there are legal firing sequences $\mathcal{X}'$ and $\mathcal{Y}'$ not containing v_0 so that $C^{(\mathcal{X}, \mathcal{X}')} = C^{(\mathcal{Y}, \mathcal{Y}')}$. However, since the only legal firing sequence not containing v_0 following a stable configuration is the empty sequence, we must have that $C^{\mathcal{X}} = C^{\mathcal{Y}}$. In particular, we have shown that there is a unique stable configuration reachable from any unstable initial configuration C. Further, it is well-defined when one can let vertex v_0 fire, i.e. if and only if no other vertex can fire.

Now we must show that there is at most one stable configuration that recurs. Let C be a recurrent configuration. That means there exists a nonzero vector $\mathbf{x}$, the score of a firing sequence $\mathcal{X}$, such that $C = C - L(G)\mathbf{x}$. Since we know that the kernel of $L(G)$ is spanned by the all ones vector, we know that such an $\mathbf{x}$ must be a positive integral multiple of the all ones vector. However, we can in fact say more strongly:

Lemma. Given stable recurrent C, i.e. $C = C - L(G)\mathbf{x}$ where $\mathbf{x}$ is the score of a legal firing sequence, then $\mathbf{x}$ may be the all ones vector.

Proof. Since C is stable, the firing sequence must start with v_0 firing. Assume that there is a legal firing sequence (from $C^{(v_0)}$) where one of the vertices appears more than once. Without loss of generality, assume that v_1 is the first vertex which appears twice in the sequence, and call the configuration $\hat{C}$ right before v_1 fires for the second time. Since C was stable, $C_1 \leq \deg(v_1) - 1$. When v_0 fires, this may increase the number of chips on v_1 (if v_0 is a neighbor of v_1), and v_1's neighbors may subsequently fire, but they each fire at most once. Since we have assumed that v_1 has already fired a first time, we have that $\hat{C}_1 = C_1 - \deg(v_1) + d_{1,v_0} + \sum_{v_i \text{ is a neighbor of } v_1 \text{ which has fired}} d_{1,i}$. However, this sum (including v_0 as a possible neighbor) is less than or equal to $\deg(v_1)$, and consequently $\hat{C}_1 \leq C_1 \leq \deg(v_1) - 1$. Thus v_1 cannot fire a second time, and we have a contradiction.

As a consequence, any legal firing sequence contains each vertex at most once. Another way of saying this is that vertex v_0 must fire a second time before any other vertex does. If v_0 and all other vertices have fired, we are done since we have the all ones vector and we clearly have returned to C. On the other hand, if we obtain a stable configuration earlier, once v_0 fires a second time, any vertex that has not already fired will never be able to catch up to v_0. But we should eventually get to a multiple of the all ones vector, so we must have been in the first case.

Note that given stable recurrent configuration C we let v_0 fire to obtain $C^{(v_0)}$ and we have shown that the subsequent firing of $(v_{i_1}, \dots, v_{i_{n-1}})$, $(i_1, i_2, \dots, i_{n-1})$ is a permutation of $\{1, 2, \dots, n-1\}$ that exactly undoes the firing of v_0. This means that C is the first stable configuration reached after reducing. Thus the first recurrent stable configuration reached will subsequently recur from then on as the only stable configuration ever reached again. Thus, it is the *only* critical configuration.

Proof of Proposition. We now end the proof of the Theorem by including the proof of the remaining Proposition, which we prove by induction on the length of $\mathcal{X} \setminus \mathcal{Y}$. Without loss of generality, assume that $\mathcal{X} \setminus \mathcal{Y}$ starts with v_1. Let $\overline{\mathcal{X}}$ denote the truncated sequence which ends right before the $(y_1 + 1)$st appearance of v_1 in $\mathcal{X}$. Clearly $\overline{\mathcal{X}}$ is legal as well, and has score $\overline{x}$ satisfying

$$\overline{x_1} = y_1, \quad \overline{x_i} \leq y_i \quad \text{for } 1 < i \leq n.$$

Thus the coordinate

$$\begin{aligned}
(C^{\mathcal{Y}})_1 &= C_1 - d_1 y_1 + \sum_{i=2}^{n} d_{i,1} y_i \geq \\
(C^{\overline{\mathcal{X}}})_1 &= C_1 - d_1 \overline{x_1} + \sum_{i=2}^{n} d_{i,1} \overline{x_i} \\
&\geq \deg(v_1)
\end{aligned}$$

since vertex v_1 is about to be fired again in sequence $\mathcal{X}$. Thus v_1 can be fired again after $\mathcal{Y}$ as well and so $(\mathcal{Y}, v_1)$ is legal as well. We now repeat the argument for $\mathcal{X}$ and $(\mathcal{Y}, v_1)$ until we have shown the entire sequence $(\mathcal{Y}, \mathcal{X} \setminus \mathcal{Y})$ is legal.

4 Chip-Firing on Directed Graphs

For more on chip-firing, including chip-firing on directed graphs, we invite the reader to look at the papers "Chip-Firing Games on Directed Graphs" by Anders Bj orner and Lásló Lovász (1992) as well as "Chip-Firing and Rotor-Routing on Directed Graphs" by Alexander Holroyd, Lionel Levine, Karola Mézáros, Yuval Peres, James Propp, and Daivd Wilson (2008).

Much of the theory that we have developed for undirected graphs G also works for digraphs D but not always. In this section, we give a glimpse of some of the aspects of chip-firing on digraphs.

First of all, the notion of sending chips to ones' neighbors is easily adapted in this case. We use the directed Laplacian matrix whose diagonal entries are the outdegrees and the off-diagonal entries are the number of *directed* edges between a given pair of vertices. As in the undirected case, we assume our graph is loop-less and finite. Recall that we may turn any undirected graph into a digraph by replacing each edge with a pair of oppositely directed edges.

Since $L(D)$ is not necessarily symmetric, it does not easily decompose into a product AA^T like in the previous section. Note that the definition of the incidence matrix also is now problematic since our edges already have orientations. However there are two hypotheses for which most of the undirected theory goes through.

Definition. We say that a digraph has a **global sink** s if all vertices $v \in V(D) \setminus \{s\}$ have an edge $v \longrightarrow s$ incident to it and s contains no outgoing edges.

Lemma. If digraph D has a global sink, then every chip configuration on D stabilizes.

Proof. Omitted, see Lemma 2.4 of [Holroyd-Levine-Mézáros-Peres-Propp-Wilson].

In fact, there are several other results of this flavor where the presence of a global sink ensures that the chip-firing process runs similar to the undirected case. The other common hypothesis we have already seen in our class.

Lemma. (Burning algorithm due to Dhar (1990)) If a digraph D is **Eulerian** (i.e. connected and balanced) with a sink vertex v_0 then a chip configuration C is recurrent if and only if the stabilization $\overline{C + \beta} = C$ where

$$\beta(v) = \text{outdedg}(v) - \text{indeg}(v) \geq 0.$$

Furthermore, if C is recurrent, then each vertex fires exactly once during the stabilization of $C + \beta$.

Remark. Here by stabilization we would first fire sink v_0 if $C + \beta$ were already stable and stop at the next stable configuration reached.