Simplicial Chip Firing

Alex St Laurent

Summer 2012

1 Introduction

There is a wide literature surrounding the theory of chip-firing on graphs and the related critical group. The literature includes equivalent models — *dollar game*, *abelian sandpile model*, etc. Recently, the work of Duval, Klivans, and Martin [1] extended the theory of these critical groups from graphs to simplicial complexes. This opened the door to an analogous model of discrete flow on these structures.

Throughout the Summer of 2012, a program was built in Java to aid in the visualization and experimentation of this discrete flow in one or two dimensions (the former case being ordinary chip-firing on a graph). With the aid of the tool, conjectures were explored that paved the way for continued research in the Fall 2012 semester.

2 The Graphical Case

Let $G = (V, E)$ be a finite, connected graph devoid of multi-edges, loop edges, or directed edges. Distinguish one vertex q to be the *bank*. A *configuration* of chips on G is a map $c : V \to \mathbb{Z}$, regarded as a column vector. A configuration is *conservative* if the entries sum to zero. In particular, the configurations are typically standardized further by requiring $c(q) \leq 0$ and $c(v) \geq 0$ for $v \neq q$. This comes without a loss of generality.

A vertex $v \neq q$ can *fire* if it has at least as many chips as its neighbors. That is, v can fire if $c(v) \geq \deg(v)$. When a vertex fires, it sends one chip to each of its neighbors, thus losing $\deg(v)$ chips.

This behavior can be defined more concretely with the notion of the Laplacian. A *Laplacian matrix* L is a $|V|$-by-$|V|$ matrix whose diagonal entries contain the degrees of the vertices and whose off-diagonals are -1 to indicate adjacency or 0 to indicate non-adjacency.

The entry L_{rc}, referring to vertices v_r and v_c, is equal to:

$$L_{rc} = \begin{cases} \deg(v_r) & r = c \\ -1 & \{v_r, v_c\} \in E \\ 0 & \{v_r, v_c\} \notin E \end{cases}$$

Note $L_{rc} = L_{cr}$, and thus L is symmetric.

Consider a sequence of vertices to be (legally) fired, denoted $\mathcal{X}$. The *representative vector* of $\mathcal{X}$ is x if $x(v)$ is equal to the number of times v is fired in $\mathcal{X}$. If the result of starting with configuration c, then firing vertices as prescribed by $\mathcal{X}$, is the configuration c', then:

$$c' = c - Lx$$

That is, each firing of v corresponds to subtracting v's respective column (or row) in L.

If no vertex $v \neq q$ can fire, it is then and only then can the bank q fire. It serves as a way to get the system "unstuck". This leads to the idea of stability: a configuration of chips if *stable* if only the bank can fire.

We would like to think stable configurations are important in some way. Well, they can be, if they possess another property: a configuration is *recurrent* if there exists a legal sequence of firings that has no net effect on the configuration. That is, the sequence brings the configuration back to where it started.

A configuration is *critical* if it is both stable and recurrent. In a sense, critical configurations are inevitable final resting places of the firing process.

An important theorem surrounding critical configurations states that, given any conservative configuration on a graph, a unique critical configuration can be reached through a sequence of legal firings. Furthermore, the equivalence classes of configurations that are all equivalent to the same critical configuration form a group structure, with the critical configurations themselves serving as appropriate representatives (much like how the integers $0 \ldots n-1$ serve as representatives in the group of integers modulo n). The *critical group* of G is denoted $K(G)$, and its order is equal to the number of spanning trees of the graph. The group operation consists of element-wise addition of configurations (and then, optionally, determining the critical configuration to which the sum is equivalent by firing).

Let σ be a 1-by-$|V|$ row-matrix consisting of all-ones, so that its kernel consists of conservative configurations (the motivation of this construction will be apparent when discussing the extension to simplicial complexes). Because the critical group is this kernel, partitioned such that two configurations are in the same class if firing can be applied to one to reach the other, we can give an algebraic definition to the critical group:

$$K(G) = \ker \sigma / \operatorname{im} L$$

Additionally, let us define the *reduced Laplacian* $\tilde{L}$ as being the Laplacian with the row/column referring to q removed. A fundamental group isomorphism is as follows:

$$K(G) \cong \mathbb{Z}^{|V|-1} / \tilde{L}$$

With this, we can ignore the idea of conservativity and the amount of chips on q.

3 The Critical Group on Simplicial Complexes

An *abstract simplicial complex*, denoted Δ, is a subset of the power set of $\{0 \ldots n\}$ which is closed under the operation of taking subsets. That is, if a particular subset is an element of Δ, then so are all of its subsets. A *simplicial complex* is Δ together with a geometric embedding into space,

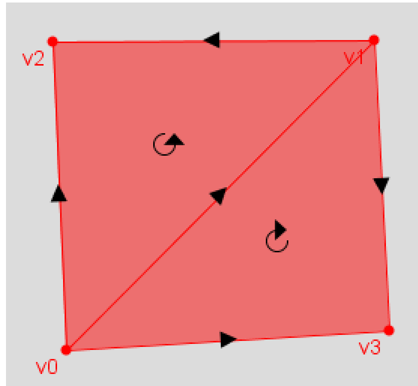

Figure 1: An oriented two-dimensional complex.

which gives rise to the thinking of the individual integers within the subsets as *vertices*, because subsets of size one are represented as points. The *dimension* of a simplicial complex is one less than the size of its largest subset, which is also the Euclidean spatial dimension of the geometric representation of this subset, i.e. the *simplex* (we will also use the term "simplex" to refer to the abstract subset itself). Note that a graph is a one-dimensional simplicial complex, with its edges representing subsets containing two elements. The simplices are triangles and tetrahedrons in two and three dimensions, respectively.

Any simplex can have an arbitrary *orientation*, which is important for defining certain matrices on a complex. In simple terms, an orientation of a simplex is an ordering of its vertices, with two orientations being equivalent if they represent the same geometric direction on the simplex. For example, the edge $\{0, 1\}$ can either be oriented $\langle 0, 1\rangle$ or $\langle 1, 0\rangle$. Geometrically, this is interpreted as annotating the edge with an arrow. The triangle $\{0, 1, 2\}$ has two possible orientations, $\langle 0, 1, 2\rangle$ or $\langle 0, 2, 1\rangle$. Any other ordering of the vertices are equivalent to one of these two — they can be thought of as "clockwise" and "counterclockwise", and are geometrically represented by a curling arrow.

Two simplices x and y of dimension $i-1$ and i, in which x is a subset of y, are *consistently oriented* if the orientation on x matches the orientation on the subset of y containing the same elements of x. Geometrically, if they are consistently oriented, the arrows will being going in the same direction. They are *inconsistently oriented* if the opposite is true. In Figure 1, the triangle $\{0, 1, 2\}$ and the edge $\{1, 2\}$ are consistently oriented, while the triangle and edge $\{0, 2\}$ are inconsistently oriented.

For the edge/vertex case, one must set a standard as to whether consistency of orientation occurs when the edge "leaves" the vertex, or when the edge "arrives" at the vertex. However, this won't affect the results of the firing process.

Now we can build up to the analogue of the graphical Laplacian. Let Δ_i denote the set of i-dimensional simplices on Δ. We shall define $\partial_i : \mathbb{Z}^{|\Delta_i|} \to \mathbb{Z}^{|\Delta_{i-1}|}$ as follows, using the row indices to refer to $(i-1)$-dimensional simplices, and the column indices to refer to i-dimensional simplices:

$$(\partial_i)_{rc} = \begin{cases} 1 & r \text{ and } c \text{ are consistently oriented} \\ -1 & r \text{ and } c \text{ are inconsistently oriented} \\ 0 & \text{neither is a subset of the other} \end{cases}$$

In the somewhat degenerate case of ∂_0, we use the technicality that Δ_{-1} is simply the empty set, with which all vertices are consistently oriented. Thus, we get a 1-by-$|\Delta_0|$ row matrix (exactly the matrix σ defined earlier in the graphical case).

With this, we can define the *i-dimensional combinatorial Laplacian* $L_i : \mathbb{Z}^{|\Delta_i|} \to \mathbb{Z}^{|\Delta_i|}$ of Δ as:

$$L_i = \partial_{i+1}\partial_{i+1}{}^T$$

In the graphical case, we find that L_0 is the same Laplacian defined earlier (one can verify that the two definitions will be equivalent).

We can now define the i-dimensional critical group of Δ:

$$K_i(\Delta) = \ker \partial_i / \operatorname{im} L_i$$

Note that in the graphical case, $K(G) = K_0(G)$.

In order to define the idea of the reduced Laplacian, we must introduce the notion of a *simplicial spanning tree*. In general, the definition rests on various topological constructs that we have little motivation of getting into. We can say that 0-dimensional spanning trees are individual vertices (they "span" the empty set) and 1-dimensional spanning trees are the usual graphical spanning trees. For more on this topic, see [1].

Given a proper i-dimensional spanning tree Γ of a d-dimensional complex (with $i < d$), the reduced Laplacian $\tilde{L_i}$ is L_i with the rows/columns referring to Γ removed. This gives rise to the following group isomorphism:

$$K_i(\Delta) \cong \mathbb{Z}^{|\Delta_i|-|\Gamma|} / \operatorname{im} \tilde{L_i}$$

The order of this group is related to a weighted enumeration of the i-dimensional spanning trees.

4 Discrete Flow on Simplicial Complexes

As defined by the critical group, we have *conservative configurations* (i.e. elements of $\ker \partial_i$) modulo the equivalence given by the combinatorial Laplacian L_i. When $i = 1$, given some configuration c, one interpretation of $c(e)$ (on some edge e) is one of "flow" in the direction the orientation of e. In higher dimensions, this interpretation holds with generalized *i-flow*; for example, 2-flow circulates about a triangle.

The condition $c \in \ker \partial_1$ means that the net flow about each vertex is zero. In general, the conservativity of a configuration means that the all i-flows about an $(i-1)$-dimensional simplex must cancel out about that simplex. In the graphical case, all the flows on the vertices (which don't have any inherit "direction") cancel out around the empty set.

By the definition of the critical group, two configurations are equivalent if subtracting columns of the Laplacian can take you from one to the other. This is analagous to the standard chip-firing model, in which the firing of a vertex corresponds to subtracting a row.

In higher dimensions, we can similarly subtract rows of the combinatorial Laplacian. A simplex, when fired, will have its flow decreased by the amount $(i+1)$-dimensional simplices of which it is a subset (which is the degree in the graphical case).

Given the group isomorphism, there is some inherent notion of being able to ignore the spanning tree when dealing with group properties or firing. In the graphical case, this is apparent given the distinguished aspect of the source vertex.

The open problem is to develop the "chip-firing" theory given these critical group constructs, and/or to identify a general method to determine "critical configurations", i.e. a set of stable and recurrent equivalence class representatives. One immediate obstacle is that, while the degree-based stability worked well because firing a vertex caused its neighbors to gain chips, in higher-dimensional firing, neighbors can actually *lose* flow. Additionally, while the spanning tree Γ in a graph has a simple, positive net effect, the spanning tree in higher dimensions might not be so simple. Firing the entire substructure at once may have a mixed effect or even no net effect.

There was time spent exploring various ideas for stability rules on various 2-dimensional complexes. Additionally, using various trees to reduce the Laplacian and various subsets of these trees to act as the bank (the portion that fires when no other edges can). It was found that, if a tree could be selected such that the edges not in the tree had a positive effect on one another when firing, and if the bank also had a positive effect on these edges, then the traditional degree rule worked in *most* cases attempted. They "worked" in the sense that every configuration was only equivalent to one critical configuration.

When continuing research in the fall, complexes such as these will be a point of study. Additionally, time will be devoted to looking at the $\ker L_i$, that is, the set of representative vectors of firing sequences that have no net effect on the configuration. This space will be important to characterize if we are to properly examine recurrent configurations.

5 The Program

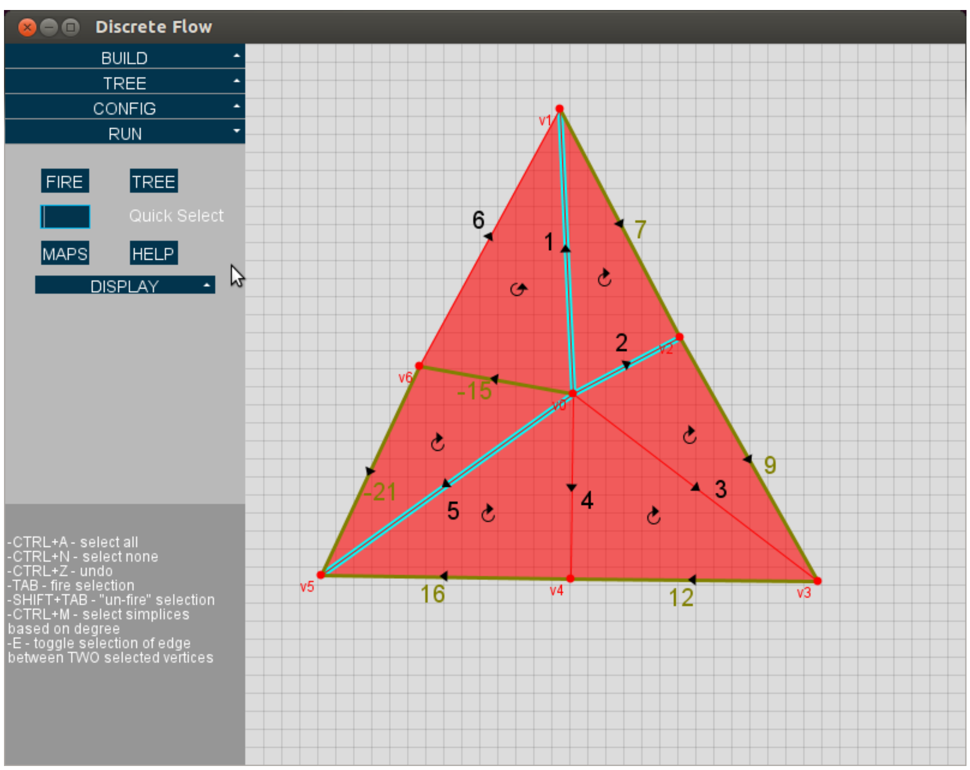

Figure 2: A screenshot of "Discrete Flow"

As mentioned, time was devoted to building a Java application to aid experimenting with complexes and firing in one and two dimensions. In Figure 2, we can see three edges highlighted in light blue, ready to be fired. Besides automating matrix construction and row calculations, a useful aspect of the program is that the user can fire entire substructures at once and immediately observe their net effects; whether they are the entire spanning tree, a particular sub-tree, or an arbitrary collection of edges, the program will aggregate the effects.

The program, called "Discrete Flow" will be available on Professor Klivan's website as a Windows executable, a Mac OS application, and a runnable Java jar file.

6 Peripheral Results

When attempting to characterize matrices/configurations and prove conjectures (that turned out to be not true in general), some other results were stumbled upon that maybe be relevant in the future. Their proofs are presented here.

Note: the symbols v, e, and τ are used to represent $(i-2)$-dimensional, $(i-1)$-dimensional, and i-dimensional simplices, respectively. While these intuitively match English geometric descriptors in the $i = 2$ case ("vertex", "edge", "triangle"), by no means are any results restricted to this case. The choice of symbols is only meant to aid visualization by aiding the reader to imagine a two-dimensional complex.

Kernels and Inner Products

On an i-dimensional complex, define vectors h and t on all $(i-1)$-dimensional simplices e, such that $h(e)$ = the number of i-faces in which e is consistently oriented (a "head"), and $t(e)$ = the number of i-faces in which e is inconsistently oriented (a "tail").

Given the 1-dimensional edges in lexicographic order $\langle 01, 02, 03, 12, 13\rangle$, in the above example, $h = (2, 0, 0, 1, 1)^T$, and $t = (0, 1, 1, 0, 0)^T$.

Additionally, for any i-dimensional simplex τ, let $H(\tau)$ represent the set of consistently oriented $(i-1)$-faces of τ (the "head" set), and let $T(\tau)$ represent the set of inconsistently oriented $(i-1)$-faces of τ (the "tail" set). For example, in Figure 1, $H(012) = \{01, 13\}$.

Let H be the multiset of all $H(\tau)$, and T be such for all $T(\tau)$. That is, H consists of all $(i-1)$-simplices e with multiplicity $h(e)$, and T consists of them with multiplicity $t(e)$. In Figure 1, $H = \{01, 01, 12, 13\}$, and $T = \{02, 03\}$.

Proposition 1. $\alpha \in \ker \partial_i^T \Leftrightarrow \langle \alpha, h\rangle = \langle \alpha, t\rangle$

Proof. $(\Rightarrow)$ We can regard the entry in ${\partial_i}^T\alpha$ corresponding to simplex τ as:

$$[{\partial_i}^T\alpha](\tau) = \sum_{e\in H(\tau)} \alpha(e) - \sum_{e\in T(\tau)} \alpha(e)$$

Because $\alpha \in \ker \partial_i^T$,the right-hand side of the above equality is zero. Because this equality true for each i-simplex τ, we can extend over the sum of all τ:

$$\sum_{e\in H(\tau)} \alpha(e) - \sum_{e\in T(\tau)} \alpha(e) = 0\,\forall\tau$$

$$\sum_{e \in H} \alpha(e) = \sum_{e \in T} \alpha(e)$$

Whenever e is a head, its corresponding $\alpha(e)$ contributes to the left-hand side of this new equality. Meaning, the summation over H is equal to the inner product $\langle \alpha, h \rangle$. Similarly, the right-hand side is equal to $\langle \alpha, t \rangle$. Thus, we have $\langle \alpha, h \rangle = \langle \alpha, t \rangle$.

($\Leftarrow$) Conversely, if we have $\langle \alpha, h \rangle = \langle \alpha, t \rangle$ for any vector α, the above equalities of summation can be derived through simple replacement, which implies $\alpha \in \ker \partial_i^T$.

□

Notice how H and T had a higher-dimensional perspective on the $(i-1)$-simplices. Now we would like to develop the analogous lower-dimensional perspective.

For starters, define h' and t' as the analogies to h and t. That is, $h'(e) =$ the number of $(i-2)$-simplices with which e is consistently oriented, and t' is similar for inconsistent orientations.

Denote the set of $(i-1)$-simplices with consistent orientation with respect to the $(i-2)$-simplex v as $[v]_H$ (e.g. the set of incoming edges to the vertex v in the graphical case), and those with inconsistent orientation as $[v]_T$. Note: this means $\forall v, e : v \in H(e) \Leftrightarrow e \in [v]_H, v \in T(e) \Leftrightarrow e \in [v]_T$.

In the same way H and T were the multisets encapsulating all $H(\tau)$ and $T(\tau)$, we shall have H' and T' refer be the multisets encapsulating all v_H and v_T. That is, H' consists of all $(i-1)$-simplices e with multiplicity $h'(e)$, and T' consists of them with multiplicity $t'(e)$.

Proposition 2. $c \in \ker \partial_{(i-1)} \Leftrightarrow \langle c, h' \rangle = \langle c, t' \rangle$

Proof. ($\Rightarrow$) We can regard the entry in $\partial_{(i-1)}c$ corresponding the simplex v as:

$$[\partial_{(i-1)}c](v) = \sum_{e \in [v]_H} c(e) - \sum_{e \in [v]_T} c(e)$$

Because $c \in \ker \partial_{(i-1)}$, the right-hand side of the above equality is zero. Because this is true for each $(i-2)$-simplex v, we can extend over the sum of all v:

$$\sum_{e \in [v]_H} c(e) - \sum_{e \in [v]_T} c(e) = 0 \; \forall v$$

$$\sum_{e \in H'} c(e) = \sum_{e \in T'} c(e)$$

Wherever e is consistently oriented with some $(i-2)$ simplex, its corresponding $c(e)$ contributes to the left-hand side of this new equality. Meaning, the summation over H' is equal to $\langle c, h' \rangle$. Similarly, the right-hand side is equal to $\langle c, t' \rangle$. Thus, we have $\langle c, h' \rangle = \langle c, t' \rangle$.

($\Leftrightarrow$) Conversely, if we have $\langle c, h' \rangle = \langle c, t' \rangle$, for any vector c, the above equalities of summation can be derived through simple replacement, which implies $c \in \ker \partial_{(i-1)}$.

□

Differences

Proposition 3. $(h - t) \in \ker \partial_{(i-1)}$

Proof. Let us denote the entries in $\partial_{(i-1)}h$ and $\partial_{(i-1)}t$ corresponding to the simplex v as $[\partial_{(i-1)}h](v)$ and $[\partial_{(i-1)}t](v)$, respectively.

In $[\partial_{(i-1)}h](v)$ each $(i-1)$-simplex $e \in [v]_H$ makes a positive contribution equal to the number of i-simplices in which it appears as a head. Similarly, each $(i-1)$-simplex $e \in [v]_T$ makes the equivalent *negative* contribution (i.e. the opposite of the number of i-simplices in which it appears as a head).

Analogously, in $[\partial_{(i-1)}t](v)$, each $(i-1)$-simplex $e \in [v]_H$ makes a positive contribution equal to the number of i- simplexes in which it is unfavorably oriented (a tail). Similarly, each $(i-1)$-simplex $e \in [v]_T$ makes the equivalent *negative* contribution (i.e. the opposite of the number of i-simplices in which it appears as a tail).

Now, let us consider the difference:

$$[\partial_{(i-1)}h](v) - [\partial_{(i-1)}t](v)$$

Instead of considering the contributions from each $(i-1)$-dimensional simplex e, and summing across those, we shall consider the contributions from each i-dimensional simplex τ, and take that sum. Namely, for each τ, we will consider the contribution of its faces e as the total contribution that the simplex τ makes to the difference. Then, we shall sum across the contributions from all τ.

For each i-simplex τ, we have the following contributions:

- $+1$ $\forall e \in H(\tau)$ s.t $e \in [v]_H$
- -1 $\forall e \in H(\tau)$ s.t. $e \in [v]_T$
- -1 $\forall e \in T(\tau)$ s.t. $e \in [v]_H$
- $+1$ $\forall e \in T(\tau)$ s.t. $e \in [v]_T$

The last two are negated due to subtracting $[\partial_{(i-1)}t](v)$.

Conveniently, each simplex τ can only contribute two of these $(i-1)$-simplices e. If this isn't immediately apparent, think of each τ being isomorphic to $[i] = \{0, 1, \dots i\}$ and each v as isomorphic to $[i-2] = \{0, 1, \dots (i-2)\}$. Each face of τ, e, of which v is a face, is isomorphic to a proper superset of $[i-2]$, *and* isomorphic to a proper subset of $[i]$. There are exactly two such sets ($[i-2] \cup \{i-1\}$ and $[i-2] \cup \{i\}$).

Let us break down the contribution from each τ by case:

Case one. Both faces $e \in [v]_H$, or both faces $e \in [v]_T$. The net contribution from these will be $+1-1$ or $-1+1$, which is zero.

Case two. One face $e \in [v]_H$, while the other face $e \in [v]_T$. These must be inconsistently oriented with respect to τ, meaning one $e \in H(\tau)$, and the other $e \in T(\tau)$. They must have opposite contribution, which sum to zero.

Thus, in any case, the net contribution from any τ is zero, which makes the sum of contributions across all τ zero. Thus, for all v:

$$[\partial_{(i-1)}h](v) - [\partial_{(i-1)}t](v) = 0$$

$$\partial_{(i-1)}(h-t) = 0$$

Therefore, $(h - t) \in \ker \partial_{(i-1)}$

□

Luckily, developing the analogue for $(h' - t')$ won't be as much work:

Proposition 4. $(h' - t') \in \ker {\partial_i}^T$

Proof. By Proposition 2, we have $\langle h - t, h' \rangle = \langle h - t, t' \rangle$ (via $c = h - t$). By manipulating this equality of inner products, we get:

$$\begin{aligned} \langle h - t, h' - t' \rangle &= 0 \\ \langle h, h' - t' \rangle - \langle t, h' - t' \rangle &= 0 \\ \langle h, h' - t' \rangle &= \langle t, h' - t' \rangle \\ \langle h' - t', h \rangle &= \langle h' - t', t \rangle \end{aligned}$$

Which, by Proposition 1, taking $\alpha = h' - t'$, implies $(h' - t') \in \ker {\partial_i}^T$

□

Proposition 5. $(h - t) = \partial_i \beta$*, where* $\beta = \vec{1}$*, the all-ones vector*

Proof. Consider each row in ∂_i, as it relates to the $(i-1)$-dimensional simplex e. In each i-simplex in which e is a head, there will be a positive entry. In each i- simplex in which e is a tail, there will be a negative entry. This means that each row sum is the head-number minus the tail-number. Therefore, if $\beta = \vec{1}$, the all-ones vector, then $(h - t) = \partial_i \beta$.

□

Note: this proof only solidifies the claim that $(h - t) \in \ker \partial_{(i-1)}$, because $0 = \partial_{(i-1)} \partial_i \beta = \partial_{(i-1)}(h - t)$. However, this does not devalue the earlier proof of the presence in the kernel, because the general method of that proof (counting contributions from intermediate boundaries) could be applied to a proof of $\partial_{(i-1)} \partial_i = 0$.

Degrees and Handshaking

In graph theory, a well-known relationship known as the handshaking lemma, relates the sum of the degrees of a graph $G = (V, E)$ and the size of its edge set:

$$\sum_{v \in V} \deg(v) = 2|E|$$

Let us now develop an analogue on simplicial complexes:

First observe that, $h + t = d$, where $d(e) =$ the degree of e (i.e. the number of i-simplices of which it is a face). Additionally, by manipulation, $\langle \alpha, h \rangle + \langle \alpha, t \rangle = \langle \alpha, h + t \rangle = \langle \alpha, d \rangle$. So, if $\langle \alpha, h \rangle = \langle \alpha, t \rangle = k$, then $\langle \alpha, d \rangle = 2k$ (where $\alpha \in \ker \partial_i^T$).

Let us to restrict to the $i = 1$ case, where h, t and d act on the vertices of a connected graph, and let $\alpha = \vec{1}$, the all-ones vector. This will be a member of (and, in fact, generates) $\ker \partial_1^T$, as described in Biggs [2] (end of section 2). Then we have $\langle \alpha, d \rangle = \sum_{v \in V} \deg(v)$, and $\langle \alpha, h \rangle = \langle \alpha, t \rangle = |E|$ (because each edge is a head exactly once and a tail exactly once).

In this manner, the (equivalent) equalities $\langle \alpha, d \rangle = 2\langle \alpha, h \rangle$ and $\langle \alpha, d \rangle = 2\langle \alpha, t \rangle$ can be thought of as a general case of the handshaking lemma.

References

[1] Art M. Duval, Caroline J. Klivans, and Jeremy L. Martin, Critical Groups of Simplicial Complexes, To appear in *Annals of Combinatorics.*

[2] Norman Biggs, Chip-firing and the Critical Group of a Graph *Journal of Algebraic Combinatorics* **9** (1999), 25-45.

[3] Chris Godsil and Gordon Royle, *Algebraic Graph Theory*, Graduate Texts in Mathematics 207, Springer-Verlag, 2001.

[4] Lionel Levine and James Propp, What is. . . a Sandpile?, *Notices of the AMS* **57**, no. 8, 976-979.