The Tutte Polynomial and the Critical Configuration Polynomial on Acyclic Digraphs

A Thesis
Presented to
The Division of Mathematics and Natural Sciences
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Maxwell S. Hallock

May 2009

Approved for the Division
(Mathematics)

David Perkinson

Acknowledgements

To my mom and dad, who supported me through the bad times and the good, to Mary Hallock for kicking my ass when I needed it, and to Tom and Sarah Symes for making my education possible.

Table of Contents

List of Figures

Abstract

This paper will present the Tutte polynomial and a number of its properties, specifically its relationship to the recurrent configurations for the abelian sandpile model. We investigate generalizations to the case of a directed graph with complete success only for acyclic digraphs.

Introduction

The field of graph theory is a relatively new one. Euler published the first paper on graph theory in 1736, in which he presented his formulation and discussion of the Seven Bridges of Konigsberg problem, but most of the work in the field has occurred since the late nineteenth century. Algebraic graph theory, the application of algebraic methods—linear and abstract algebra, in particular—to graph theoretic problems, has only become a serious field of study within the last thirty years.

William T. Tutte (1917-2002) made a number of interesting advances in the field including the development of what we now know as the Tutte polynomial, $T(\Gamma; x, y)$, a polynomial in two variables, that encodes a substantial amount of interesting information about the graph. The Tutte polynomial is the subject of study for a number of recent papers ([4] [5] [6] and many others). Among these are numerous papers by Criel Merino, published over the last eleven years.

Of particular interest is Merino's first paper, "Chip-Firing and the Tutte Polynomial", published in the *Annals of Combinatorics* in 1997 [6]. In this paper, he proves an astonishing result—that for all undirected graphs, evaluated along the line $x = 1$, the Tutte polynomial is equal to the critical configuration polynomial $P(\Gamma; y)$ for the abelian sandpile model. These definitions and results are discussed in Chapter 2. However, these results were obtained for undirected graphs and no analogous result yet exists for directed graphs.

Fascinated by the $T(\Gamma; 1, y) = P(\Gamma; y)$ relationship, I took it upon myself to explore the possibility that a similar relationship exists between graph invariants for directed graphs. First, I took a similar approach to the problem that Criel does in his proof that $T(\Gamma; 1, y) = P(\Gamma; y)$, using the generating function for recurrent configurations of directed graphs, H_Γ as defined in Chapter 3.1, as my directed analogue for the critical configuration polynomial, but this proved fruitless.

Then, I happened upon a paper published by Chung and Graham in 1995 entitled, "On the Cover Polynomial of a Digraph", in which they define the cover polynomial, a graph invariant for directed graphs that, while not a perfect analogue to the Tutte polynomial in the directed case, does have a number of comparable properties. My interest piqued by the properties of this function, I explored the possibility that it, in conjunction with H_Γ, might have an analogous relationship to that of the Tutte and critical configuration polynomials. However, I quickly found a number of counterexamples for this relationship. These results are discussed in Chapter 3.1–2.

Unable to find any other directed graph algorithms analogous to the Tutte polynomial, I resolved to construct one of my own. The result of my efforts is the pruning polynomial, $X(\Gamma; r)$, which is defined and discussed in Chapter 3.3. I proved that $X(\Gamma; r) = H_\Gamma$ in the acyclic case. Furthermore, I provide a counter-example for the general directed case at the end of Chapter 3.

Chapter 1

Basic Graph Theoretic Concepts and Definitions

Before any serious discussion of the Tutte polynomial and the other graph algorithms to be addressed herein can take place, we must establish some fundamental understanding of algebraic graph theory. For the sake of expediency, we will be assuming a reasonable knowledge of abstract algebra and some basic graph theory.

Let $\Gamma = (V, E)$ be a directed graph with vetex set V and edge set E. Then, for $v, w \in V$, define $\deg(v, w)$ to be the number of edges connecting v to w. Now, let $v \in V$. Then,

$$\begin{aligned} d_v = \text{out-degree}(v) &= \sum_{w \in V} \deg(v, w) = \text{ the number of edges leaving } v. \\ \text{in-degree}(v) &= \sum_{w \in V} \deg(w, v) = \text{ the number of edges entering } v. \end{aligned}$$

A vertex $s \in V$ is a *sink* if $d_s = 0$. If, in addition, each $v \in V$ has a directed path to s, then s is called the *global sink*. If a global sink exists it is necessarily unique. Let $\tilde{V}$ be $V \backslash \{s\}$. Unless otherwise stated, Γ will be assumed to have a global sink.

We define a *sandpile* or *configuration* $c = (c(v_1), c(v_2), ..., c(v_n)) \in \mathbb{N}^n$ on Γ where $n = |\tilde{V}|$, and $c(v_i)$ is interpreted as the number of "grains of sand" on vertex v_i. See Figure 1.1 for an example of this.

A configuration c is *stable at* $v \in \tilde{V}$ if $c(v) < d_v$. Otherwise, c is *unstable* at v. A configuration c is *stable* if it is stable at all $v \in \tilde{V}$. The *maximal stable configuration* on Γ is given by

$$c_{\max}(v) = d_v - 1 \text{ for all } v \in \tilde{V}.$$

The *Laplacian*, Δ, on Γ is the operator

$$\begin{aligned} \Delta : \mathbb{Z}^V &\to \mathbb{Z}^V \\ v &\mapsto d_v v^* - \sum_{w \in V} \deg(w, v) w^*, \end{aligned}$$

where v^* is the configuration with only one grain of sand on w.

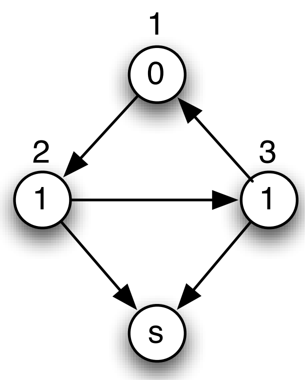

Figure 1.1: Configuration (0,1,1) on a digraph Γ with sink s.

The Laplacian can also be expressed as a matrix in the following way. First, choose an ordering of the vertices of the graph. Then,

$$\Delta_{i,j} = \begin{cases} d_{v_i} - \deg(v_i, v_i) & \text{if } i = j \\ -\deg(v_i, v_j) & \text{if } i \neq j \end{cases} \tag{1.1}$$

The *reduced Laplacian*, $\tilde{\Delta}$, is the Laplacian matrix with the row and column representing the sink removed.

Example 1. *Let Γ be the graph from Figure 1.1. Then,*

$$\Delta_\Gamma = \begin{pmatrix} 1 & -1 & 0 & 0 \\ 0 & 2 & -1 & -1 \\ -1 & 0 & 2 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

and

$$\tilde{\Delta}_\Gamma = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 2 & -1 \\ -1 & 0 & 2 \end{pmatrix}.$$

Configurations can be added to and subtracted from each other. If the number of grains of sand on a vertex, v, equals or exceeds its out-degree, we can run the *stabilization operator*, $(c(v))^0$. The vertex then "fires", sending one grain of sand to each adjacent vertex until the weight of the vertex is less than its out-degree. Please note that the stabilization operator can be applied to entire configurations. In this case, all vertices continue to fire until the configuration is stable. Since each vertex has a path to the sink, every configuration will eventually stabilize and one may show the resulting stable configuration is independent of the sequence of firings that led to it. For a formal treatment of this property of the abelian sandpile model, see Lemma 2.5 and Corollary 2.6 in [7].

A configuration, c, on Γ is *recurrent* or *critical* if and only if there exists a configuration d such that $c = (d + c_{\max})^0$.

Example 2. *Let* Γ *be the graph as seen in Figure 1.1. Then, the configuration (0,1,1), as demonstrated in the aforementioned figure, is both stable and recurrent. In fact, it is the maximum stable configuration of* Γ.

A reference for the next few definitions and theorems is Propp et al. [7]. The set of recurrent configurations on Γ forms a group, $\mathcal{S}(\Gamma)$, and

Theorem 3.

$$\mathcal{S}(\Gamma) = \frac{\mathbb{Z}^n}{\mathrm{im}(\tilde{\Delta}^t)}.$$

Another important concept is that of the *burning algorithm*. D. Dhar describes this algorithm for directed graphs in his 1999 paper, "The abelian sandpile and related models" [3]. Given a directed graph Γ and a configuration c, the algorithm has us "fire the sink" (i.e., add one grain of sand to every vertex adjacent to the sink), then recursively scan through the graph and fire any unstable vertices. If every non-sink vertex fires once and only once, then c is recurrent. We call the order of vertex firings caused by the burning algorithm the *burning sequence* for that configuration.

Please note that all of these definitons and constructions can be discussed in relation to undirected graphs after designating one of the vertices as the sink and ignoring its out-edges.

The following theorem will prove important later in the thesis.

Theorem 4. *The following statements are equivalent:*
(a) Γ *is an acyclic digraph.*
(b) The 0 *configuration is recurrent.*
(c) Every stable configuration is recurrent.

Proof. (a) $\Rightarrow$ (b): For the 0-configuration to be recurrent, there must exist some finite stable configuration c such that

$$0 = (c + c_{\max})^0.$$

Since the maximal stable configuration has $d_v - 1$ grains on every vertex, we can easily find c by adding one grain at a time until all the sand clears from the graph starting with the vertices that are farthest from the sink. The graph will eventually clear because it is acyclic and has a global sink, so no sand will get trapped on the graph. Thus, the 0-configuration is recurrent.

(b) $\Rightarrow$ (c): If the 0-configuration is recurrent, then every stable configuration is clearly reachable, so every stable configuration is recurrent.

(c) $\Rightarrow$ (a): Suppose Γ is not acyclic. Then, Γ necessarily has a cycle. Once the cycle in Γ gains sand, it will always contain at least some sand, regardless of configurations added and stabilized. Thus, the 0-configuration is not reachable. However, this contradicts our assumption that every stable configuration is recurrent. Hence, Γ is an acyclic digraph. $\square$

Chapter 2

$T(\Gamma; 1, y) = P(\Gamma; y)$ in the Undirected Case

We naturally begin our discussion of the Tutte polynomial and the critical configuration polynomial with some definitions and preliminary theorems. Please note that until otherwise indicated, we will be dealing with undirected, connected graphs.

Let Γ be a graph. Let E be the edge set of Γ and let V be the vertex set of Γ. Then, for some $e \in E$, if by removing e, the graph becomes disconnected, then e is called an *isthmus*. If both end vertices of e are the same, i.e., e connects a vertex to itself, e is called a *loop*.

Now, define Γ/e to be the *contraction* of edge e and $\Gamma - e$ to be the *deletion* of edge e. To construct Γ/e, identify the end vertices of e and then remove e. To construct $\Gamma - e$, merely remove e.

Example 5. *Let Γ be the complete graph on three vertices. Choose an edge e. The deletion and contraction of edge e is shown in Figure 2.1.*

2.1 The Tutte Polynomial

Now, we have enough background to present the linear recursion definition of the Tutte Polynomial. The following definitions and propositions are adapted from those presented by J. A. Ellis-Monaghan and C. Merino in their paper, "Graph Polynomials and Their Applications I: The Tutte Polynomial" [4].

Definition 6. *The Tutte polynomial is defined recursively such that for an undirected graph Γ and edge e,*

$$T(\Gamma; x, y) = T(\Gamma/e; x, y) + T(\Gamma - e; x, y)$$

if e is neither an isthmus nor a loop. Otherwise, Γ consists of i isthmi and j loops and

$$T(\Gamma; x, y) = x^i y^j.$$

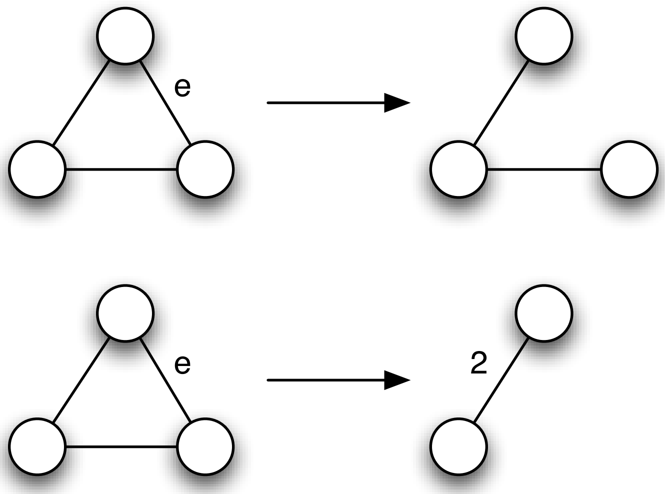

Figure 2.1: Deletion and contraction of an edge e.

It should be noted that if $E = \emptyset$, then $T(\Gamma; x, y) = 1$.

Example 7. *Let Γ again be the complete graph on three vertices. Then, $T(\Gamma; x, y) = x^2+x+y$. The construction of the Tutte polynomial on Γ can be seen in its "recursion tree" form in Figure 2.2.*

We define a *one-point-join*, $\Gamma * \Gamma'$, of two graphs Γ and Γ' to be formed by identifying a vertex u of Γ and a vertex v of Γ' into a single vertex w of $\Gamma * \Gamma'$. The disjoint union of Γ and Γ' is denoted $\Gamma \cup \Gamma'$. Then,

Proposition 8. *If Γ and Γ' are acyclic digraphs, then*

$$T(\Gamma * \Gamma'; x, y) = T(\Gamma; x, y)T(\Gamma'; x, y)$$

and

$$T(\Gamma \cup \Gamma'; x, y) = T(\Gamma; x, y)T(\Gamma'; x, y).$$

Proof. J. A. Ellis-Monaghan and C. Merino provide a proof of this in [4]. □

Theorem 9. *Let I be a graph consisting of a single isthmus and let L be a graph consisting of a single loop. Then,*

$$T(I; x, y) = x$$

and

$$T(L; x, y) = y.$$

Proof. This follows trivially from Definition 4. □

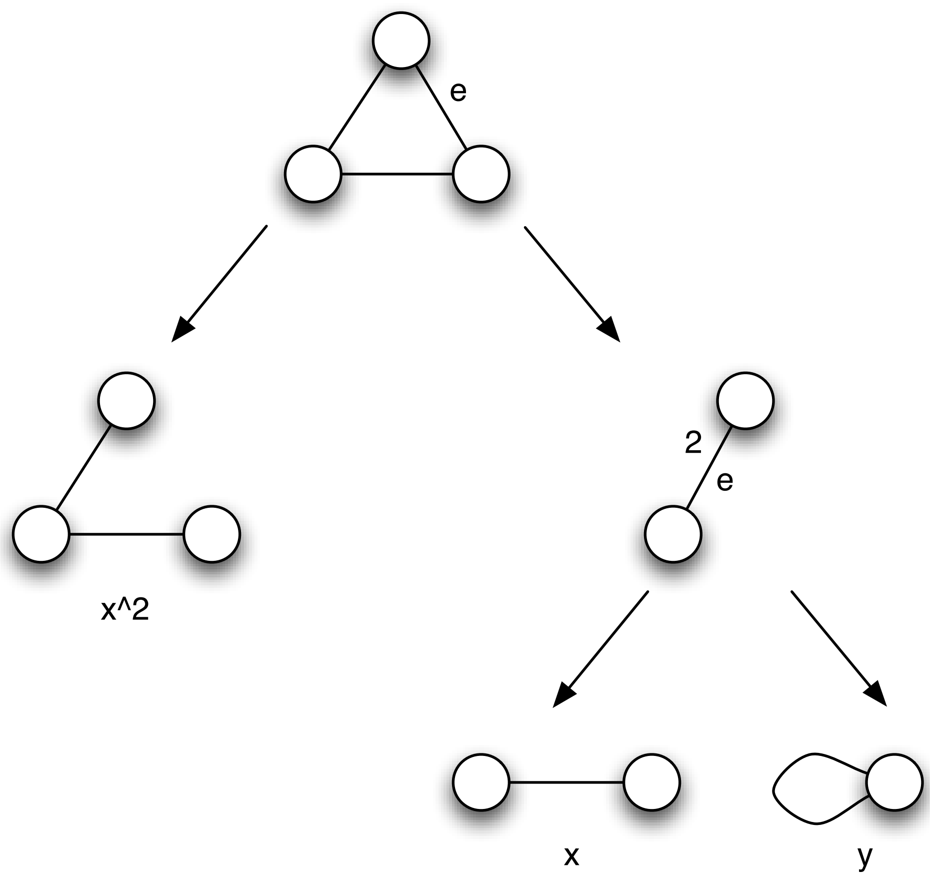

Figure 2.2: Construction of the Tutte polynomial.

2.2 Level

The remaining definitions and propositions in this chapter are adapted from those presented by C. Merino in his paper "Chip Firing and the Tutte Polynomial" [6].

Let Γ be an undirected connected graph and choose a vertex $s \in V$ to be the sink. For a configuration c and vertex $v \in \tilde{V}$,

$$c(v) = \# \text{ of grains of sand on } v.$$

Then, we can define the *total weight* of c to be

$$\mathrm{wt}(c) = \sum_{v \neq s} c(v).$$

Next, for a critical configuration c, we define the *level* of c to be

$$\mathrm{level}(c) = \mathrm{wt}(c) - |E(\Gamma)| + \deg(s),$$

where $\deg(s)$ indicates the total number of edges leaving s. To futher clarify these concepts, let us look at an example.

Example 10. *Let Γ be the complete graph on three vertices as in the previous examples. Choose a sink, s, arbitrarily. Then, the maximal stable configuration of Γ, $c_{\max}$, can be described by the vector $(1,1)$ (the sink is of course excluded). Then,*

$$\mathrm{wt}(c_{\max}) = 2$$

and

$$\mathrm{level}(c) = 2 - 3 + 2 = 1.$$

Theorem 11. *Let Γ be a graph as above and let c be a critical configuration of Γ. Then,*

$$0 \leq \mathrm{level}(c) \leq |E(\Gamma)| - |V(\Gamma)| + 1.$$

Proof. For the right inequality, consider the greatest the level can be, which would be with respect to the maximal stable configuration. So,

$$\mathrm{wt}(c_{\max}) \leq \sum_{v \neq s} (\deg(v) - 1) = 2|E(\Gamma)| - \deg(s) - |V(\Gamma)| + 1.$$

Then, for all critical c,

$$\begin{aligned}
\mathrm{level}(c) &\leq \mathrm{level}(c_{\max}) \\
&\leq 2|E(\Gamma)| - \deg(s) - |V(\Gamma)| + 1 - |E(\Gamma)| + \deg(s) \\
&= |E(\Gamma)| - |V(\Gamma)| + 1.
\end{aligned}$$

For the left inequality, it has been observed (by Björner, Lovász, and Shor [1]) that for a critical configuration c, $\mathrm{wt}(c) \geq |E(\Gamma)| - \deg(s)$. So,

$$\mathrm{wt}(c) + \deg(s) \geq |E(\Gamma)|.$$

Thus, $\mathrm{level}(c) \geq 0$. □

2.3 The Critical Configuration Polynomial

Now, we have all the background necessary to define the critical configuration polynomial.

Definition 12. *Let Γ be an undirected connected graph with a specified sink. For all $i \geq 0$, define c_i to be the number of critical configurations with level i. Then, the generating function for the critical configurations of Γ is the polynomial*

$$P(\Gamma; y) = \sum_{i=0}^{|E|-|V|+1} c_i y^i.$$

Example 13. *Let Γ be the complete graph on three vertices as in the previous examples. Then, the critical configurations are $(1,1)$, $(1,0)$, and $(0,1)$. The levels of these configurations are 1, 0, and 0, respectively. Then,*

$$P(\Gamma; y) = y + 2.$$

Theorem 14. *Let Γ be an undirected, connected graph with sink s. Then, along the line $x = 1$, the critical configuration polynomial is the evaluation of the Tutte polynomial. That is,*

$$T(\Gamma; 1, y) = P(\Gamma; y).$$

Please note that this result is independent of the choice of s. Let us illustrate this theorem with an example:

Example 15. *Let Γ be the same as in the previous examples. We know*

$$T(\Gamma; x, y) = x^2 + x + y$$

and

$$P(\Gamma; y) = y + 2.$$

Thus,

$$T(\Gamma; 1, y) = 1^2 + 1 + y = y + 2 = P(\Gamma; y).$$

Proof. Suppose Γ only has one edge, e. There are two possibilities: e is an isthmus or e is a loop. If e is an isthmus, then Γ looks like the graph in Figure 2.3 and has only one critical configuration, c, which has 0 grains on v. Then,

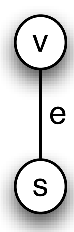

Figure 2.3: Γ such that e is an isthmus.

$$\text{level}(c) = -|E(\Gamma)| + \deg(s) = -1 + 1 = 0.$$

Thus, $P(\Gamma; y) = 1 = T(\Gamma; 1, y)$.

If e is a loop, then Γ looks like the graph in Figure 2.4 and has no configurations. Then,

$$\text{level}(c) = -|E(\Gamma)| + \deg(s) = -1 + 2 = 1.$$

Thus, $P(\Gamma; y) = y = T(\Gamma; 1, y)$.

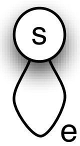

Figure 2.4: Γ such that e is a loop.

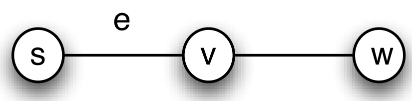

Figure 2.5: Γ such that e is a isthmus.

Now, suppose Γ has at least two edges and let $e = (v, s)$ be an edge to the sink. There are now three possibilities: e can be an isthmus, a loop, or a normal edge.

First, suppose e is an isthmus. Then, Γ will look something like the graph in Figure 2.5.

Let c' be a critical configuration on Γ/e with sink $v = s$. Then, we can get a configuration c on Γ by defining $c(w) = c'(w)$ for all $w \neq v, s$ and $c(v) = \deg(v) - 1$. Now, fire the sink (i.e., add one grain to each vertex connected to s) and run the burning algorithm. After v fires, every other vertex will fire once, since v is essentially firing into c', which is critical. Thus, c is critical.

The map defined above maps elements of $\mathcal{S}(\Gamma/e)$ into elements of $\mathcal{S}(\Gamma)$. Now, we want to show that this map is a bijection between these two sets.

Let $c \in \mathcal{S}(\Gamma)$. Define c' on Γ/e such that $c'(w) = c(w)$ for all $w \in V(\Gamma/e)\backslash\{s\}$. If we now replace the first two steps of the burning sequence (i.e., s and v) by s, we get a complete burning sequence for c', so c' is a critical configuration. Since every critical configuration of Γ assigns a value of $\deg(v) - 1$ to v, as we defined above, this map is a bijection.

Furthermore,

$$\begin{aligned}
\text{level}(c') &= \sum_{w \neq s,v} c'(w) - |E(\Gamma/e)| + \deg_{\Gamma/e}(v) \\
&= \sum_{w \neq s,v} c(w) - (|E(\Gamma)| - 1) + (\deg_\Gamma(s) + \deg_\Gamma(v) - 2) \\
&= \sum_{w \neq s,v} c(w) + \deg_\Gamma(v) - 1 - |E(\Gamma)| + \deg_\Gamma(s) \\
&= \sum_{w \neq s} c(w) - |E(\Gamma)| + \deg_\Gamma(s) \\
&= \text{level}(c).
\end{aligned}$$

Thus, $P(\Gamma; y) = P(\Gamma/e; y)$.

Second, suppose Γ is a loop. Then, Γ will look something like the graph in Figure 2.6.

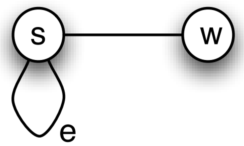

Figure 2.6: Γ such that e is a loop.

Given a critical configuration c' on $\Gamma - e$, we define a configuration c on Γ by $c(w) = c'(w)$ for all $w \in V(\Gamma)\backslash\{s\}$. Then, as in the first case, we have that c is a critical configuration and that the map between $\mathcal{S}(\Gamma - e)$ and $\mathcal{S}(\Gamma)$ is a bijection. Furthermore,

$$\begin{aligned}
\text{level}(c') &= \sum_{w \neq s} c'(w) - |E(\Gamma - e)| + \deg_{\Gamma - e}(s) \\
&= \sum_{w \neq s} c(w) - (|E(\Gamma)| - 1) + (\deg_{\Gamma}(s) - 2) \\
&= \text{level}(c) - 1.
\end{aligned}$$

Thus, $P(\Gamma; y) = yP(\Gamma - e; y)$.

Finally, suppose e is neither an isthmus nor a loop. Then, Γ will look something like the graph in Figure 2.7.

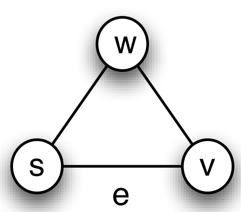

Figure 2.7: Γ such that e is a normal edge.

Now, partition $\mathcal{S}(\Gamma)$ into two sets,

A, the set of critical configurations, c, on Γ such that $c(v) = \deg(v) - 1$

and

$$A' = \mathcal{S}(\Gamma)\backslash A.$$

For a configuration $c \in A$, we associate a configuration c' on Γ/e with sink $v = s$ by defining $c'(w) = c(w)$ for all $w \in V\backslash\{v, s\}$. To determine whether or not c' is critical, fire the sink in Γ and run the burning algorithm on c and since $c(v) = \deg(v) - 1$, we can choose the first firing to be v. After we fire v, we are essentially in Γ/e, so we can reproduce the firing sequence in Γ/e with the same result. Thus, we arrive at the initial configuration c', which is critical.

Now, reverse the construction. From a critical configuration c' on Γ/e, we get c in Γ such that $c(w) = c'(w)$ for all $w \neq v$ and $c(v) = \deg(v) - 1$. The final configurations of c and c' will agree in every vertex except v. Now, fire s in Γ. Then, $c(v) = \deg(v) - 1 + \#(s,v)$, where $\#(s,v)$ represents the number of edges between s and v. The vertex v fires, loosing $\deg(v)$ grains, then all its neighbors fire, giving v $\deg(v) - \#(s,v)$ grains. Thus, after all firings, $c(v) = \deg(v) - 1$. Since v fires only once and $c(w)$ is critical for all $w \neq v$, c must be critical. Thus, we have a bijection between A and $\mathcal{S}(\Gamma/e)$.

Furthermore,

$$\begin{aligned}
\text{level}(c) &= \sum_{w\neq s} c(w) - |E(\Gamma)| + \deg_\Gamma(s) \\
&= \sum_{w\neq s,v} c'(w) + \deg_\Gamma(v) - 1 - (|E(\Gamma/e)| + 1) + \deg_\Gamma(s) \\
&= \sum_{w\neq s,v} c'(w) - |E(\Gamma/e)| + \deg_{\Gamma/e}(s) \\
&= \text{level}(c').
\end{aligned}$$

Similarly, there is a bijection between A' and $\mathcal{S}(\Gamma - e)$. First, for a critical configuration $c \in A'$ on Γ, define c' on $\Gamma - e$ such that $c'(w) = c(w)$ for all $w \neq s$. Now, fire s in $\Gamma - e$. Since the only vertex that has changed is v, any burning sequence that does not include v will still be legal in $\Gamma - e$ and v will fire once and only once when it comes up in the burning sequence because even though it has one less grain than it "needs," it also has one less out-edge.

The argument is almost identical in the other direction. Starting with a critical configuration c' on $\Gamma - e$, construct c on Γ such that $c(w) = c'(w)$ for all $w \neq s$. Now, fire s in Γ. Since v will now gain the grain of sand it lacked in c', any legal burning sequence in $\Gamma - e$ is legal in Γ. Thus, there is a bijection between A' and $\mathcal{S}(\Gamma - e)$.

Then for $c \in A'$ and its corresponding c',

$$\begin{aligned}
\text{level}(c) &= \sum_{w\neq s} c(w) - |E(\Gamma)| + \deg_\Gamma(s) \\
&= \sum_{w\neq s} c'(w) - (|E(\Gamma - e)| + 1) + (\deg_{\Gamma - e}(s) + 1) \\
&= \text{level}(c').
\end{aligned}$$

Hence, $P(\Gamma;y) = P(\Gamma - e;y) + P(\Gamma/e;y)$.

Now, we use induction on the number of edges with respect to the recursions and trivial cases we found above. Let N be the number of edges of Γ. Then, for $N = 1$, we get the trivial cases:

$$P(I;y) = 1 = T(I;1,y) \text{ and } P(L;y) = y = T(L;1,y).$$

Suppose the proposition holds for $N = n$. Let Γ^+ be a graph with $n+1$ edges and choose an edge e connecting the sink, s, to some vertex, v, which could be s. Then, there are three possibilities: e is either an isthmus, a loop, or a normal edge.

If e is an isthmus, then

$$\begin{aligned} P(\Gamma^+; y) &= 1 \cdot P(\Gamma; y) \\ &= 1 \cdot T(\Gamma; 1, y) \\ &= T(\Gamma^+; 1, y). \end{aligned}$$

If e is a loop, then

$$\begin{aligned} P(\Gamma^+; y) &= y \cdot P(\Gamma; y) \\ &= y \cdot T(\Gamma; 1, y) \\ &= T(\Gamma^+; 1, y). \end{aligned}$$

If e is a normal edge, then

$$\begin{aligned} P(\Gamma^+; y) &= P(\Gamma^+/e; y) + P(\Gamma^+ - e; y) \\ &= P(\Gamma^+/e; y) + P(\Gamma; y) \\ &= T(\Gamma^+/e; y) + T(\Gamma; y) \\ &= T(\Gamma^+/e; y) + T(\Gamma^+ - e; y). \end{aligned}$$

Thus, the proposition holds for $N = n+1$. Hence,

$$P(\Gamma; y) = T(\Gamma; 1, y).$$

□

Chapter 3

The Directed Case

Unlike the Tutte polynomial, the critical configuration polynomial does have an almost perfect analogue in the directed case. For the following discussion, let $\Gamma = (V, E)$ be a directed graph with vertex set V and edge set E.

Definition 16. *The generating function for the weights of the recurrent elements of* Γ *is*

$$H_\Gamma(y) = \sum_{i=0}^{n} a_i y^i,$$

where a_i *is the number of recurrent configurations with* $\mathrm{wt}(c) = i$ *and* $n = \mathrm{wt}(c_{\max})$.

3.1 $\tilde{\Gamma}$

Our goal now is to show some of the problems with trying to extend the method of proof of Theorem 14 in the case of a directed graph.

Definition 17. *For digraph* Γ *with isthmus* e *connecting a vertex,* v*, into the global sink,* s*, let* $\tilde{\Gamma}$ *be the graph such that the out-edges of* v*, the edge* e*, and the sink* s *are removed, effectively making* v *the sink.*

For an example of this see Figure 3.1.

Theorem 18. *Let* Γ *be a digraph with global sink,* s*. Suppose there is an isthmus connecting a vertex,* v*, into* s*. Then,*

$$\mathcal{S}(\Gamma) \approx \mathcal{S}(\tilde{\Gamma}),$$

i.e., Γ *and* $\tilde{\Gamma}$ *have isomorphic collections of recurrent configurations.*

Proof. Given the transpose of the reduced Laplacian of Γ, we apply row operations U and V such that

$$\tilde{\Delta}^t = \begin{pmatrix} | & | & | & | \\ | & | & | & | \\ | & | & | & | \\ | & | & | & d_v \end{pmatrix} \xrightarrow{U} \begin{pmatrix} | & & | & | \\ | & & | & | \\ | & & | & | \\ 0 & \cdots & 0 & 1 \end{pmatrix} \xrightarrow{V} \begin{pmatrix} | & & | & 0 \\ | & & | & \vdots \\ | & & | & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix}.$$

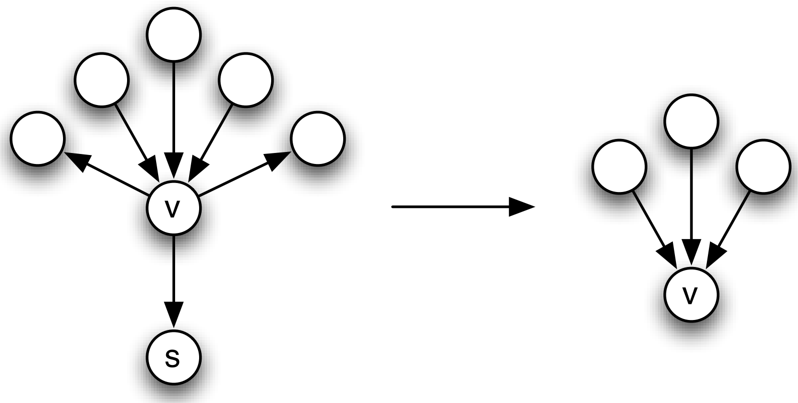

Figure 3.1: $\Gamma \to \tilde{\Gamma}$

In this first step, we add all the other rows to the bottom row, leaving us with $(0, \ldots, 0, 1)$ in the last row. In the second step, we add or subtract the last row from the rows with non-zero elements in the last column to get $(0, \ldots, 0, 1)^t$ in the last column. In other words, $VU\tilde{\Delta}^t = (\Delta')^t$. Also,

$$(\Delta')^t = \begin{pmatrix} (\Delta'')^t & \vec{0}^t \\ \vec{0} & 1 \end{pmatrix}.$$

Note that Δ'' is the reduced Laplacian of $\tilde{\Gamma}$. Since elementary row operations preserve algebraic structure, $\mathcal{S}(\Gamma) \approx \mathcal{S}(\tilde{\Gamma})$. □

Now that we know Γ and $\tilde{\Gamma}$ have the same groups of recurrent configurations, we would like to be able to find the recurrents of $\tilde{\Gamma}$ and see how they relate to the recurrents of Γ. We would like to make the isomorphism in the proof of Theorem 18 explicit.

For $a \in \mathbb{Z}^n$,

$$a = \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \xrightarrow{U} \begin{pmatrix} a_1 \\ \vdots \\ a_{n-1} \\ a_1 + \cdots + a_n \end{pmatrix} \xrightarrow{V} \begin{pmatrix} a_1 + \mathrm{wt}(v_n, 1)(a_1 + \cdots + a_n) \\ \vdots \\ a_{n-1} + \mathrm{wt}(v_{n-1}, 1)(a_1 + \cdots + a_n) \\ (a_1 + \cdots + a_n) \end{pmatrix} \geq 0.$$

Now, let $\epsilon_i = \mathrm{wt}(v_n, v_i)$ and $\mathrm{wt}(a) = \sum_{i=1}^n a_i$. Then,

$$VUa = \begin{pmatrix} a_1 + \epsilon_1 \mathrm{wt}(a) \\ \vdots \\ a_{n-1} + \epsilon_{n-1} \mathrm{wt}(a) \\ \mathrm{wt}(a) \end{pmatrix}.$$

Removing the last coordinate gives the isomorphism

$$\begin{aligned} \mathcal{S}(\Gamma) &\to \mathcal{S}(\tilde{\Gamma}) \\ a &\mapsto \begin{pmatrix} a_1 + \epsilon_1 \mathrm{wt}(a) \\ \vdots \\ a_{n-1} + \epsilon_{n-1}\mathrm{wt}(a) \end{pmatrix}. \end{aligned}$$

Defining $\tilde{a} = (a_1, ..., a_{n-1})$, dropping the last coordinate, and letting $\epsilon = (\epsilon_1, ..., \epsilon_{n-1})$ we can write this isomorphism as

$$\begin{aligned} \mathcal{S}(\Gamma) &\to \mathcal{S}(\tilde{\Gamma}) \\ a &\mapsto \tilde{a} + \epsilon\, \mathrm{wt}(a). \end{aligned}$$

Now, suppose we have a recurrent configuration a on Γ. Then, is $\tilde{a}+\epsilon\mathrm{wt}(a)$ recurrent? Sadly, no. Consider the maximal configuration, $c_{\max}$, on $\tilde{\Gamma}$. In general, the image of $c_{\max}$ will not even be stable on $\tilde{\Gamma}$. We illustrate this with an example:

Example 19. *Consider the Γ and $\tilde{\Gamma}$ presented in Figure 3.2.*

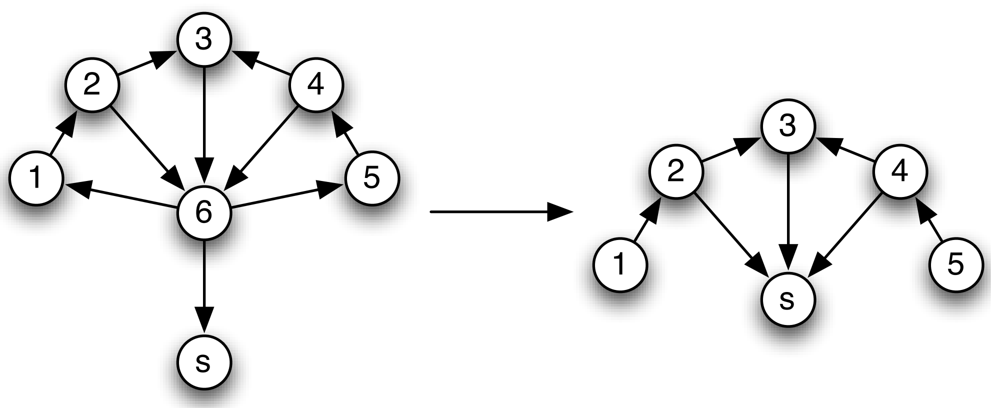

Figure 3.2: $\Gamma \to \tilde{\Gamma}$

Using the mathematics software SAGE [8], we find the recurrents for Γ to be (0,1,0,1,0,2), (0,1,0,0,0,2), (0,1,0,1,0,1), and (0,0,0,1,0,2) and the recurrents for $\tilde{\Gamma}$ to be (0,1,0,1,0), (0,0,0,1,0), (0,0,0,0,0), and (0,1,0,0,0). So, as one can see, the number of recurrent configurations is the same, as predicted. However,

$$H_\Gamma(y) = y^4 + 3y^3 \textit{ and } H_{\tilde{\Gamma}}(y) = y^2 + 2y + 1.$$

These discrepancies between the recurrent configurations and recurrent generating polynomials of Γ and $\tilde{\Gamma}$ indicate that even though $\mathcal{S}(\Gamma)$ and $\mathcal{S}(\tilde{\Gamma})$ are isomorphic as groups, there does not exist a bijection between their recurrent configurations preserving weights. Hence, the operation $\Gamma \to \tilde{\Gamma}$ does not yield the desired relation, $H_\Gamma(y) = H_{\tilde{\Gamma}}(y)$, in the directed case.

3.2 The Cover Polynomial

In 1993, Chung and Graham published a paper entitled "On the Cover Polynomial of a Digraph". This paper describes the cover polynomial, which, while it is "not exactly the directed analogue of the Tutte polynomial, it does have a number of properties which are comparable to $T(G; x, y)$" [2].

Chung and Graham define the operations of deletion and contraction slightly differently, so we will state the new definitions here. Given a directed graph, Γ, and an edge, $e = (u, v)$, $\Gamma - e$ is constructed by merely removing e and Γ/e is constructed by identifying the vertices u and v and removing all edges of the form (u, x) and (y, v). The definition of the cover polynomial follows [2].

Definition 20. *Given a digraph Γ, if Γ is composed of n independent vertices and no edges, then Γ has cover polynomial*

$$C(\Gamma) = x(x-1)(x-2)\cdots(x-n+1).$$

In the special case where $n = 0$, the corresponding digraph Γ having no vertices or edges has cover polynomial

$$C(\Gamma) = 1.$$

If e is an edge of Γ which is not a loop, then

$$C(\Gamma) = C(\Gamma - e) + C(\Gamma/e)$$

and if e is a loop, then

$$C(\Gamma) = C(\Gamma - e) + yC(\Gamma/e).$$

Sadly, this definition does not lead to the desired relationship of

$$C(\Gamma; 1, y) = H_\Gamma(y)$$

in any case! Next, we present a pair of counterexamples for the cyclic and acyclic cases.

For example, in the general directed case,

Example 21. *Consider the graph, Γ in Figure 3.3. Then,*

$$C(\Gamma; x, y) = x^4 - 2x^3 + 4x^2 - 2x + xy$$

and

$$C(\Gamma; 1, y) = y + 1.$$

However, $H_\Gamma(y) = y \neq C(\Gamma; 1, y)$.

In the acyclic case,

Example 22. *Consider the graph, Γ, in Figure 3.4. Then,*

$$C(\Gamma; x, y) = x^3$$

and

$$C(\Gamma; 1, y) = 1.$$

However, $H_\Gamma(y) = 1 + y \neq C(\Gamma; 1, y)$.

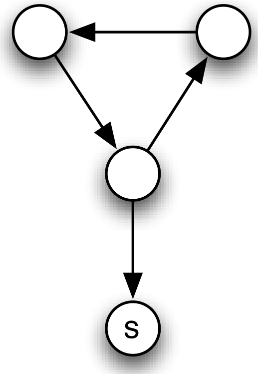

Figure 3.3: Γ

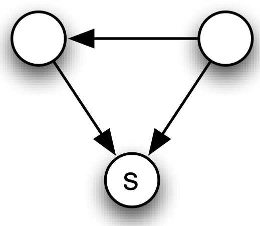

Figure 3.4: Γ

3.3 The Pruning Polynomial

Since we know the cover polynomial is insufficient, we must devise a new algorithm that has the desired relationship with the generating function, H_Γ, we defined at the beginning of this chapter. In this section, we restrict ourselves to the case of acyclic graphs. Consider the following, which I will dub the pruning polynomial (so-called because each iteration of the recursion "trims down" the graph):

Definition 23. *Given a directed, acyclic graph Γ, choose a vertex w and an out-edge e of w. $\Gamma - e$ remains the same as before, but we redefine Γ/e to be Γ with all the out-edges of w removed.*

If out-degree$(w) > 1$, *we define the pruning polynomial of Γ to be*

$$X(\Gamma; y) = yX(\Gamma - e; y) + X(\Gamma/e; y).$$

If out-degree$(w) = 1$,

$$X(\Gamma; y) = X(\Gamma/e; y).$$

In the special case where Γ has no edges,

$$X(\Gamma; y) = 1.$$

We will illustrate the process of the pruning polynomial with the following example.

Example 24. *Consider the graph, Γ, in Figure 3.4. In Figure 3.5, we see how it is acted on by the pruning polynomial in its "recursion tree" form where r represents "delete" and c represents "contract". By the definition of the pruning polynomial, we have*

$$X(\Gamma; y) = 1 + y,$$

which in this case happens to be equal to $H_\Gamma(y)$. However, we will prove this relation for all acyclic digraphs.

Our first goal is to show that the pruning polynomial is well-defined for acyclic digraphs. In other words, we would like to show that the output of the pruning polynomial is independent of the choice of edges. However, if we can show that $X(\Gamma; y) = H_\Gamma(y)$, then the proof that the pruning polynomial is well-defined is trivial. So, we will proceed.

Theorem 25. *Let Γ be an acyclic digraph. Then,*

$$X(\Gamma; y) = H_\Gamma(y).$$

Proof. Let $\mathrm{Stab}(\Gamma')$ be the number of stable "configurations" on any graph $\Gamma' = (V', E')$, which may not necessarily be connected or have a global sink. By a stable configuration, c on Γ', we mean any vector $c = (c(v_1), \ldots, c(v_n))$ such that $c(v) < d_v$ for all $v \in V'$ and $c(v) = 0$ if v is a sink. Then, for any $c \in \mathrm{Stab}(\Gamma')$, let $\mathrm{wt}(c) = \sum_{v \in V'} c(v)$, as before.

Define $\tilde{H}_{\Gamma'}(y)$ to be

$$\tilde{H}_{\Gamma'}(y) = \sum_{i=0}^{n} a_i y^i,$$

where $a_i = \#\{c \in \mathrm{Stab}(\Gamma') : \mathrm{wt}(c) = i\}$.

If Γ is an acyclic digraph, then $\tilde{H}_\Gamma(y) = H_\Gamma(y)$ since the recurrent configurations on Γ are the same as the stable configurations. Now, we will prove that $X(\Gamma; y) = \tilde{H}_\Gamma(y)$ for any graph Γ by induction on the number of edges.

Base Case: Suppose Γ has no edges. In this case, we define $\tilde{H}_\Gamma(y) = X(\Gamma; y) = 1$, where the one configuration is the zero configuration. Now, suppose Γ has at least one edge, $e = (w, w')$. This leaves us with two cases.

Case 1: Suppose out-degree$(w) = 1$. Then, given a configuration, c, on Γ/e, there exists a mapping

$$\begin{aligned} \phi : \mathrm{Stab}(\Gamma/e) &\to \mathrm{Stab}(\Gamma) \\ c(v) &\mapsto c(v) \end{aligned}$$

for all $v \in V_{\Gamma/e}$. Since $c(w) = 0$, necessarily, Γ has no more stable configurations than Γ/e and thus ϕ is a bijective mapping. The mapping ϕ also preserves configuration weights. Thus, $\tilde{H}_{\Gamma/e}(y) = \tilde{H}_\Gamma(y)$. Hence, $X(\Gamma; y) = \tilde{H}_\Gamma(y)$.

Case 2: Now, suppose out-degree$(w) > 1$. Then, there exists a mapping

$$\begin{aligned} \phi : \mathcal{S}(\Gamma/e) &\hookrightarrow \mathcal{S}(\Gamma) \\ c(v) &\mapsto c(v) \end{aligned}$$

for all $v \in V_{\Gamma/e}$. Note that ϕ is injective.

Given a stable configuration $c \in \mathcal{S}(\Gamma - e)$, there exists a mapping

$$\begin{aligned} \phi' : \mathcal{S}(\Gamma - e) &\hookrightarrow \mathcal{S}(\Gamma) \\ c(v) &\mapsto c(v) + w^* \end{aligned}$$

for all $v \in V_\Gamma$. So, stable configurations on $\Gamma - e$ correspond to stable configurations on Γ with one more grain of sand on w and, thus, one more weight overall.

We know that

$$|\mathrm{Stab}(\Gamma - e)| = (d_w - 1) \prod_{v \neq w} d_v$$

and

$$|\mathrm{Stab}(\Gamma/e)| = \prod_{v \neq w} d_v.$$

Thus,

$$|\mathrm{Stab}(\Gamma - e)| + |\mathrm{Stab}(\Gamma/e)| = |\mathrm{Stab}(\Gamma)| = \prod_{v \in V} d_v.$$

The images of ϕ and ϕ' are disjoint because

$$\mathrm{im}(\phi) = \{c \in \mathrm{Stab}(\Gamma) : c(w) = 0\}$$

whereas

$$\mathrm{im}(\phi') = \{c \in \mathrm{Stab}(\Gamma) : c(w) > 0\}.$$

So, together, they account for all of the stable configurations on Γ. Furthermore,

$$X(\Gamma; y) = yX(\Gamma - e; y) + X(\Gamma/e; y),$$

thus accounting for the loss in overall weight between the elements of $\mathrm{Stab}(\Gamma - e)$ and $\mathrm{Stab}(\Gamma)$. Hence, $X(\Gamma; y) = \tilde{H}_\Gamma(y)$.

Thus, by induction, $X(\Gamma; y) = \tilde{H}_\Gamma(y)$ for all acyclic digraphs Γ. □

Corollary 26. *The pruning polynomial is well-defined.*

Proof. This follows trivially from the preceeding theorem. □

Recall the *one-point-join*, $\Gamma * \Gamma'$, and *disjoint union*, $\Gamma \cup \Gamma'$, constructions. The disjoint union of acyclic digraphs is not feasable because the resulting graph would have no global sink, so we will ignore it. Then,

Proposition 27. *If Γ and Γ' are acyclic digraphs, then*

$$X(\Gamma * \Gamma'; y) = X(\Gamma; y)X(\Gamma'; y).$$

Proof. One-point-joins between acyclic digraphs can only occur between the sink of one subgraph and any vertex of the other digraph. The result is then trivial by the definition of the pruning polynomial. $\square$

Theorem 28. *Let Γ be an acyclic digraph. Then,*

$$X(\Gamma; y) = \prod_{v \in V} (1 + y + \ldots + y^{d_v - 1})$$

where we take the factor corresponding to v to be 1 if $d_v \leq 1$. Letting V_2 denote the vertices with out-degree at least 2, we can write this as

$$\begin{aligned} X(\Gamma; y) &= \prod_{v \in V_2} \frac{1 - y^{d_v}}{1 - y} \\ &= \frac{\prod_{v \in V_2} 1 - y^{d_v}}{(1 - y)^{|V_2| - 1}}. \end{aligned}$$

Proof. It is clear that

$$\prod_{v \in V} (1 + y + \ldots + y^{d_v - 1}) = \prod_{v \in V_2} \frac{1 - y^{d_v}}{1 - y} = \frac{\prod_{v \in V_2} 1 - y^{d_v}}{(1 - y)^{|V_2| - 1}},$$

so we have only to show that the pruning polynomial is equivalent to one of these formulations. Specifically, we will show that

$$X(\Gamma; y) = \prod_{v \in V} (1 + y + \ldots + y^{d_v - 1}).$$

Each factor of $\prod_{v \in V} (1 + y + \ldots + y^{d_v - 1})$ encodes information about the possible amounts of sand on each vertex based on its out-degree. A vertex with an out-degree of 1, for example, can have at most 0 grains of sand on it, so the factor that represents that vertex would be $y^0 = 1$. Similarly, a vertex with out-degree of 2 will have a factor of $1 + y$ because it can have either 1 grain or 0 grains. Since Γ is an acyclic digraph, we can construct a configuration by choosing one monomial from each factor and multiplying them together. This multiplication will render a monomial of the form y^i, where i is the total weight of that configuration on Γ. Repeating this for every combination of monomials from each factor and summing them together will render a polynomial that encodes the number of configurations of each weight. We know all possible configurations are represented because we have encoded all possible weights on every vertex. Hence,

$$X(\Gamma; y) = \prod_{v \in V} (1 + y + \ldots + y^{d_v - 1}).$$

$\square$

Thus, we have a closed form definition of the pruning polynomial.

It is also important to note that while the pruning polynomial functions exactly as we would like it to on acyclic digraphs, it consistently fails in the general directed case. The following counter-example demonstrates this.

Example 29. *Consider the graph, Γ, in Figure 3.4. Then Γ has three recurrent configurations:* $(1,1),(1,0)$, *and* $(0,1)$. *Thus, calculating $X(\Gamma;y)$ as if it were an acyclic graph,*

$$X(\Gamma;y) = 1 + 2y + y^2 \neq 2y + y^2 = H_\Gamma(y).$$

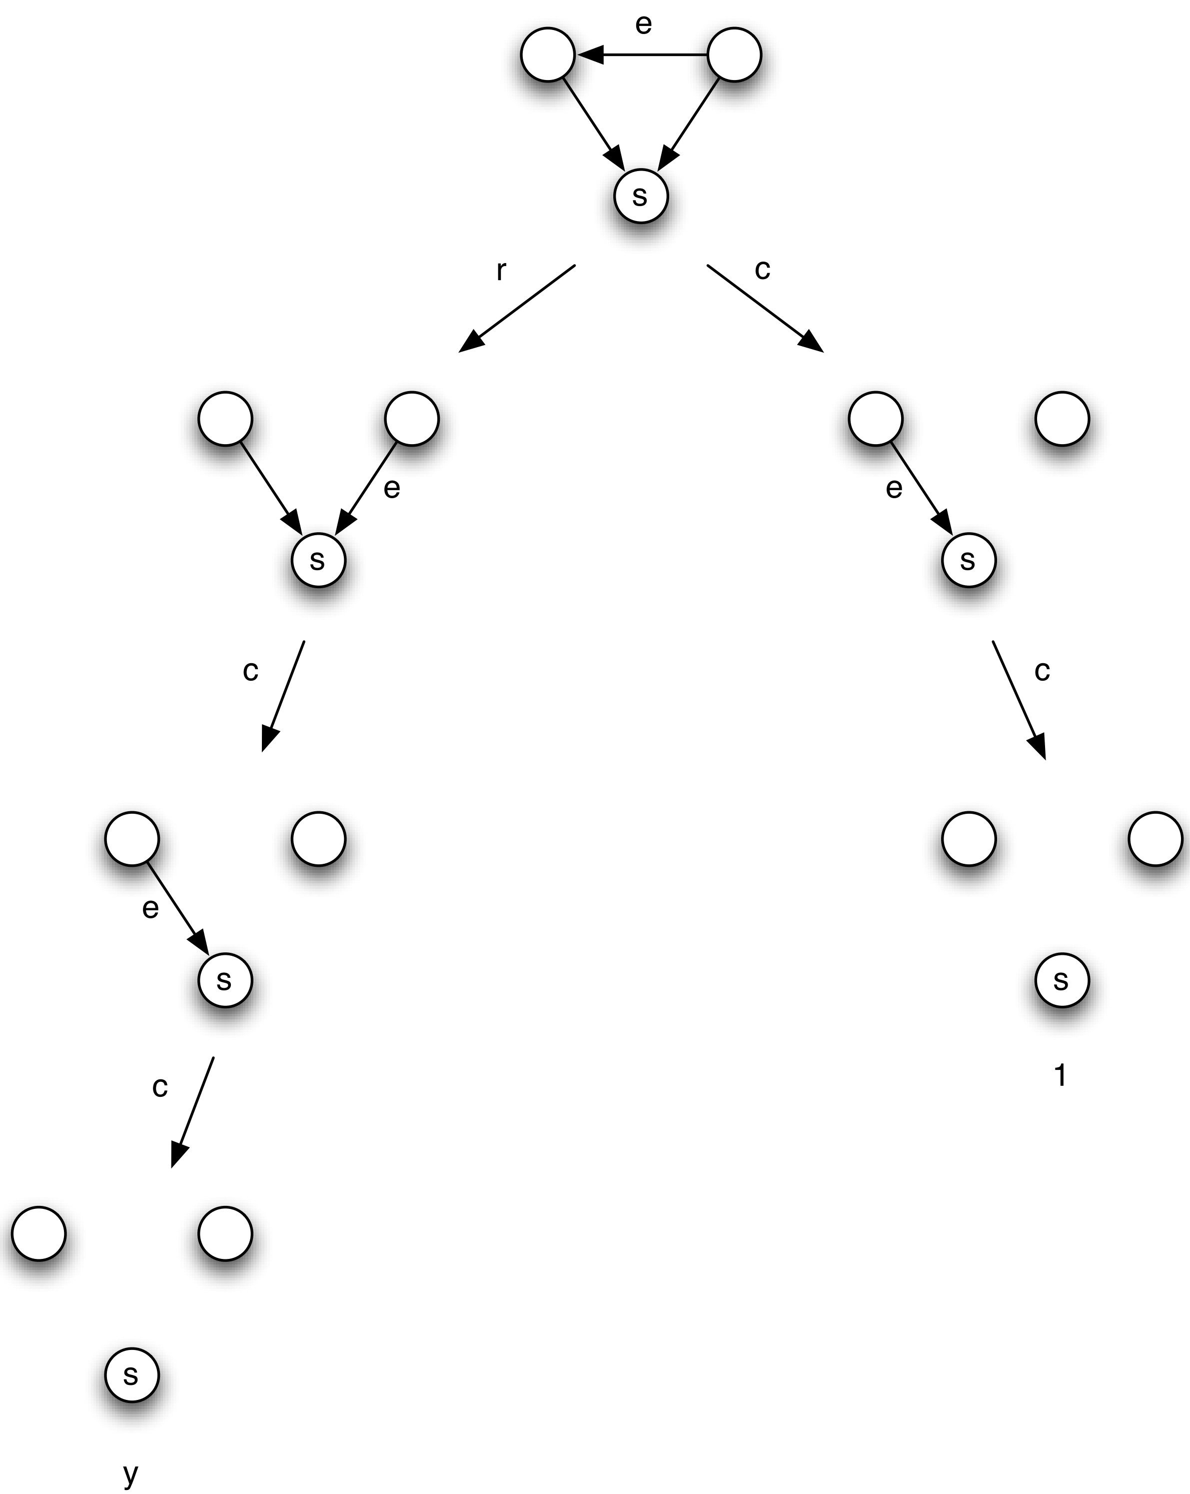

Figure 3.5: Construction of the pruning polynomial on the acyclic digraph, Γ.

Conclusion

Our final result, that there is an analogous relationship to $T(\Gamma; 1, y) = P(\Gamma; y)$ in the acyclic directed case, while interesting, is not particularly significant. The most significant part of the result is that it provides a base case from which others might hope to generalize. Specifically, we have shown that there is a nice way to calculate the generating function for recurrent configurations in the directed acyclic case. The goal of future research on this particular subject would be in part to show that such a result exists in the generalized directed case.

References

[1] A. Björner, L. Lovász, and P. W. Shor. Chip-firing games on graphs. *European J. Combinatorics*, 12:283–291, 1991.

[2] F. R. K. Chung and R. L. Graham. On the cover polynomial of a digraph. *Journal of Combinatorial Theory*, pages 273–290, January 1995.

[3] Deepak Dhar. The abelian sandpile and related models. *Physica A*, 4(263):1–24, 1999.

[4] Joanna A. Ellis-Monaghan and Criel Merino. Graph polynomials and their applications 1: The tutte polynomial. *arXiv:0803.3079v2 [math.CO] 29 Jun 2008*, pages 1–42, June 2008.

[5] Andrew Goodall. Fourier analysis on finite abelian groups: some graphical applications. In *Combinatorics, complexity, and chance*, volume 34 of *Oxford Lecture Ser. Math. Appl.*, pages 103–129. Oxford Univ. Press, Oxford, 2007.

[6] Criel Merino López. Chip firing and the Tutte polynomial. *Annals of Combinatorics*, 1(3):253–259, 1997.

[7] James Propp et al. Chip-firing and rotor-routing on directed graphs. *In and Out of Equilibrium*, 2(60):331–364, 2008.

[8] W. A. Stein et al. *Sage Mathematics Software (Version 3.4)*. The Sage Development Team, 2009. `http://www.sagemath.org`.