Discrete Mathematics 340 (2017) 1287–1299

Contents lists available at ScienceDirect

Discrete Mathematics

journal homepage: www.elsevier.com/locate/disc

On the roots of all-terminal reliability polynomials

Jason Brown [a], Lucas Mol [b,*]

[a] Department of Mathematics and Statistics, Dalhousie University, Halifax, Nova Scotia, Canada B3H 3J5
[b] University of Winnipeg, Department of Mathematics & Statistics, Winnipeg, Manitoba, Canada R3B 2E9

ARTICLE INFO

Article history:
Received 11 March 2016
Received in revised form 23 January 2017
Accepted 27 January 2017
Available online 25 February 2017

Keywords:
Graph polynomials
All-terminal reliability
Polynomial roots
Brown–Colbourn conjecture

ABSTRACT

Given a graph G in which each edge fails independently with probability $q \in [0, 1]$, the *all-terminal reliability* of G is the probability that all vertices of G can communicate with one another, that is, the probability that the operational edges span the graph. The all-terminal reliability is a polynomial in q whose roots (*all-terminal reliability roots*) were conjectured to have modulus at most 1 by Brown and Colbourn. Royle and Sokal proved the conjecture false, finding roots of modulus larger than 1 by a slim margin. Here, we present the first nontrivial upper bound on the modulus of any all-terminal reliability root, in terms of the number of vertices of the graph. We also find all-terminal reliability roots of larger modulus than any previously known. Finally, we consider the all-terminal reliability roots of simple graphs; we present the smallest known simple graph with all-terminal reliability roots of modulus greater than 1, and we find simple graphs with all-terminal reliability roots of modulus greater than 1 that have higher edge connectivity than any previously known examples.

1. Introduction and background

Let $G = (V, E)$ be an undirected, finite, loopless, connected (multi)graph in which each edge fails independently with probability $q \in [0, 1]$ and vertices are always reliable. The *all-terminal reliability* of G, denoted $\mathrm{Rel}(G; q)$, is the probability that all vertices of G can communicate with one another; that is, the probability that the operational edges span the graph. All-terminal reliability is a well-studied model of network robustness, and much research has been carried out on a variety of algorithmic and theoretical aspects of the model including algorithmic complexity, polynomial time algorithms for restricted families, efficient bounding procedures, the existence of optimal graphs, and analytic properties (when the all-terminal reliability of a graph is viewed as a function of q). See [8], for example, or [3] for a more recent survey on all-terminal reliability. Note that all-terminal reliability is often studied in terms of $p = 1-q$, the probability that each edge is operational, but our results on all-terminal reliability are easier to state and prove in terms of q, so we deal exclusively in the variable q in this article.

The all-terminal reliability of a connected graph G with edge set E, denoted $\mathrm{Rel}(G; q)$, is always a polynomial in q of degree (at most) $m = |E|$, as a subgraph with operational edges $E' \subseteq E$ arises with probability

$$(1-q)^{|E'|}q^{|E|-|E'|}.$$

Summing this probability over all sets E' for which all vertices of G can communicate gives the all-terminal reliability of G. The fact that the polynomial has degree exactly m will be seen later from the H-form of the polynomial.

* Corresponding author.
E-mail address: lucas.mol@dal.ca (L. Mol).

http://dx.doi.org/10.1016/j.disc.2017.01.024

It is natural to inquire about the nature and location of the roots of all-terminal reliability polynomials, called *all-terminal reliability roots* or *ATR roots* henceforth. ATR roots were noted to have modulus at most 1 (in q) for small graphs, and it was conjectured in [2] that this was the case for all graphs. This contrasts sharply with what is known for other graph polynomials, such as chromatic polynomials [15], independence polynomials [5], and domination polynomials [6], where the roots are dense in the complex plane. Despite some results and generalizations in the affirmative [7,16], the conjecture for ATR roots was shown to be false in [14]. However:

- The ATR roots provided were only outside of the unit disk by a slim margin; the largest modulus of an ATR root found was approximately 1.04.
- The simple graphs with ATR roots outside of the unit disk were quite large, with the smallest having over 1500 vertices and over 3000 edges.
- All of the simple graphs with ATR roots outside of the unit disk had many vertices of degree 2, and it is unclear whether all simple graphs with ATR roots outside of the unit disk have such low edge connectivity.

Finally, although ATR roots of modulus greater than 1 were found, no general upper bound on the modulus of an ATR root was given.

In this article, we continue the exploration of the location of ATR roots. In Section 2, we find a nontrivial (though non-constant) bound on the modulus of any ATR root of a graph G in terms of the order of G (this extends a weaker result in [9]). In Section 3.1, we study graphs with ATR roots of modulus greater than 1, finding graphs with ATR roots of greater modulus than any previously known. Finally, in Section 3.2 we consider simple graphs with ATR roots of modulus greater than 1. We find a smaller example of a simple graph with ATR roots outside of the unit disk, and we find simple graphs that have ATR roots outside of the unit disk and have much higher edge connectivity than any previously known examples.

We shall need some background on reliability for the next section. For a graph G of order n and size m (that is, with n vertices and m edges), we can express the all-terminal reliability polynomial of G as

$$\mathrm{Rel}(G;q) = \sum_{i=0}^{m-n+1} F_i q^i (1-q)^{m-i},$$

where F_i denotes the number of subsets of E of cardinality i whose removal leaves the graph connected. The collection of all such subsets of G is called the *cographic matroid* of G, and the sequence $(F_0, \ldots, F_{m-n+1})$ is called the *F-vector* of the cographic matroid of G (see [8] for more details). The generating polynomial of the F-vector of G is called the *F-polynomial* of G, denoted $F(G;x)$. It is known that one can rewrite the polynomial in its *H-form* as

$$\mathrm{Rel}(G;q) = (1-q)^{n-1} \sum_{k=0}^{m-n+1} H_k q^k.$$

The sequence $(H_0, \ldots, H_{m-n+1})$ is called the *H-vector* of the cographic matroid of G. Moreover, the generating polynomial

$$H(G;x) = \sum_{k=0}^{m-n+1} H_k x^k$$

of the H-vector turns out to be an evaluation of the well-known two-variable *Tutte polynomial* (see [12]):

$$T(G;1,x) = H(G;x). \tag{1}$$

There is a relatively new interpretation to the coordinates of the H-vector of a cographic matroid that we shall find particularly useful. We describe the *chip-firing game* that yields this new interpretation. Let $G = (V, E)$ be a connected multigraph without loops, and let w denote a special vertex of G. A *configuration* of G is a function $\theta : V \to \mathbb{Z}$ for which $\theta(v) \geq 0$ for all $v \neq w$ and $\theta(w) = -\sum_{v \neq w} \theta(v)$. For $v \neq w$, the number $\theta(v)$ represents the number of chips on vertex v. We imagine that the special vertex w has infinitely many chips. In configuration θ, a vertex $v \neq w$ is *ready to fire* if $\theta(v) \geq \deg(v)$; vertex w is ready to fire if and only if no other vertex is ready. *Firing* vertex u changes the configuration from θ to θ', where

$$\theta'(u) = \theta(u) - \deg(u)$$

and for $v \neq u$

$$\theta'(v) = \theta(v) + l(u, v),$$

where $l(u, v)$ is the number of edges between u and v in G. A configuration is *stable* when $\theta(v) < \deg(v)$ for all $v \neq w$; that is, if and only if w is ready to fire.

A *firing sequence* $\Theta = (\theta_0, \theta_1, \ldots, \theta_k)$ is a sequence of configurations in which θ_i is obtained from θ_{i-1} by firing one vertex that is ready to fire for each $i \in \{1, \ldots, k\}$. It is *nontrivial* when $k > 0$. We write $\theta_0 \to \theta_k$ when some nontrivial firing sequence starting with θ_0 and ending with θ_k exists. Configuration θ is *recurrent* if $\theta \to \theta$. Stable, recurrent configurations are called *critical*. For a critical configuration θ, a *critical sequence* is a legal firing sequence of minimal length that makes θ recur. Merino [11,12] proved the following surprising result which connects the critical configurations to all-terminal reliability: