

# Geodesic cycle length distributions in fictional character networks

Alex Stivala<sup>1</sup>

March 22, 2023

## Abstract

A geodesic cycle in a graph is a cycle with no shortcuts, so that the shortest path between any two nodes in the cycle is the path along the cycle itself. A recently published paper used random graph models to investigate the geodesic cycle length distributions of a unique set of delusional social networks, first examined in an earlier work, as well as some other publicly available social networks. Here I test the hypothesis, suggested in the former work, that fictional character networks, and in particular those from works by a single author, might have geodesic cycle length distributions which are extremely unlikely under random graph models, as the delusional social networks do. The results do not show any support for this hypothesis. In addition, the recently published work is reproduced using a method for counting geodesic cycles exactly, rather than the approximate method used originally. The substantive conclusions of that work are unchanged, but some differences in the results for particular networks are described.

**Keywords**— Geodesic cycle, Isometric cycle, Exponential random graph model, ERGM,  $dk$ -series random graphs, Social networks, Fictional character networks

## 1 Introduction

In order to investigate the hypothesis that people tend to conceive of social networks in local and spatial terms, Martin (2017) used random graphs to analyze the structure of a unique set of social networks. These three networks are social networks of alternative personalities described by a patient, “Patricia”, undergoing therapy for dissociative identity disorder, whose drawings of these networks are reproduced in David et al. (1996). One of the most striking findings in Martin (2017) is the presence in one of these networks of a large “hollow ring”, a cycle with no “shortcuts”, which is much larger than those expected under the random graph models considered. As discussed in Stivala (2020a) and recapitulated below, this “hollow ring” is in fact a *geodesic cycle* in graph theory terms.

In Martin (2017), the random graph null model used was the  $dk$ -series (Mahadevan et al., 2006; Orsini et al., 2015) family of graph distributions, and only the distribution of the largest geodesic cycle lengths is considered. Stivala (2020a) extends this work by using exponential random graph models (ERGM) (Lusher et al., 2013) as well as the  $dk$ -series random graphs, and also considering the geodesic cycle length distributions, in addition to the maximum geodesic cycle lengths. That work also examined a selection of other publicly available social networks, as well as the original “Patricia” networks, finding that, of the networks examined, only the “Patricia” networks seem to have geodesic cycle lengths that are not reproduced by random graph models.

Commentary on Stivala (2020a) mostly concerned the [mis-]interpretation of ERGM parameters and the trend (or otherwise) in the social networks field to be moving towards a monoculture of using parametric models, and in particular ERGMs, for everything (Martin, 2020; Stivala, 2020b). What has not yet been addressed, is to test the hypothesis suggested in Stivala (2020a), that fictional character networks (see Labatut and Bost (2019) for an overview), and specifically those with only a single author (unlike the one fictional character network examined in that work), might have anomalous geodesic cycle length distributions, like the “Patricia” networks.

The primary purpose of this work is to test this hypothesis, by examining how well the geodesic cycle length distributions in some fictional character networks — mostly from works by a single author — can be reproduced by random graph models, specifically the  $dk$ -series random graphs, and, in some cases, exponential random graph models.

<sup>1</sup>Università della Svizzera italiana, Via Giuseppe Buffi 13, 6900 Lugano, Switzerland. Email: alexander.stivala@usi.ch

Further, it was implicitly assumed in Stivala (2020a) that the method used there for counting geodesic cycles was exact. In fact, however, that method is only an approximation (more precisely, a lower bound). And so a second purpose of this work is to reproduce all of the results in Stivala (2020a) and compare with a method that counts geodesic cycles exactly.

## 2 Geodesic cycles

In this and the following sections, only undirected graphs are considered.

As discussed in Stivala (2020a), the structure that Martin (2017) describes as a “hollow ring” — a cycle with no “shortcuts”, so that the shortest path between any two nodes in the cycle is along the cycle itself — is known in graph theory as a *geodesic cycle* (Negami and Xu, 1986; Li and Shi, 2018). Another way of putting this is that a geodesic cycle is a cycle such that, for any two vertices on the cycle, the distance between them in the cycle is equal to their distance in the graph (Amaldi et al., 2009; Lokshantov, 2009).

As Mitchell (1974, p. 279) writes of the “proliferation of concepts and terms” in social network analysis, it is a “terminological jungle, in which any newcomer may plant a tree” (Barnes (1972, p. 8) as quoted in Mitchell (1974, p. 279)), and as noted in Stivala (2020a), a geodesic cycle has also been termed an *isometric cycle* (Lokshantov, 2009), and an *atomic cycle* (Gashler and Martinez, 2012). Adding yet another tree to the terminological jungle, an isometric cycle has also been termed a *short cycle* (Klemm and Stadler, 2006).

Unlike geodesic cycle, *chordless cycle* is a well-known standard term in graph theory, although there are conflicting conventions as to whether or not cycles of length three are included in the definition (Weisstein, 2023). A *chord* is an edge connecting two otherwise non-adjacent nodes in a cycle, and a chordless cycle is a cycle that has no chords. A chordless cycle is also known as an *induced cycle* or sometimes a *hole*, for chordless cycles of length at least four. Every geodesic cycle is chordless, but not every chordless cycle is geodesic. An illustrated example of these definitions is given in Stivala (2020a, Fig. 1).

In Stivala (2020a), the `find_large_atomic_cycle` algorithm (Gashler, 2011; Gashler and Martinez, 2012) was used to count geodesic cycles. However, despite the implicit assumption in that work that this algorithm enumerates geodesic cycles (that is, finds all of them), in fact, this is not guaranteed. The proof of correctness in Gashler and Martinez (2012) proves that any cycle it returns is geodesic, and that it will find at least one of the specified minimum length if it exists.<sup>1</sup> However, despite being described as “iterating over all the atomic cycles in the graph” and that “atomic cycles in the graph are enumerated” (Gashler and Martinez, 2012, p. 6), the proof does not show that all atomic (geodesic) cycles are enumerated (or iterated over), and in fact it is not guaranteed to do so. Hence it does not exactly count geodesic cycles, but only approximates the count. More precisely, it gives a lower bound on the number of geodesic cycles of each length. This is perfectly suitable for the purpose for which this algorithm was designed (finding such cycles and “cutting” them to improve the performance of a machine learning algorithm), but if we are going to count geodesic cycles, we would like an algorithm that can be guaranteed to do so exactly (even if only to get an idea of how close the, more easily obtained, lower bound is to the exactly value).

In the case of the three “Patricia” networks considered in Martin (2017); Stivala (2020a), in fact the counts are exact, which was verified by manual inspection. However this is not the case for all networks, as will be shown below (see Section 6.1 and Appendix B).

In this work, in order to count geodesic cycles exactly, an algorithm `isCycleGeodesic` is used. This algorithm tests if a given cycle in a graph is geodesic directly by the definition. That is, for every pair of nodes in the cycle, it tests if the distance between those nodes in the cycle is equal to the shortest path distance (geodesic) between them in the graph. It is a precondition of this function that the cycle is specified as a list of nodes in order along the cycle, so the cycle distance is efficiently computed as  $\min\{|i - j|, k - |i - j|\}$ , where  $i$  and  $j$  are the indices of the nodes in the cycle,  $i, j \in \{0 \dots k - 1\}$ , and  $k$  is the length of the cycle. The distance matrix (giving shortest path length for every pair of nodes) for the graph can be pre-computed, or, more efficiently, a single-source shortest paths algorithm can be used on demand, and the results memoized for reuse. This algorithm is implemented by the `filterCyclesGeodesic.py` Python script, which takes a graph and, on the input stream, a list of cycles in the graph, and writes to the output stream only those cycles from the input which are geodesic.

Note that, because all geodesic cycles are chordless, it is more efficient to give as input to this algorithm only chordless cycles, which can be enumerated more efficiently than having to enumerate all cycles (Uno and Satoh, 2014).

---

<sup>1</sup>However, as can be seen, from, for example, Fig. B6, where `find_large_atomic_cycle` does not find a geodesic cycle of length 7, four of which are found by the method described in this section, either this proof is not correct or there is an error in the implementation.

### 3 Exponential random graph models and $dk$ -series random graphs

As in Stivala (2020a), I use two different families of random graphs: exponential random graph models (ERGMs), and the  $dk$ -series family of random graphs.

ERGMs are a random graph model, widely used in the social sciences to model social networks (Lusher et al., 2013; Amati et al., 2018; Koskinen, 2020). An ERGM is a probability distribution over graphs, where the probability of a graph is a function of a set of graph statistics with their associated parameters. The statistics are counts of “configurations”, which represent local structures in the graph, such as edges, stars, or triangles. These statistics need not be purely structural, but can also include nodal attributes, in order to model homophily, for example. Given an observed graph, the parameters corresponding to the configurations in the model can be estimated by maximum likelihood, with their sign and significance indicating over- or under- representation, given all the other parameters in the model, of the corresponding configuration.

In this work, I do not make any interpretation of the estimated parameters, but simply use them to simulate a set of graphs from the estimated model, for use as a null model in comparing the observed geodesic cycle distributions to those generated from the model.

In contrast, the  $dk$ -series random graphs (Orsini et al., 2015) are a series of random graph null models of increasing complexity, forming a nested hierarchy. For each level of the hierarchy, an ensemble of random graphs are generated, for which particular statistics are fixed at the values in the observed graph. The simplest level of the hierarchy is  $0k$ , in which the density is fixed. This is just the Erdős–Rényi random graph model. This is followed by  $1k$ , in which the degree distribution is also fixed, again, a well-known random graph null model (Newman et al., 2001), used for “motif” detection in biological networks, for example (Milo et al., 2002). The next level is  $2k$ , which also fixes the joint degree distribution, thus reproducing the degree correlations (assortativity) of the observed graph. For practical use in generating ensembles of graphs that resemble empirical networks in relevant statistics, Orsini et al. (2015) define the next two distributions in the hierarchy as  $2.1k$  and  $2.5k$ . These are the  $2k$ -random graphs with fixed values of average local clustering and average local clustering by degree (clustering spectrum), respectively.

An advantage of the  $dk$ -series approach is that we do not need to contend with the potential difficulties in estimating ERGM parameters, where for some networks, such as most of the fictional character networks considered here, it can be difficult or perhaps even impossible (without adding a lot of extra constraints, such as fixing the presence or absence of particular ties) to find a converged model. Further discussion of the relative merits and shortcomings of these approaches can be found in Martin (2017); Stivala (2020a); Martin (2020); Stivala (2020b).

### 4 Data

The nine networks used in Stivala (2020a) include the three “Patricia” networks from David et al. (1996) originally used in Martin (2017), as well as six other networks from a variety of domains (Lind, 2012; Weissman, 2019; Leavitt and Clark, 2014; Lusseau et al., 2003; Lazega and Pattison, 1999; Lazega, 2001; Zachary, 1977; Kapferer, 1972; Mastrandrea et al., 2015), as described in Stivala (2020a).

The seven fictional character networks used in this work consist of four chapter co-occurrence networks from classic novels, one character interaction network from the *Star Wars* movie series, and two other types of character interaction networks from the *Harry Potter* books. The four chapter co-occurrence networks considered are from *Anna Karenina* by Leo Tolstoy, *David Copperfield* by Charles Dickens, *Adventures of Huckleberry Finn* by Mark Twain, and *Les Misérables* by Victor Hugo.

The *Anna Karenina*, *David Copperfield*, and *Huckleberry Finn* character interaction networks (Knuth, 1993) were downloaded from <https://people.sc.fsu.edu/~jburkardt/datasets/sgb/sgb.html> [accessed 1 September 2019]. Data in the Stanford GraphBase format (Knuth, 1993) was parsed using Python code adapted from the charnet project (Holanda et al., 2019) from <https://github.com/ajholanda/charnet>. These networks are illustrated in Figures A1, A2, and A3. The *Les Misérables* character interaction network (Knuth, 1993) was downloaded from <http://www-personal.umich.edu/~mejn/netdata/lesmis.zip> [accessed 10 July 2019]. This network is illustrated in Fig. A4.

The two *Harry Potter* networks are “Dumbledore’s Army”, and the peer support network for the first six *Harry Potter* books. The “Dumbledore’s Army” network (Everton et al., 2022) from the *Harry Potter* books was downloaded from [https://core-dna.netlify.app/publication/harry\\_potter\\_dumbledores\\_army](https://core-dna.netlify.app/publication/harry_potter_dumbledores_army) [accessed 13 January 2023]. A binary network was constructed by including only edges (representing trust or mutual understanding)

Table 1: Summary statistics of the networks.

Network	N	Components	Mean degree	Density	Clustering coefficient	Assortativity coefficient	Mean path length
Anna Karenina	138	1	7.14	0.05215	0.26827	-0.35010	2.45
David Copperfield	87	1	9.33	0.10853	0.35124	-0.25814	1.95
Huckleberry Finn	74	3	8.14	0.11144	0.48756	-0.17297	2.14
Les Misérables	77	1	6.60	0.08681	0.49893	-0.16523	2.64
Dumbledore’s Army	29	1	4.90	0.17488	0.54404	0.19288	2.57
Harry Potter	64	29	3.62	0.05754	0.52844	-0.19461	1.90
Star Wars	111	2	8.00	0.07273	0.35099	-0.20863	2.61

All networks are undirected. “Clustering coefficient” is the global clustering coefficient (transitivity), and “Assortativity coefficient” is the degree assortativity. All statistics were computed using the `igraph` (Csárdi and Nepusz, 2006) library in R (R Core Team, 2016).

with a tie strength of 4 or 5, resulting in the network shown in Everton et al. (2022, Fig. 2) and Fig. A5.

The Harry Potter peer support network (Bossart and Meidert, 2013) for the first six Harry Potter books was downloaded from <http://www.stats.ox.ac.uk/~snijders/siena/HarryPotterData.html> [accessed 12 July 2019]. The network is directed, with a directed edge  $(i, j)$  indicating that character  $i$  gives peer support to character  $j$ . The network was converted to undirected by changing any directed edge to an undirected edge (mutual directed edges result in a single undirected edge). This network also contains node attributes for school year, gender, and wizarding house, and is illustrated in Fig. A6.

The “merged” Star Wars character interaction network was used, meaning Anakin and Darth Vader are the same node (Gabasova, 2015, 2016). This network is illustrated in Fig. A7.

Summary statistics of these networks are shown in Table 1.

## 5 Methods

ERGM models were estimated, and networks simulated from the models, with the `ergm` R package in the `statnet` software suite (Handcock et al., 2008; Hunter et al., 2008; Hummel et al., 2012; Handcock et al., 2016, 2022; Krivitsky et al., 2023). Networks were simulated from the  $dk$ -series random graph models using the `RandNetGen` software<sup>2</sup> (Mahadevan et al., 2006; Colomer-de Simón et al., 2013; Colomer-de Simón and Boguñá, 2014; Orsini et al., 2015). The `RandNetGen` program does not simulate graphs from the  $0k$  distribution (Erdős–Rényi random graphs), so, as in Martin (2017); Stivala (2020a), the `igraph` R package (Csárdi and Nepusz, 2006) was used to simulate these networks.

For each ERGM model or  $dk$ -series distribution, 100 networks are simulated, and the box plots, generated with the `ggplot2` R package (Wickham, 2016) represent data from these 100 networks.

For the networks considered in Stivala (2020a), the same sets of simulated networks from that work were used. Further details of the ERGM models in that work can be found in Stivala (2020a,b).

To count geodesic cycles in a network, the `CYPATH` software<sup>3</sup> (Uno and Satoh, 2014) was used to generate all of the chordless cycles, which were then given as input to the `filterCyclesGeodesic.py` script (see Section 2), implemented in Python with the `igraph` library, which outputs only those cycles that are geodesic.

To find the longest geodesic cycle in a graph, Algorithm 4.1, “LIC – Longest Isometric Cycle” from Lokshtanov (2009) was implemented in Python in the `longestIsometricCycle.py` script. It is important to note that, as described in Catrina et al. (2021), Lemma 3.6 of Lokshtanov (2009) is erroneous, and that therefore the algorithm is only correct for even-length cycles. For odd-length cycles, the conditions of this Lemma may be met for a particular cycle of length  $k$  (for an odd value of  $k$ ), and yet there is no cycle of length  $k$  in the graph, and therefore the algorithm will incorrectly find that there is a geodesic cycle of length  $k$ . There is an example showing this case in Catrina et al. (2021, Fig. 5), and it also occurs, for example, in the “Patricia” graph for 1990, in which the conditions of Lokshtanov (2009, Lemma 3.6) are met for a geodesic cycle of length 11, and yet there is no such cycle: the longest geodesic cycle is of length 10 (see Martin (2017); Stivala (2020a), and Fig. B1). In order to have the algorithm work correctly for both

<sup>2</sup><https://github.com/polcolomer/RandNetGen>

<sup>3</sup><http://research.nii.ac.jp/~uno/code/cypath.html>

odd and even values of  $k$ , the auxiliary bipartite graph construction described by Catrina et al. (2021, Observation 5.4) was implemented in `longestIsometricCycle.py`. It is an assumption of this algorithm that the input graph is connected, so the implementation handles graphs that are not connected by running the algorithm on each connected component, and returning the size of the largest geodesic cycle found in any component. This is the longest geodesic cycle in the graph, since any cycle must be entirely contained within a single connected component.

For each set of 100 simulated networks, the script to count geodesic cycles was run with an elapsed time limit of 48 hours, on an Intel Xeon E5-2650 v3 2.30GHz processor on a Linux compute cluster node. The iteration over the 100 simulated networks was done serially, that is, one network at a time. (This serial counting process for each set of 100 networks was done in parallel over the networks simulated from different distributions, however). If the time limit was reached, then only the results for the networks for which the counting could be completed within the time limit are included. In such cases, where this number was less than 100, this is noted in the relevant figure caption. In some cases, for the  $0k$ -distribution, the count cannot be completed within the time limit for even the first network in the set of simulated networks, and so there are no results, and the results for this distribution are omitted.

The lengths of the largest geodesic cycles, as shown in the plots in Section 6.2 and Appendix B, are simply the largest geodesic cycles in the distributions, computed as described above, that have a nonzero count. In order to further verify that these are correct, the `longestIsometricCycle.py` script described above was used to find the length of the longest geodesic (i.e. isometric) cycle in each network, and I verified that these were the same as those found from `filterCyclesGeodesic.py`

## 6 Results and discussion

### 6.1 Reproducing earlier approximate results with exact counts

In Appendix B, the results from Stivala (2020a), where geodesic cycles were counted only approximately (lower bound) using the `find_large_atomic_cycle` algorithm (Gashler, 2011; Gashler and Martinez, 2012) are plotted on the same plots as those where the geodesic cycles are counted exactly, as described Section 2. For Patricia’s 1990 (Fig. B1) and 1992 (Fig. B2) networks, the results are not any different from those described in Stivala (2020a): in the 1990 network, the largest geodesic cycle (length 10) is extremely unlikely under any of the models (ERGM or  $dk$ -series), while in the 1992 network, that the largest geodesic cycle is only of length 4 is extremely unlikely under any of the models (in this case, larger maximum geodesic cycle lengths are expected). The results for the geodesic cycle length distributions are also unchanged, with the fit being poor for the 1990 and 1992 networks, but acceptable, in the case of the ERGM only, for the 1993 network (Fig. B3). The situation for maximum geodesic cycle lengths in the 1993 network is slightly different, with the exact counting showing that the largest geodesic cycle (length 8) is more clearly unlikely than is apparent when the approximate (lower bound) method was used. Along with Fig. B4, this shows that, just as discussed in Stivala (2020a), the ERGM shows a better fit to the geodesic cycle length distribution in this network than any of the  $dk$ -series distributions.

For all three “Patricia” networks, the observed maximum geodesic cycle length is found to be the same by both methods; indeed the fact that the `find_large_atomic_cycle` method found the correct largest geodesic cycle lengths was verified manually on these networks.

Figures B5 – B10 show the results for the other networks considered in Stivala (2020a). When using the approximate (lower bound) counting method, the ERGM shows a good fit for the largest geodesic cycle size; however when using the exact counting method, this is not the case for the dolphin social network (Fig. B6) or the Kapferer tailor shop network (Fig. B9). In Stivala (2020a, p. 53) it was stated that the ERGM shows an acceptable fit to the geodesic cycle length distributions for all the non-“Patricia” networks considered, as does the  $dk$ -series  $2.5k$  distribution (the latter with the exception of the Grey’s Anatomy and high school friendship networks). However when using the exact counting method, we can see that this is not true for the dolphin social network (Fig. B6), the Lazega law firm network (Fig. B7), the Zachary karate club network (Fig. B8), and the Kapferer tailor shop network (Fig. B9). In the case of the dolphin social network, the Zachary karate club network, and the Kapferer tailor shop network, the  $dk$ -series  $2.5k$  distribution also shows a poor fit to the geodesic cycle length distribution when using the exact counting method. However in the case of the Lazega law firm network, the  $dk$ -series  $2.5k$  distribution shows an acceptable fit to the geodesic cycle length distribution when using the exact counting method, while the ERGM does not.

In the Grey’s Anatomy network, the exact counting method shows the same results discussed in Stivala (2020a): both the ERGM and  $2.5k$  distributions fit the maximum geodesic cycle length well, but the ERGM also fits the geodesic




Figure 1: Largest clique sizes for the Dumbledore’s Army network. The dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series or ERGM as labeled on the  $x$ -axis. Largest clique sizes were computed with the `igraph` R package (Csárdi and Nepusz, 2006).

cycle length distribution well, while the  $2.5k$  distribution does not (Fig. B5). As noted in Stivala (2020a, p. 53), this is because the ERGM can reproduce the lack of odd-length cycles in this network by using node attribute information (most relationships in this network are heterosexual), while the  $dk$ -series distributions cannot. It is, however, stated in Stivala (2020a, p. 60) that of the seven four-cycles in this network, of which all seven are also chordless, only five are also geodesic. This is not correct: exact counting of geodesic cycles show that all seven are geodesic (Fig. B5).

Just as discussed in Stivala (2020a), in the high school friendship network (Fig. B10), the ERGM fits the maximum geodesic cycle length distribution better than the  $2.5k$  distribution does.

## 6.2 Fictional character networks

As discussed in Stivala (2020a), it can be difficult to estimate ERGM parameters for character interaction networks, not least because of the presence of a main character, often linked to most of the other characters, forming a high degree “hub” node. As noted in Stivala (2020a, p. 60), I was unable to find good converged ERGM for the character interaction networks for *Les Misérables*, *David Copperfield*, *Anna Karenina*, *Huckleberry Finn*, and hence only  $dk$ -series distributions are used here for these networks. This is also the case for the *Star Wars* network.

ERGM models for the two Harry Potter networks (Dumbledore’s Army and the Harry Potter peer support network) are shown, along with goodness-of-fit plots, in Appendix C. Parameter estimates of a simple model of the Dumbledore’s Army network are shown in Table C1. Note that, despite the geometrically weighted edge-wise shared partners (`gwesp`) term being included in the model, the model fit to the edge-wise shared partners distribution is not very good; in particular it fails to reproduce the pronounced peak in the distribution at 6, or the peak in the degree distribution at 9 (Fig. C1). This could well be because of the difficulty of an ERGM fitting the clique of size 8 clearly visible in Fig. A5 — a clique that includes, as we would expect, Harry Potter himself (Everton et al., 2022, Fig. 2). Fig. 1 shows that, indeed, the ERGM does not reproduce the large (8 node) clique in the observed network; only the  $2.5k$  distribution is close to doing so.

Parameter estimates for an ERGM for the Harry Potter peer support network are shown in Table C2. I was unable to find a converged model when the *gwesp* or *gwdsp* (geometrically weighted dyad-wise shared partners) terms were included, and so this model is a “dyad-independent” model (Koskinen and Daraganova, 2013; Krivitsky et al., 2023), as it contains only the edge parameter, and nodal attribute parameters relating to a single node (*nodefactor*) or dyad (*nodematch*). Unsurprisingly, given the lack of a model term (such as *gwesp*) to account for transitivity, the goodness-of-fit plot shows a poor fit to the edge-wise shared partner distribution (Fig. C2).

In attempting to use graph statistics of character networks for book genre classification, Holanda et al. (2019) find that all of the networks they examine are disassortative, meaning that characters with high degree interact preferentially with those of low degree. They therefore concluded that assortativity could not distinguish these genres. The works they examined comprised three genres: biographical, legendary, and fiction (the four books they classified as fiction included two of those examined in this work: *Huckleberry Finn* and *David Copperfield*). Here we find, similarly, that six of the seven networks (all of which are fiction) are disassortative (Table 1). The exception, with positive assortativity coefficient, is the *Dumbledore’s Army* network from the Harry Potter books. This network is notably different from the others in that it is neither a chapter co-occurrence nor character interaction (or narrative peer support) network, but a network representing trust or mutual understanding, between only a subset of the characters. Specifically it is an informal “dark network” (Cunningham et al., 2016) of a resistance movement against the Death Eaters (Everton et al., 2022). The relatively large clique in this network contributes to it having a positive assortativity.

Figures 2 – 8 show the largest geodesic cycle size and geodesic cycle size distributions for the fictional character networks considered in this work. For the four classic novels *Anna Karenina* (Fig. 2), *David Copperfield* (Fig. 3), *Huckleberry Finn* (Fig. 4), and *Les Misérables* (Fig. 5), it can be seen that the  $2.5k$  distribution fits both the largest geodesic cycle size and the geodesic cycle size distribution well, as does the  $2.1k$  distribution in most cases (it does not fit the geodesic cycle length distribution for the *Anna Karenina* network, Fig. 2, so well).

For the *Dumbledore’s Army* network (Fig. 6), none of the distributions fit the largest geodesic cycle size particularly well (although we might consider the fit for  $2.1k$  acceptable, and the distributions are not as far from the observed value as in the “Patricia” 1990 and 1992 networks). The  $dk$ -series  $2.5k$  distribution fits the observed geodesic cycle length distribution acceptably, while the fit for the ERGM is not quite as good (having a worse fit to geodesic cycles of length 5, for example). So the ERGM model does not do as badly at reproducing the geodesic cycle length distribution as it does for the largest clique size (Fig. 1), relative to the  $dk$ -series  $2.5k$  distribution.

For the Harry Potter peer support network (Fig. 7), the  $dk$ -series  $2.5k$  distribution fits the geodesic cycle length distribution very well, while the ERGM, particularly for geodesic cycles of length 5 or less, fits it extremely poorly. The ERGM also does not reproduce the largest geodesic cycle size, while the  $2.5k$  and  $2.1k$  distributions do. As already discussed, the ERGM model for this network is a dyad-independent model, and does not fit the edge-wise shared partner distribution well. So in this case we can see that distributions that reproduce transitivity, specifically  $2.1k$  (average local clustering) and  $2.5k$  (clustering by degree) also reproduce the geodesic cycle length distributions. However distributions that do not reproduce the transitivity of the observed network, such as  $0k$ ,  $1k$ ,  $2k$ , and the dyad-independent ERGM, also do not reproduce the geodesic cycle length distributions.

For the Star Wars character interaction network (Fig. 8), both the  $2.1k$  and  $2.5k$  distributions fit the maximum geodesic cycle length well, and also fit the geodesic cycle length distribution well, with the exception that they both produce significantly more geodesic cycles of length 5 than are present in the observed network.

## 7 Conclusions and future work

Having compared the results in Stivala (2020a) with the results when counting geodesic cycles exactly instead of with the lower bound approximation used originally, we see that the results for the three “Patricia” networks are unchanged. However for some of the other networks considered, what seemed like good fits to the geodesic cycle length distributions for both the  $dk$ -series  $2.5k$  and ERGM distributions are actually not very good fits for either model when the exact counting method is used. In one case (the *Lazega law firm* network), the  $2.5k$  distribution has a good fit to the geodesic cycle length distribution while the ERGM does not. The essential conclusion remains, however, that none of the networks examined have maximum geodesic cycle lengths that are extremely unlikely under any of the random graph distributions considered.

The hypothesis suggested in Stivala (2020a), that fictional character networks, and specifically those from works with a single author, might have anomalous geodesic cycle length distributions, like the “Patricia” networks, is not supported on any of the seven fictional character networks examined in this work (six of which are from single author

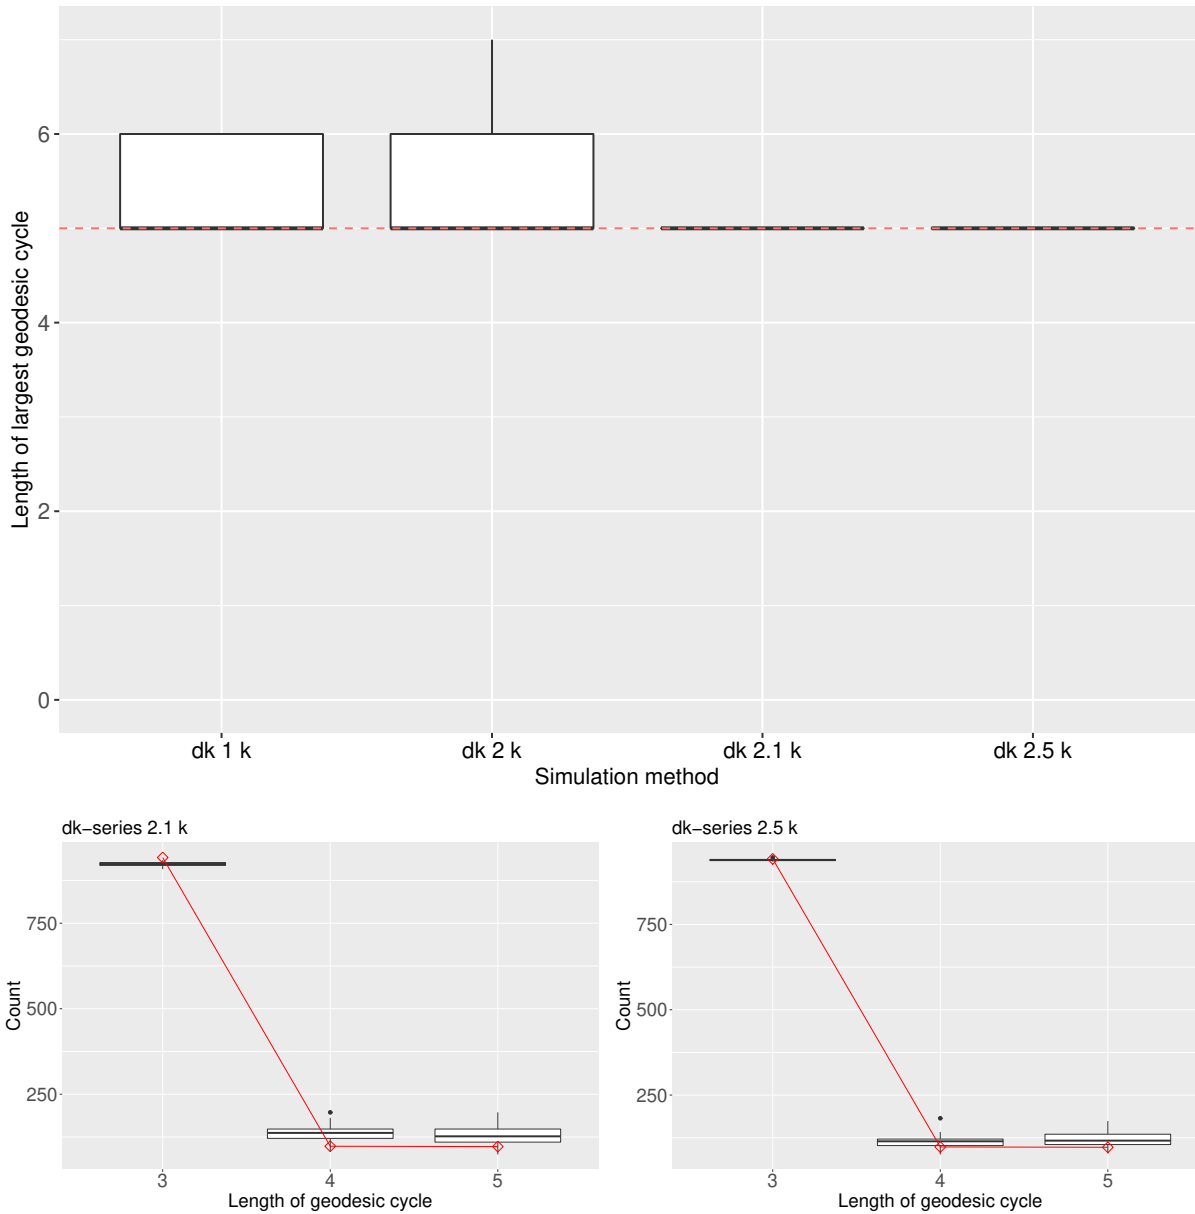


Figure 2: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Anna Karenina network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series 2.1 $k$  distribution (left) and the 2.5 $k$  distribution (right).



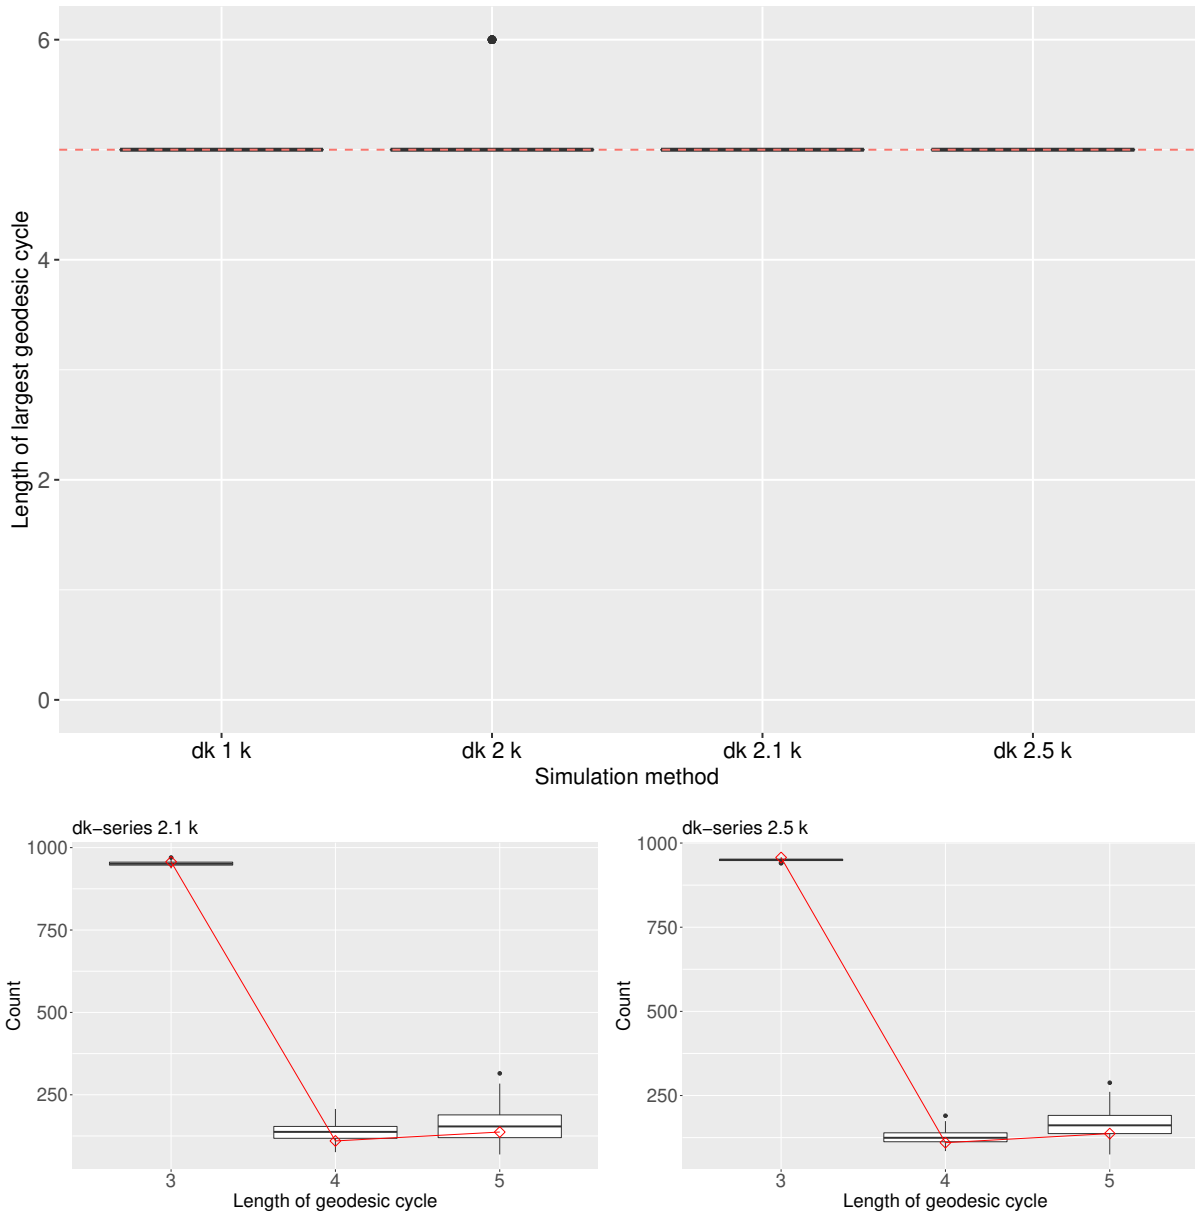


Figure 3: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the David Copperfield network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series 2.1 $k$  distribution (left) and the 2.5 $k$  distribution (right).

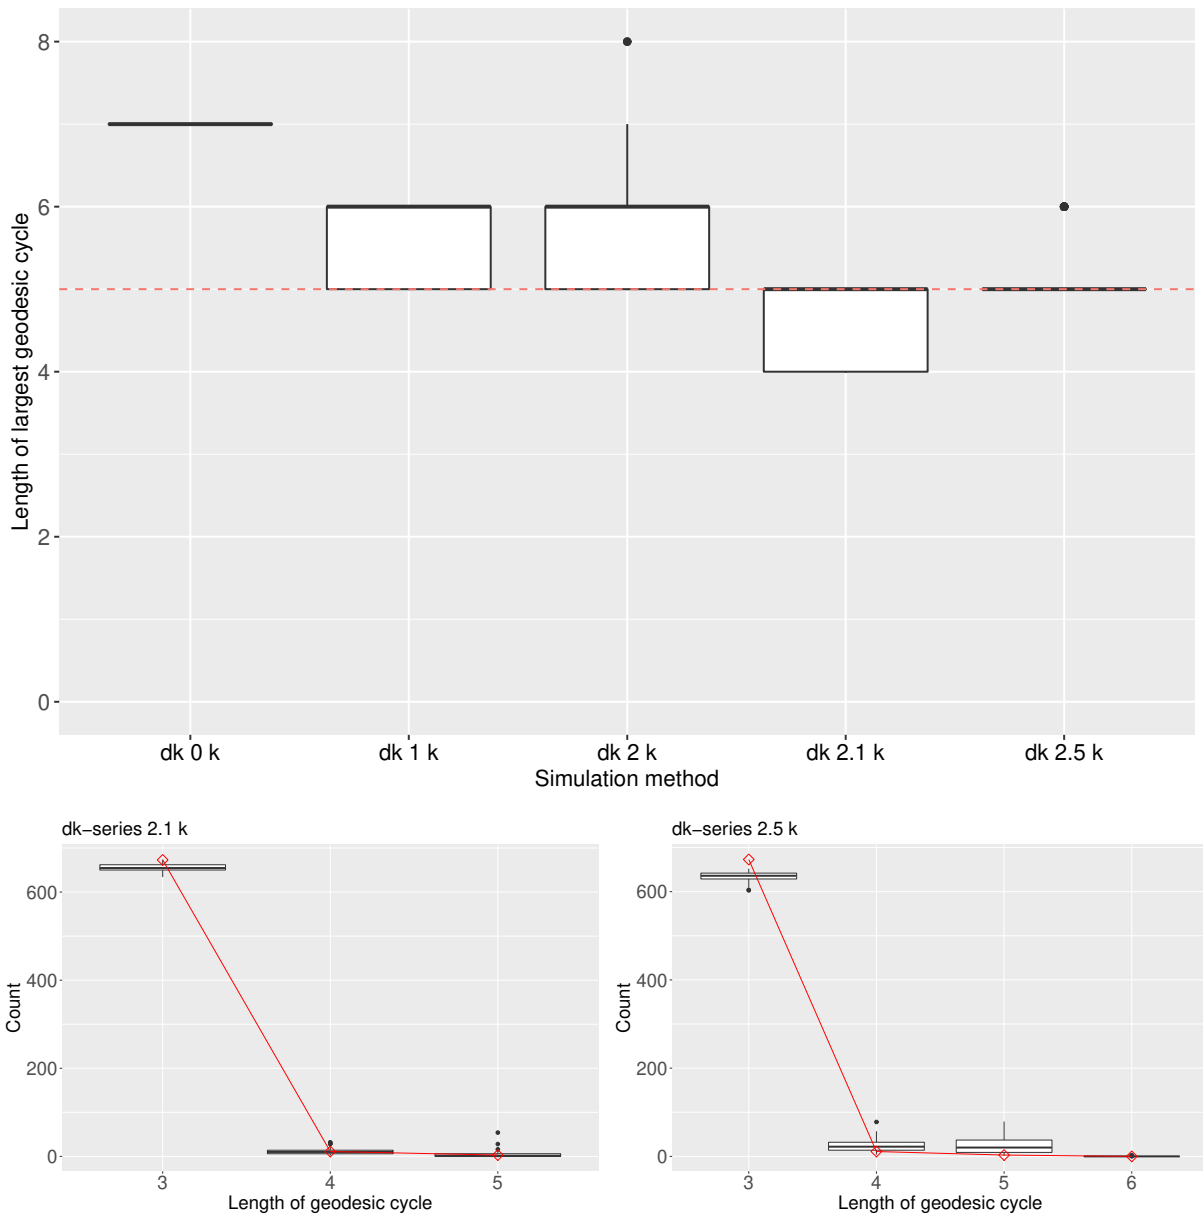


Figure 4: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Huckleberry Finn network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks (but only 16 networks for the  $0k$  distribution) simulated from the  $dk$ -series as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series  $2.1k$  distribution (left) and the  $2.5k$  distribution (right).

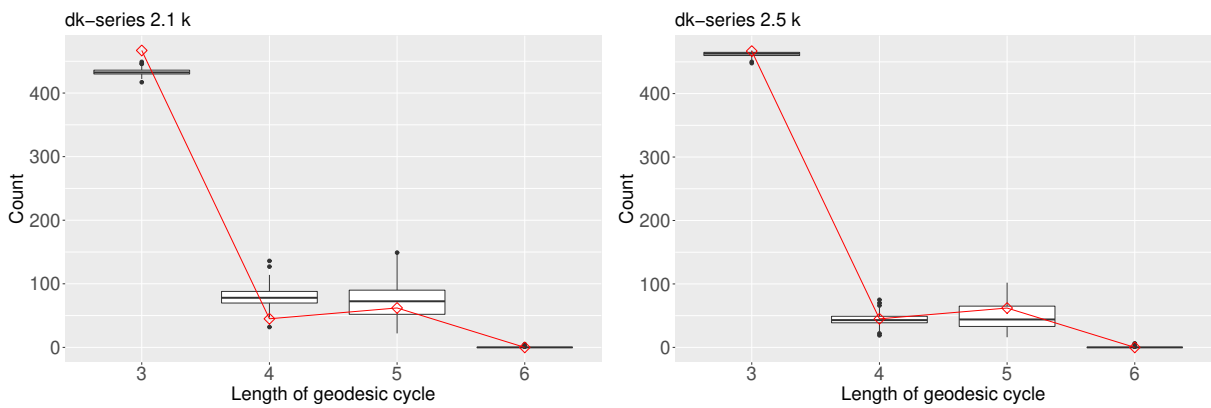



Figure 5: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Les Misérables network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks (but only 11 networks for the  $0k$  distribution) simulated from the  $dk$ -series as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series  $2.1k$  distribution (left) and the  $2.5k$  distribution (right).

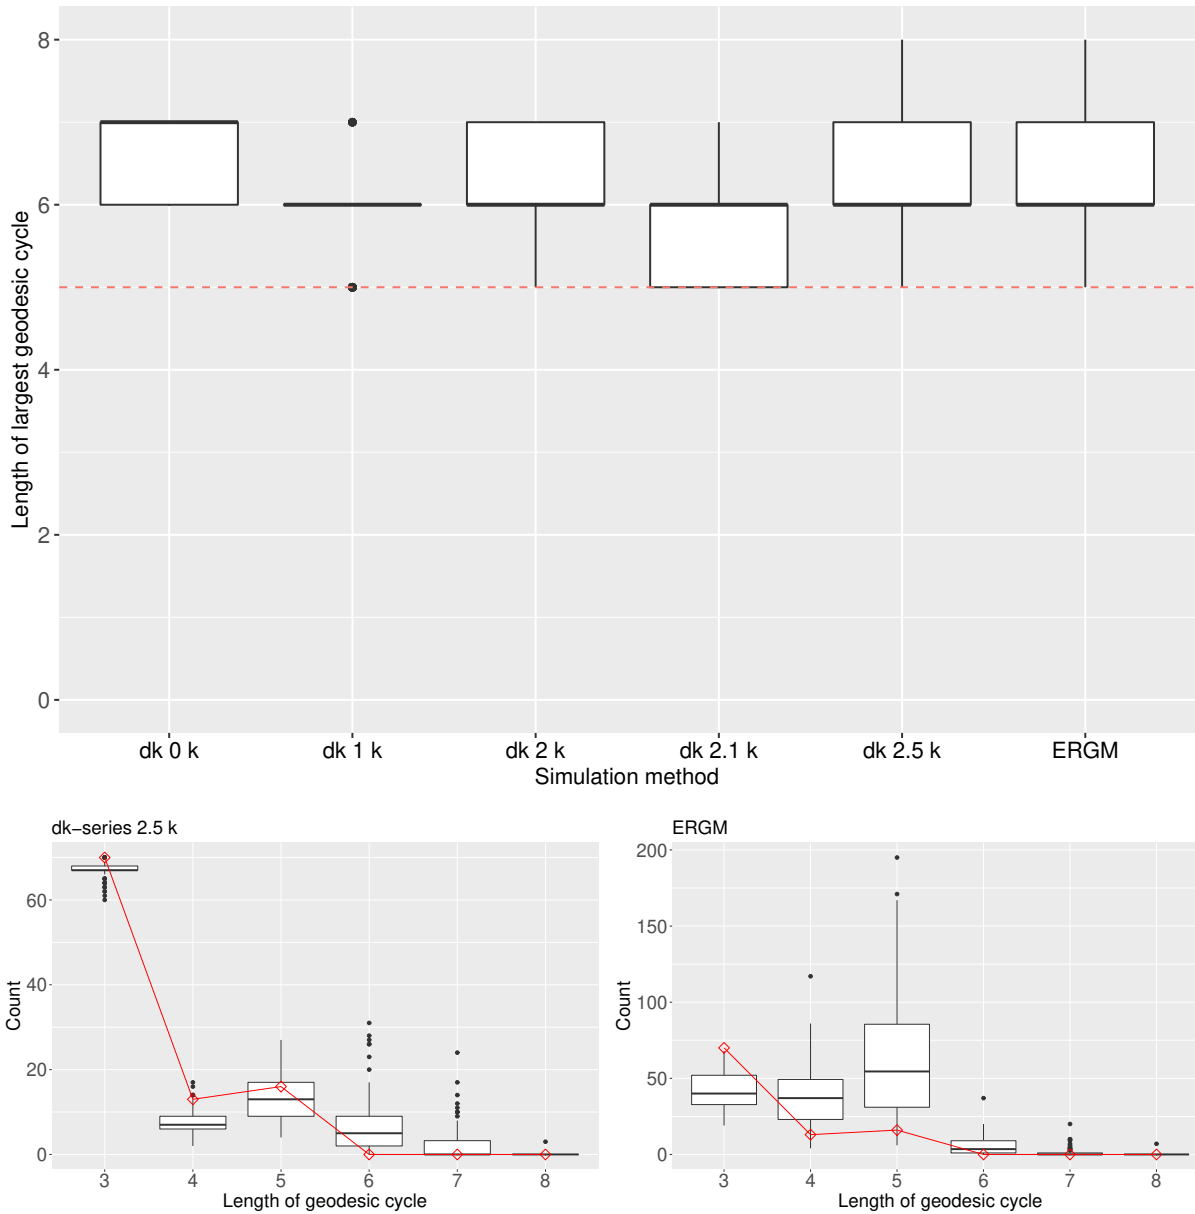


Figure 6: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Dumbledore's Army network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series or ERGM as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series 2.5 $k$  distribution (left) and the ERGM (right).

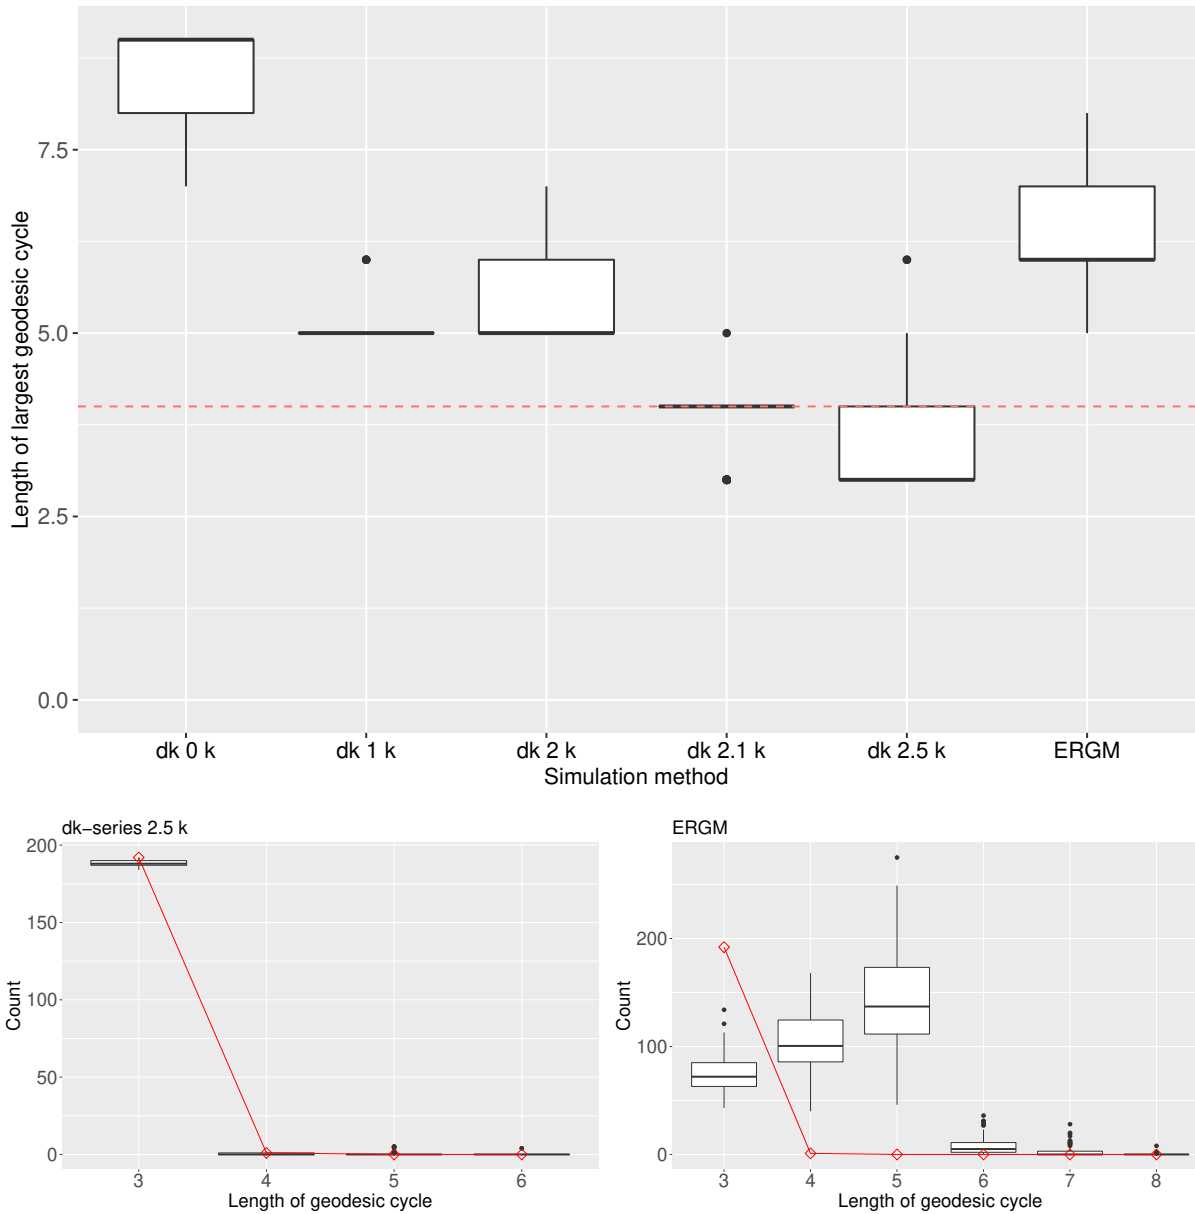


Figure 7: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Harry Potter peer support network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series or ERGM as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series 2.5k distribution (left) and the ERGM (right).

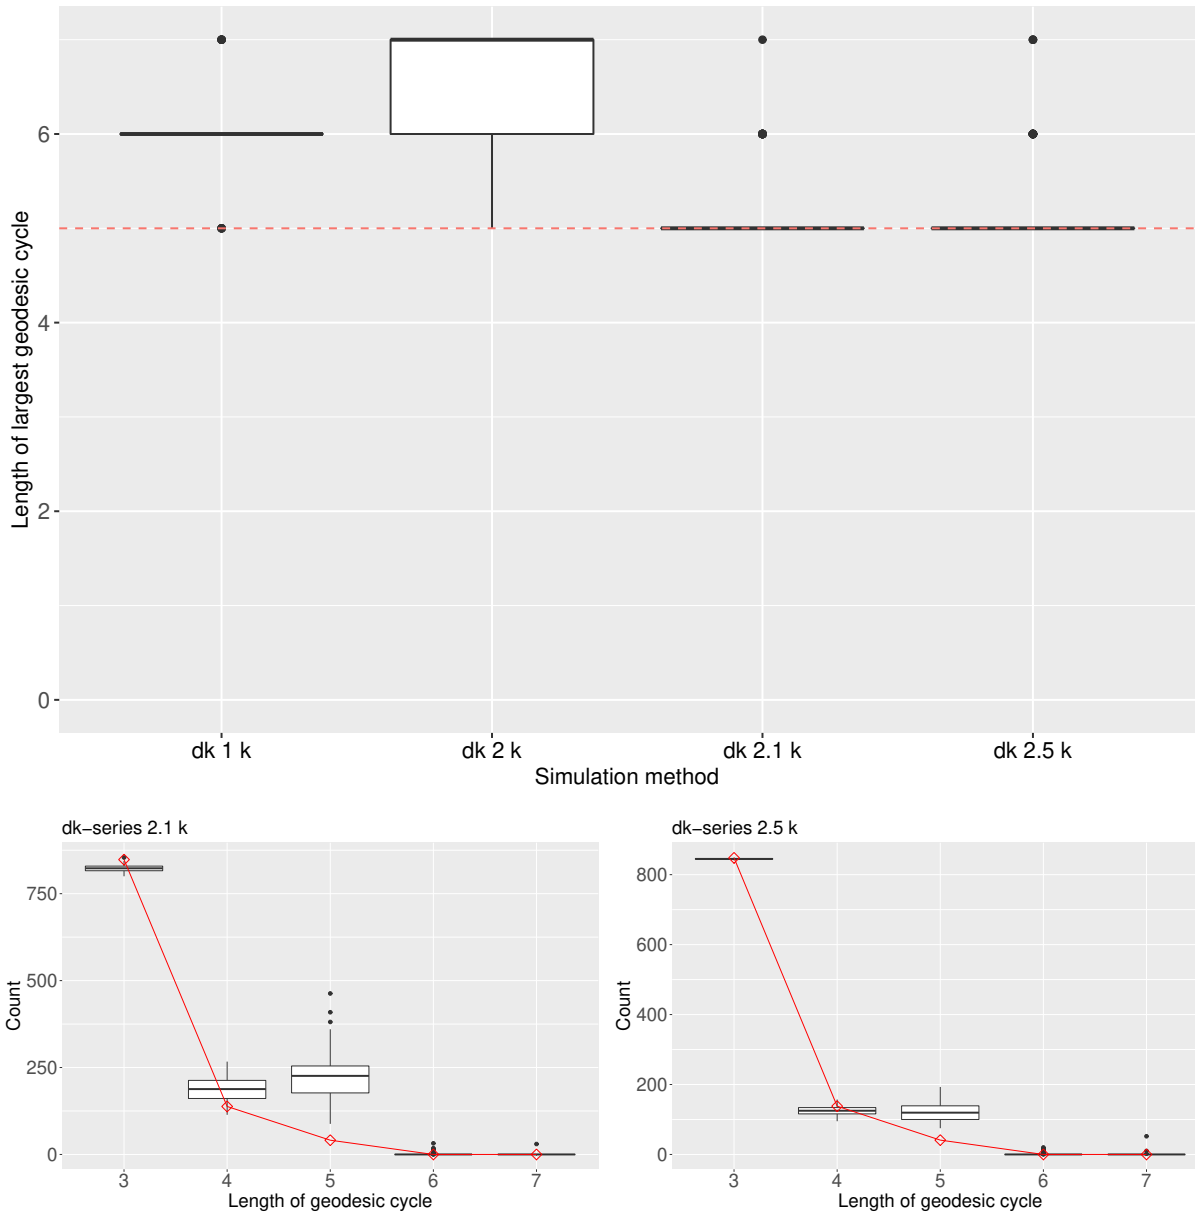


Figure 8: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Star Wars network. In the top plot, the dashed red line is the value in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series or ERGM as labeled on the  $x$ -axis. In the bottom plots, the points shown as red diamonds joined by the red line are the values in the observed network, with the box plots showing the values in 100 networks simulated from the  $dk$ -series 2.1 $k$  distribution (left) and the 2.5 $k$  distribution (right).

works). In all cases, at least one of the  $dk$ -series models or ERGM reproduces the geodesic cycle length distributions acceptably well. Hence it seems that there really is something special about the “Patricia” networks in this regard, which is not shared by fictional character networks. Just as with the other networks considered in Stivala (2020a), it seems that fictional character networks, even those from works by a single author, have a macro-level structure, the geodesic cycle length distribution, which can arise naturally from micro-level structures modeled by ERGM or fit exactly by the  $dk$ -series random graphs.

The method used here to count geodesic cycles exactly by reading an enumeration of all chordless cycles and testing whether or not each one is geodesic, is rather unsatisfactory in that it involves the enumeration of a potentially very large number of candidate chordless cycles, the vast majority of which are not geodesic, in most cases. It would obviously be preferable to directly enumerate geodesic cycles. Unlike the problem of enumerating chordless cycles, there appears to be no published work specifically on enumerating geodesic cycles. For example, the former is listed in an enumeration of enumeration algorithms (Wasa, 2016), and the latter is not. The algorithm in Lokshtanov (2009) for finding the length of the longest geodesic cycle, while polynomial in the size of the input graph, does not appear to be suitable for enumerating all geodesic cycles, and was rather designed to prove the (rather surprising) fact that finding the longest geodesic cycle is in  $\mathcal{P}$ , unlike most variants of the longest cycle problem, such as finding the longest cycle or longest chordless cycle, which are  $\mathcal{NP}$ -complete (Lokshtanov, 2009).<sup>4</sup> Hence the former is computationally tractable, and the latter are not (unless  $\mathcal{P} = \mathcal{NP}$ ).

In the absence of a more elegant and efficient algorithm, and its implementation, for directly enumerating geodesic cycles, the method used here could potentially be made more efficient (if not more elegant or satisfactory from a theoretical point of view) simply by using a more efficient method of enumerating chordless cycles. The Uno and Satoh (2014) algorithm for enumerating chordless cycles was used in this work, not least because it is the only one for which a publicly available (and efficiently implemented, in the C programming language) implementation appears to exist. Although this algorithm has been described as the “most notable and elegant listing algorithm” (Ferreira et al., 2014, p. 419) for this problem, at least one algorithm with better asymptotic complexity has been described (Ferreira et al., 2014). The enumeration of enumeration algorithms (Wasa, 2016) lists in addition the algorithm of Wild (2008), although this has no known guaranteed performance (Ferreira et al., 2014), and neither does the algorithm of Sokhn et al. (2013), although Dias et al. (2014) claims to have an improved algorithm, with a parallel implementation for GPU (graphics processing unit) described in Jradi et al. (2015).

Although there appear to be no published papers specifically on enumerating geodesic cycles, one set of papers addresses this issue as part of another problem. The work of Amaldi et al. (2009, 2010, 2011), improves on the Horton (1987) algorithm for finding a minimum cycle basis of a graph, by restricting the candidates to isometric (geodesic) cycles. Amaldi et al. (2009) describes “an efficient  $O(nm)$  procedure which allows to detect a single representation of each isometric cycle without explicitly constructing the non-isometric cycles” (Amaldi et al., 2010, p. 400). This suggests that this procedure from Amaldi et al. (2009) could in fact be used to efficiently enumerate all geodesic cycles. The creation of a publicly available implementation of such an algorithm could greatly increase the practicality (particularly for large graphs) of examining geodesic cycle length distributions of empirical and random graphs for testing hypotheses about geodesic cycles or for use as an additional goodness-of-fit test for random graph models.

## Acknowledgements

I used the high performance computing cluster at the Institute of Computing, Università della Svizzera italiana, for all data processing and computations.

## Data availability statement

All code, scripts, and data are available from [https://github.com/stivalaa/geodesic\\_cycles](https://github.com/stivalaa/geodesic_cycles).

---

<sup>4</sup>As noted in Catrina et al. (2021), there is an error in part of this proof, in the case of odd-length cycles, but the proof can be corrected using the even-length case with an auxiliary graph construction.

## Appendix A Visualizations of networks




Figure A1: Anna Karenina character interaction network. Visualization created using the network R package (Butts, 2008, 2015).






Figure A2: David Copperfield character interaction network. Visualization created using the network R package (Butts, 2008, 2015).




Figure A3: Huckleberry Finn character interaction network. Visualization created using the network R package (Butts, 2008, 2015).




Figure A4: Les Misérables character interaction network. Visualization created using the network R package (Butts, 2008, 2015).




Figure A5: Dumbledore's Army network. Visualization created using the network R package (Butts, 2008, 2015).



Figure A6: Harry Potter peer support network. Visualization created using the network R package (Butts, 2008, 2015). Nodes are coloured according to the house the character is in.




Figure A7: Star Wars character interaction network. Visualization created using the network R package (Butts, 2008, 2015).

## Appendix B Geodesic cycle length distributions for networks from Stivala (2020a)

The plots in this section reproduce, on the same plots, the results from Stivala (2020a), where the geodesic cycles were counted only approximately (lower bound) using the `find_large_atomic_cycle` algorithm (Gashler, 2011; Gashler and Martinez, 2012), labeled `findAtomicCycles` on the plots, along with the results counting geodesic cycles exactly, labeled `countGeodesicCycles` on the plots. In the plots of largest geodesic cycle lengths (top plot on figures), the observed largest geodesic cycle lengths are shown as horizontal lines, while in the plots of geodesic cycle length distributions (bottom plots on figures), the observed values are shown as diamonds (joined by lines as a visual aid only). The box plots represent values from 100 (except where otherwise noted) networks simulated from the distributions as labeled on the plots.

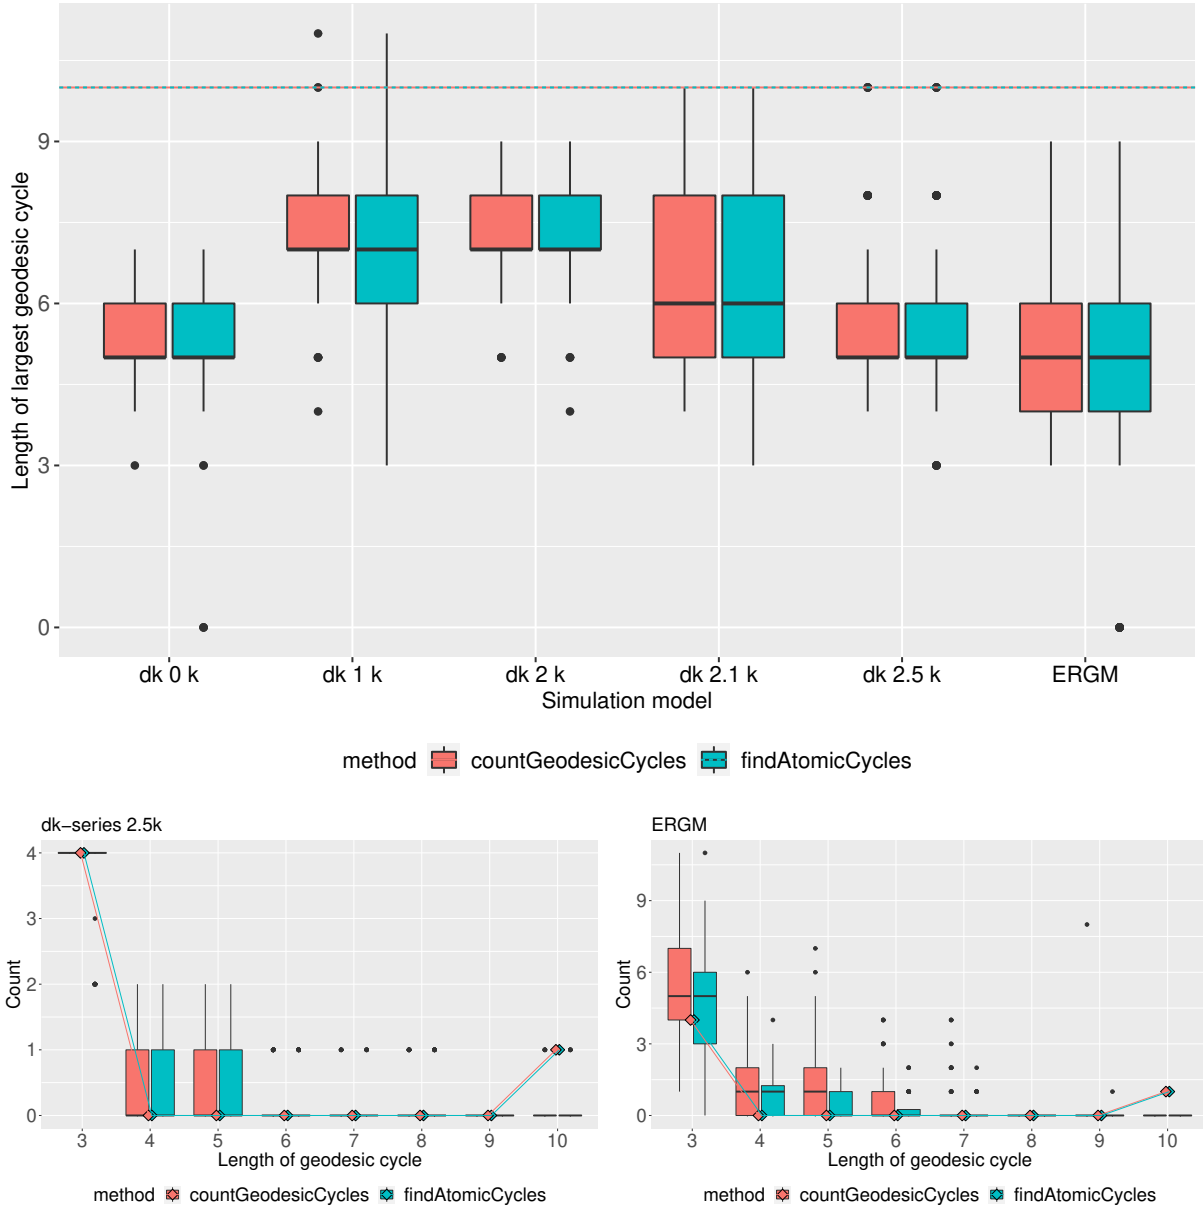


Figure B1: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for Patricia’s 1990 network, corresponding to Fig. 5 in Stivala (2020a). The bottom two plots show results for the *dk*-series 2.5*k* distribution (left) and from the ERGM (right).

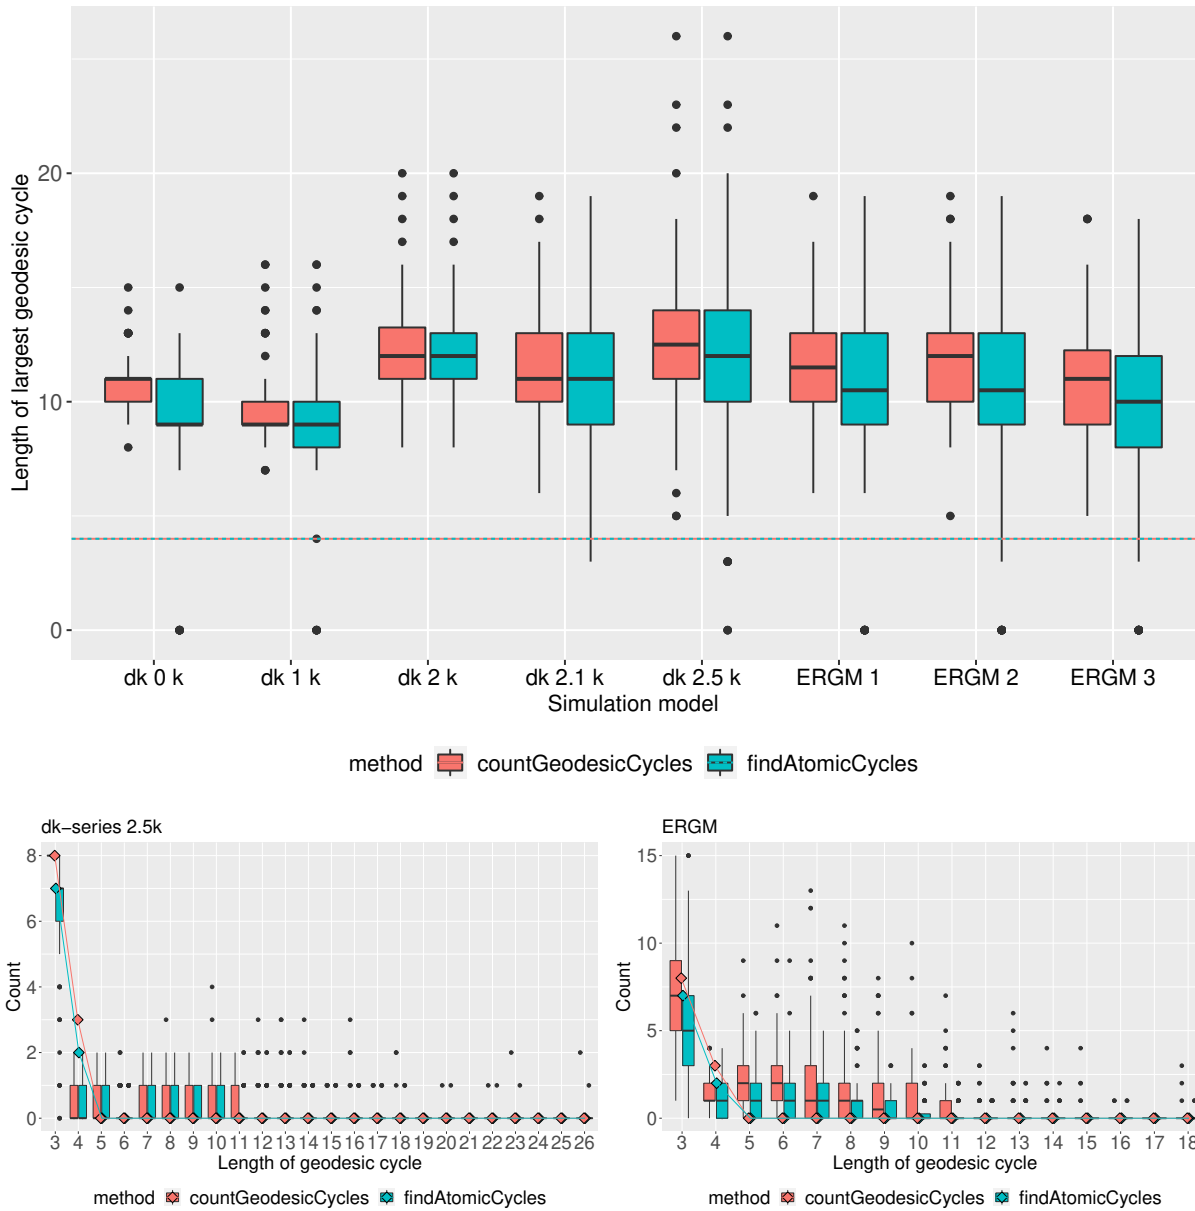


Figure B2: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for Patricia's 1992 network, corresponding to Fig. 6 in Stivala (2020a). The bottom two plots show results for the *dk*-series 2.5k distribution (left) and from the ERGM (right).

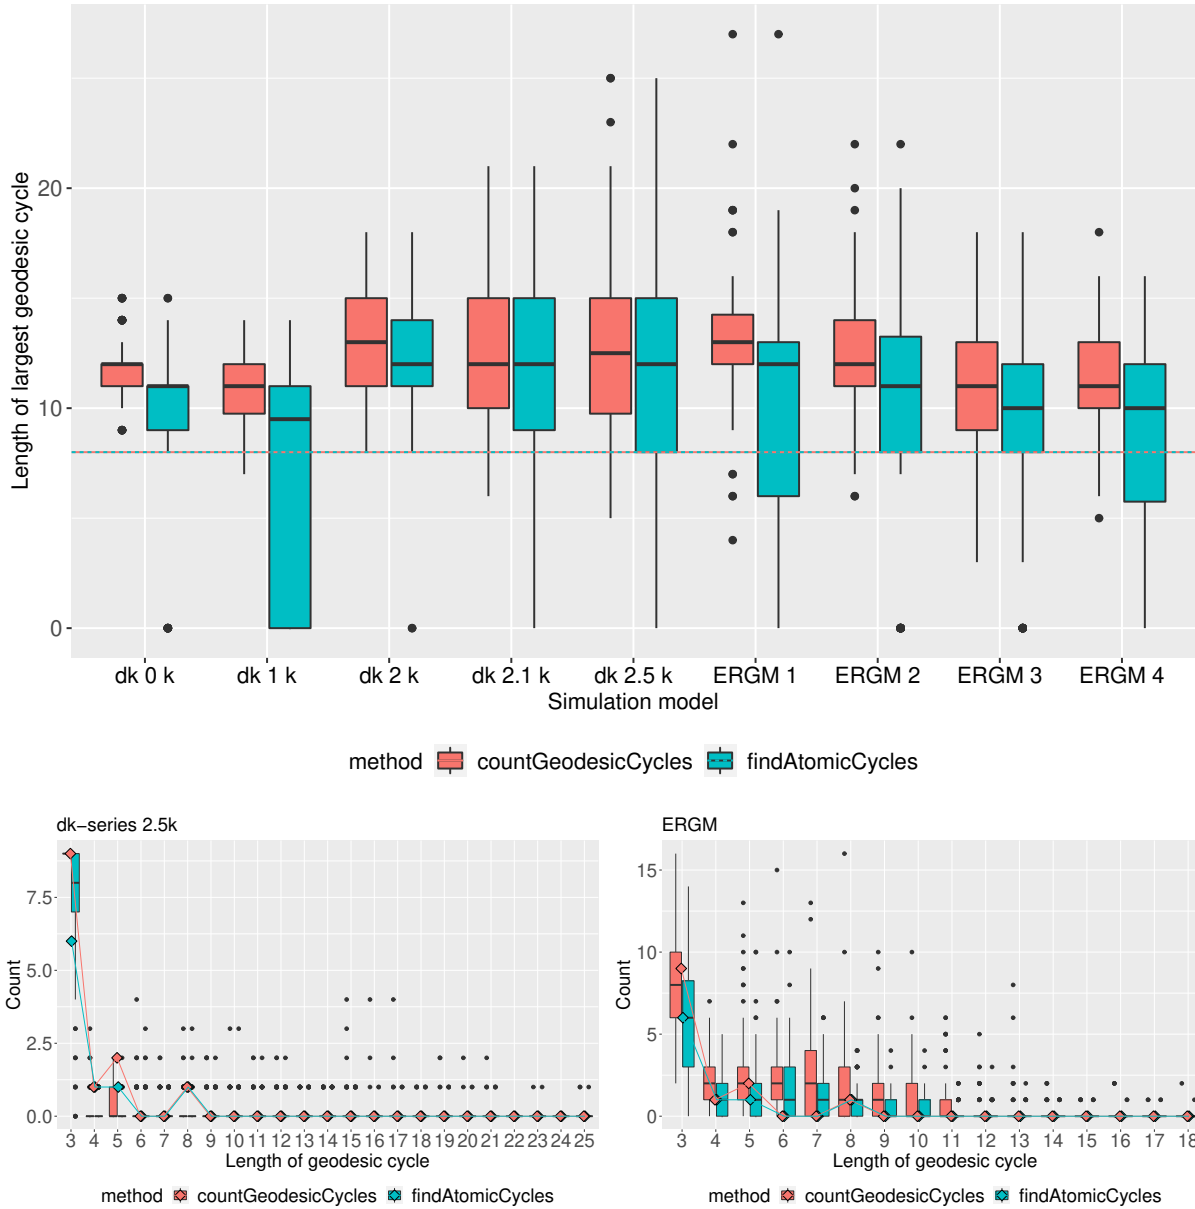


Figure B3: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for Patricia's 1993 network, corresponding to Fig. 7 in Stivala (2020a). The bottom two plots show results for the *dk*-series 2.5k distribution (left) and from the ERGM Model 3 (right).

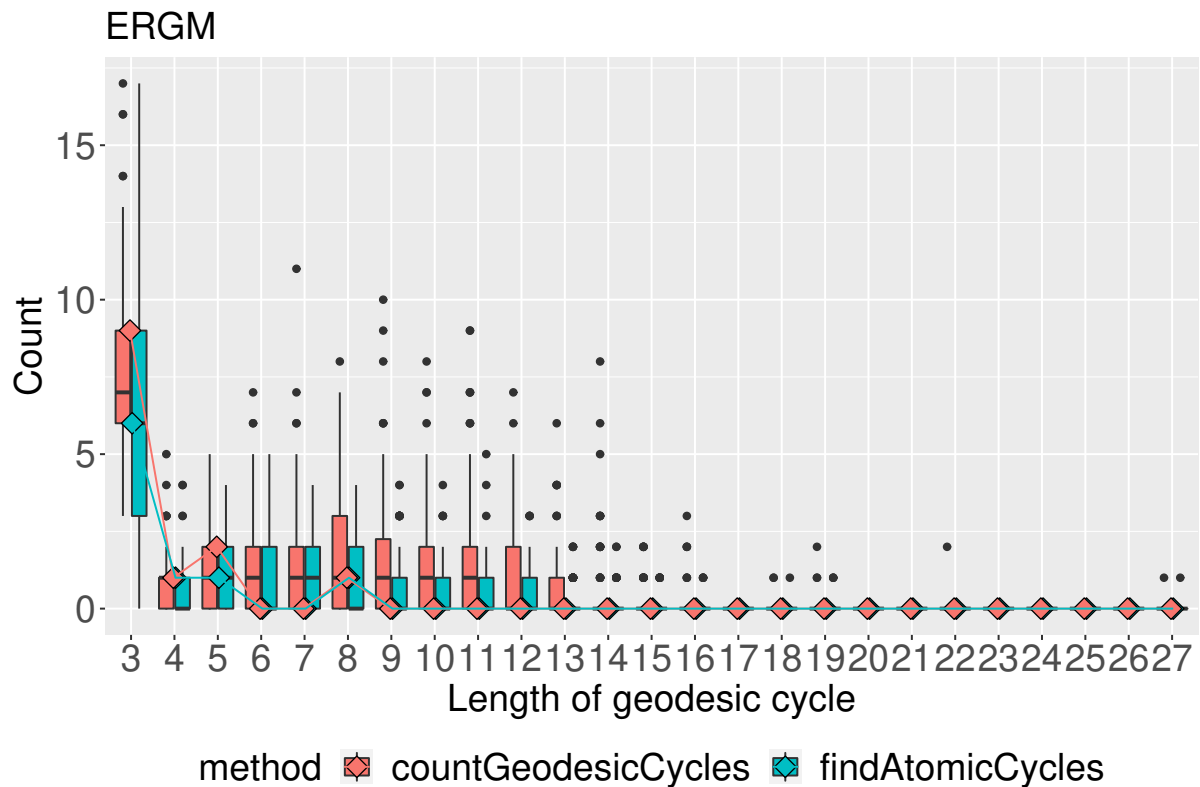
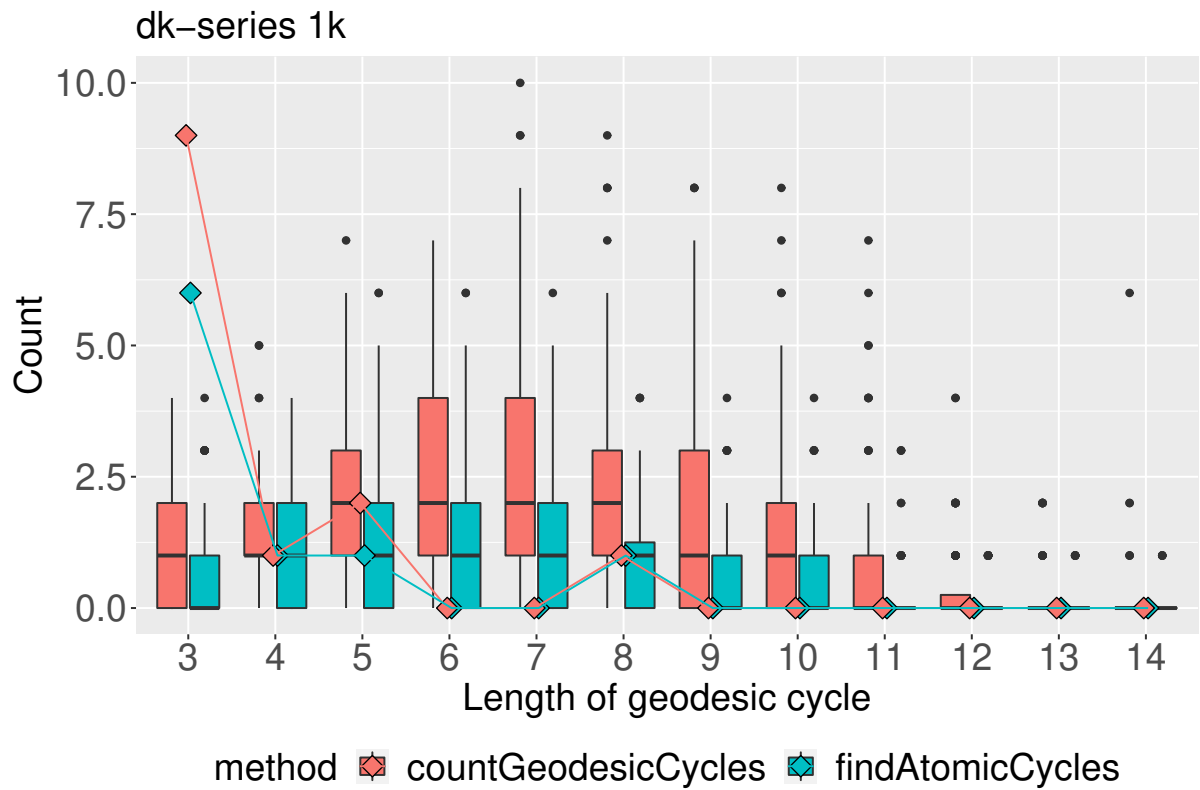


Figure B4: Distribution of geodesic cycle sizes for Patricia's 1993 network, corresponding to Fig. 8 in Stivala (2020a). These results are for the *dk*-series 1k distribution (top) and ERGM Model 1 (bottom).

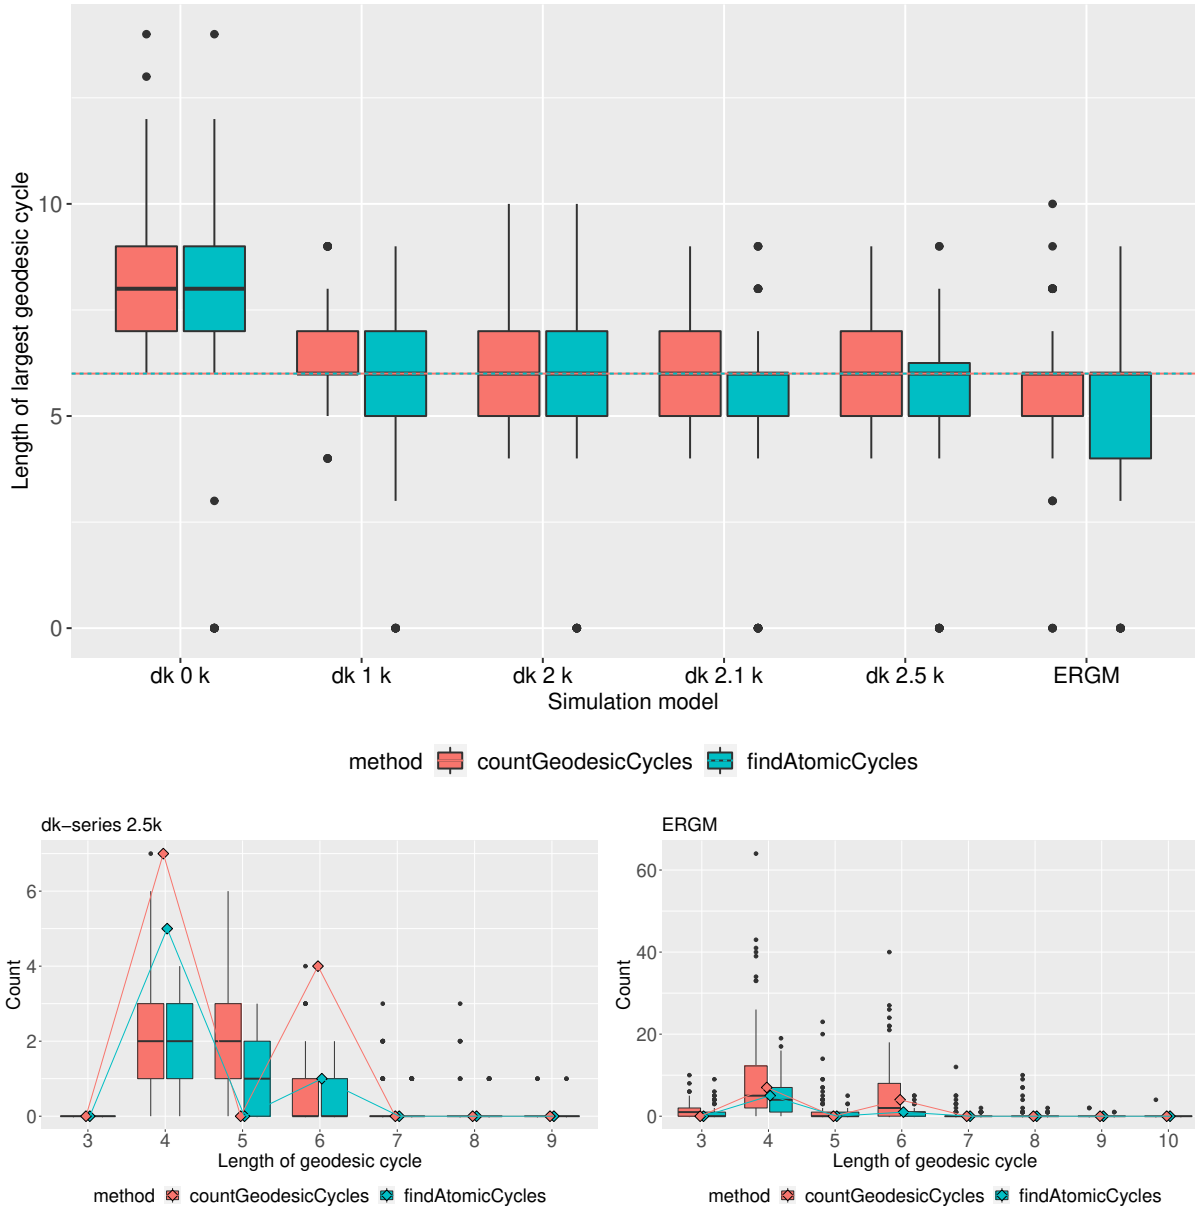


Figure B5: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Grey's Anatomy sexual contact network, corresponding to Fig. 9 in Stivala (2020a). The bottom two plots show results for the *dk*-series 2.5k distribution (left) and from the ERGM (right).



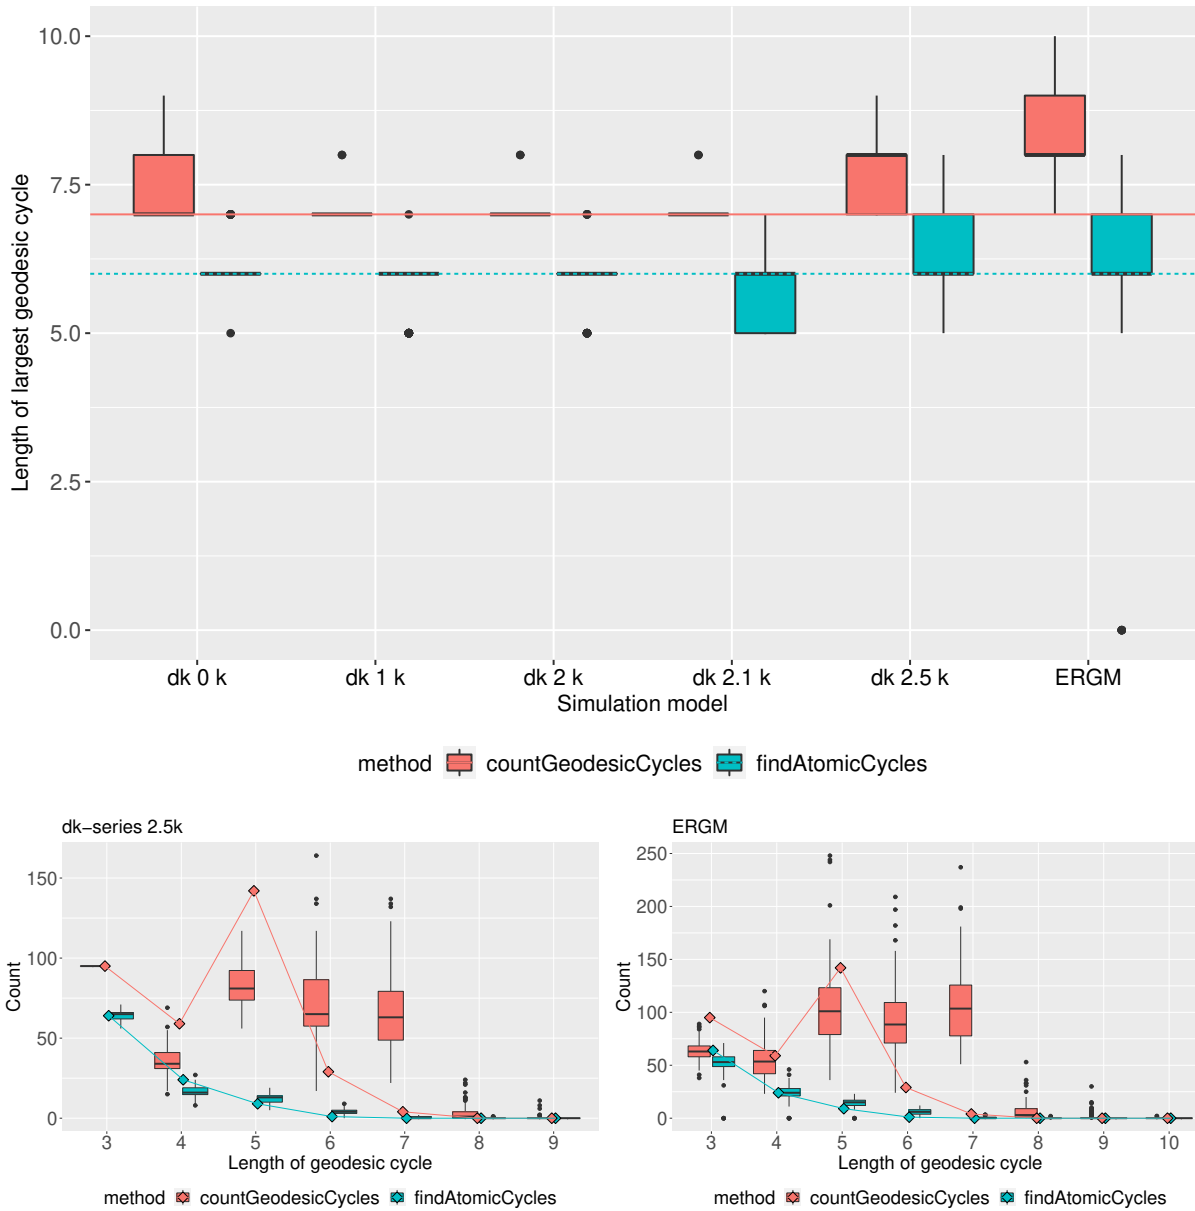


Figure B6: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the dolphin social network, corresponding to Fig. 10 in Stivala (2020a). The bottom two plots show results for the  $dk$ -series 2.5k distribution (left) and from the ERGM (right).

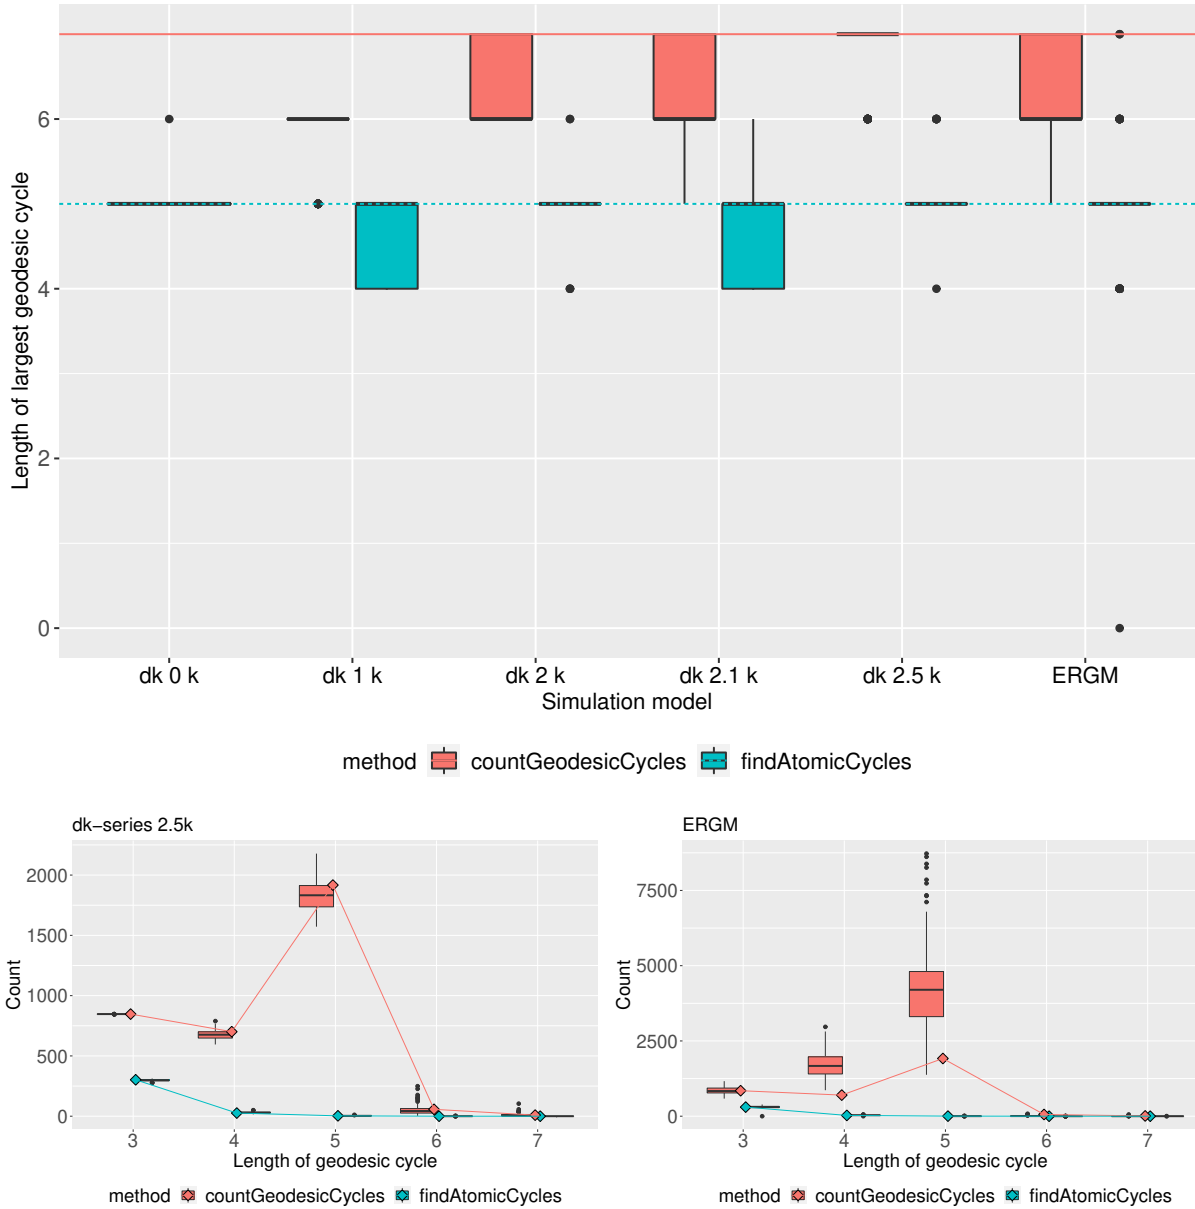


Figure B7: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Lazega law firm friendship network, corresponding to Fig. 11 in Stivala (2020a). The bottom two plots show results for the  $dk$ -series 2.5k distribution (left) and from the ERGM (right).

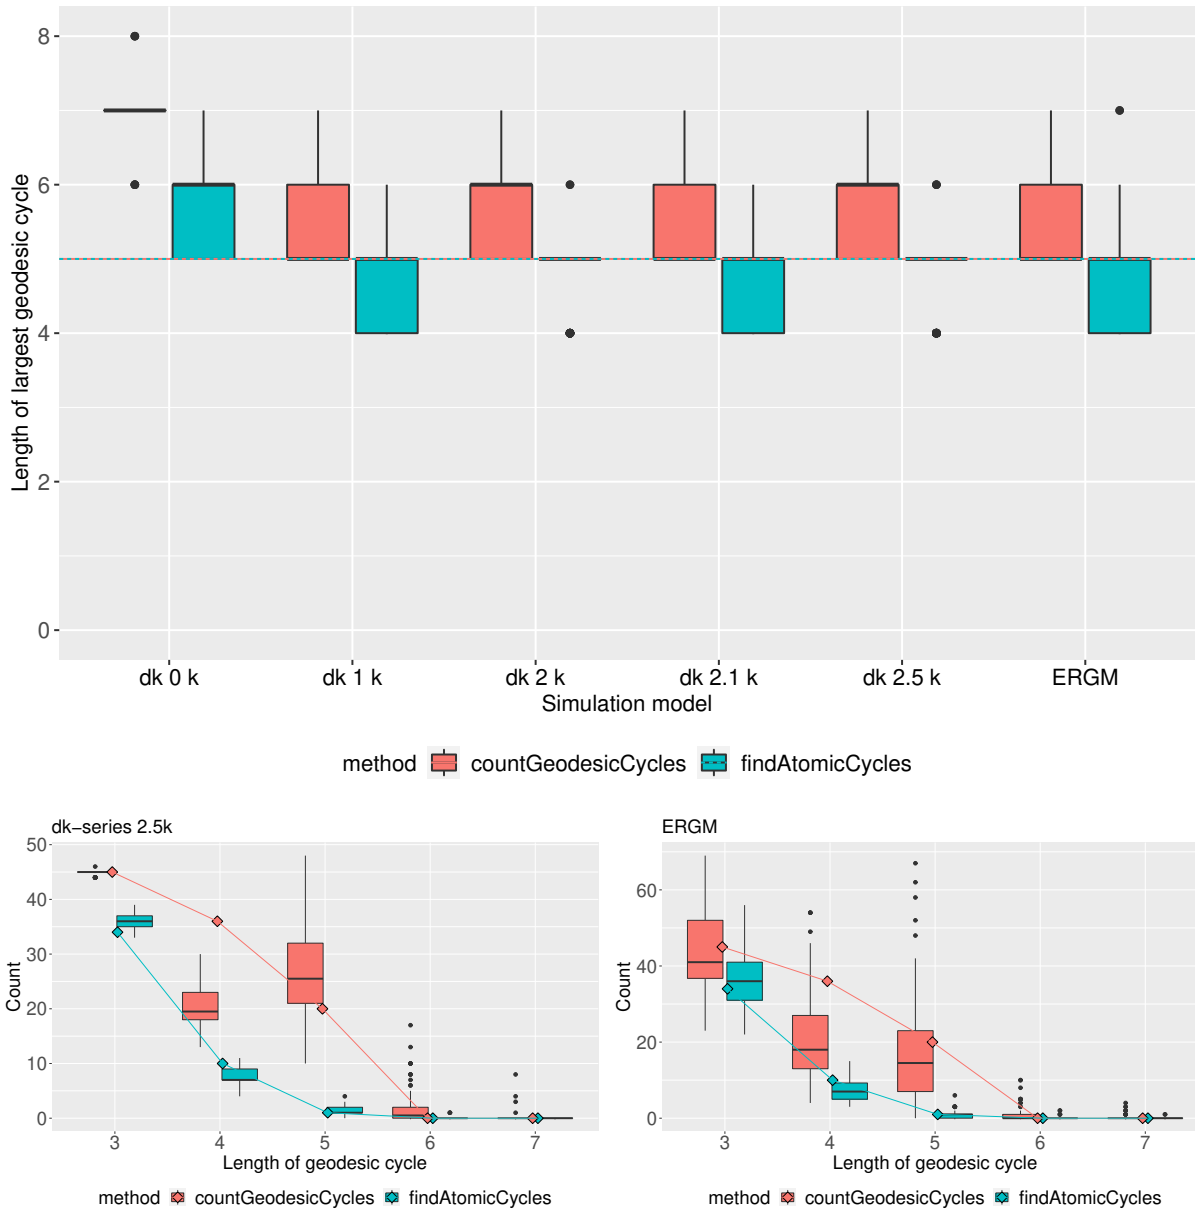


Figure B8: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Zachary karate club network, corresponding to Fig. 12 in Stivala (2020a). The bottom two plots show results for the  $dk$ -series 2.5k distribution (left) and from the ERGM (right).

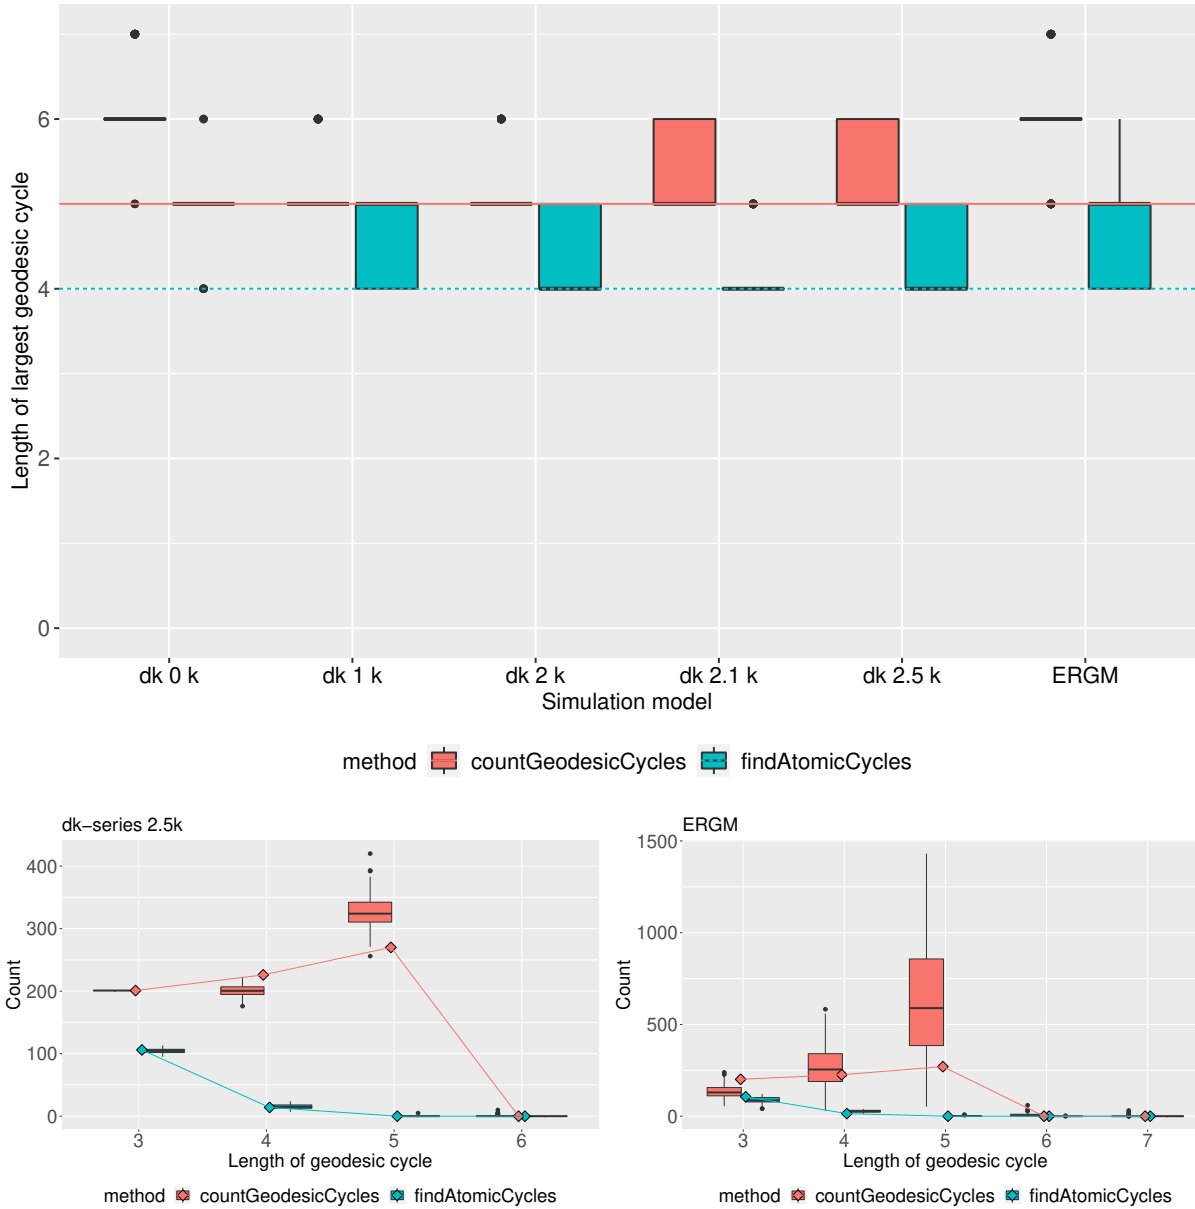


Figure B9: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the Kapferer tailor shop network, corresponding to Fig. 13 in Stivala (2020a). The bottom two plots show results for the  $dk$ -series  $2.5k$  distribution (left) and from the ERGM (right).

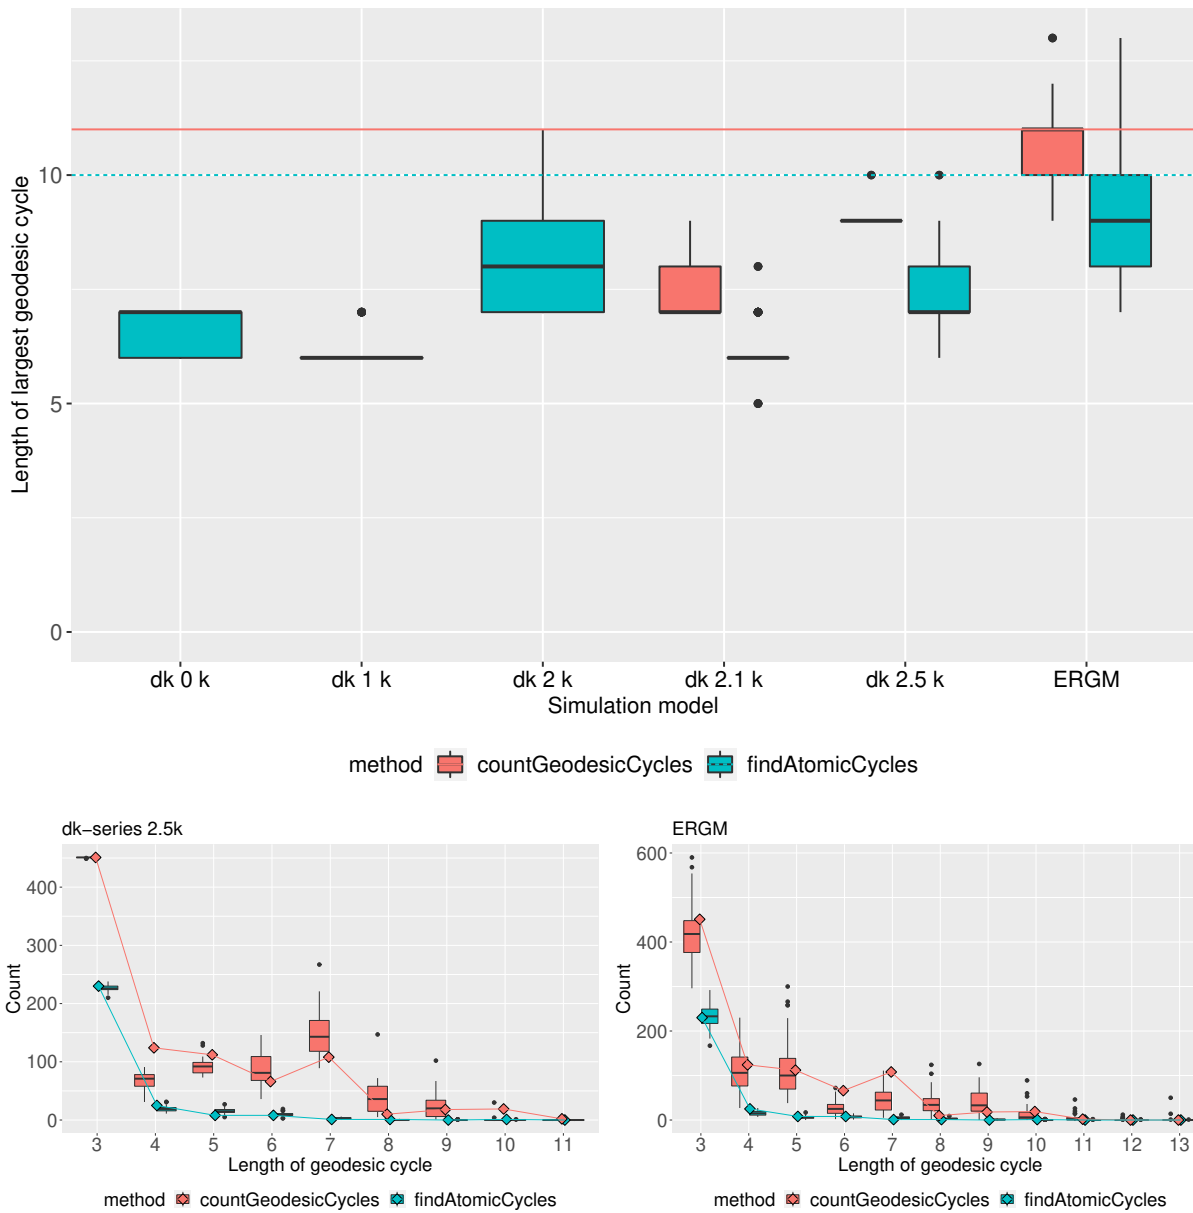


Figure B10: Largest geodesic cycle size (top), and distribution of geodesic cycle sizes (bottom) for the high school friendship network, corresponding to Fig. 14 in Stivala (2020a). The bottom two plots show results for the  $dk$ -series 2.5k distribution (left) and for the ERGM (right). These plots represent only the 93, 17, and 55 networks (from the 100 simulated for each) for  $dk$  2.1 k,  $dk$  2.5 k, and ERGM, respectively, for which the cycles could be counted within the 48 hour time limit.

## Appendix C ERGM models

Table C1: ERGM model of the Dumbledore's Army network.

Effect	Estimate	std. error	p-value	
edges	-3.0729	0.4863	< 0.0001	***
gwdeg.fixed.0.5	1.0245	1.0997	0.351507	
gwesp.fixed.0.5	0.7965	0.2297	0.000526	***

Note: \*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$ .

Table C2: ERGM model of the Harry Potter peer support network.

Effect	Estimate	std. error	p-value	
edges	-3.0046	0.4370	< 0.0001	***
nodefactor.gender.2	-0.0823	0.1920	0.66839	
nodefactor.house.2	-0.6507	0.2420	0.00716	**
nodefactor.house.3	-1.0031	0.2391	< 0.0001	***
nodefactor.house.4	-1.7512	0.2927	< 0.0001	***
nodefactor.schoolyear.1986	1.0385	1.0957	0.34326	
nodefactor.schoolyear.1987	-0.3320	0.3851	0.38868	
nodefactor.schoolyear.1988	-0.1574	0.7422	0.83203	
nodefactor.schoolyear.1989	-0.1282	0.2341	0.58389	
nodefactor.schoolyear.1990	-0.5506	0.3505	0.11619	
nodefactor.schoolyear.1992	0.2222	0.3519	0.52778	
nodefactor.schoolyear.1993	-1.6409	0.6248	0.00863	**
nodefactor.schoolyear.1994	-2.4096	0.5896	< 0.0001	***
nodematch.gender	-0.0533	0.2448	0.82777	
nodematch.house	2.1908	0.2767	< 0.0001	***
nodematch.schoolyear	2.0168	0.3120	< 0.0001	***

Note: \*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$ .

## Goodness-of-fit diagnostics

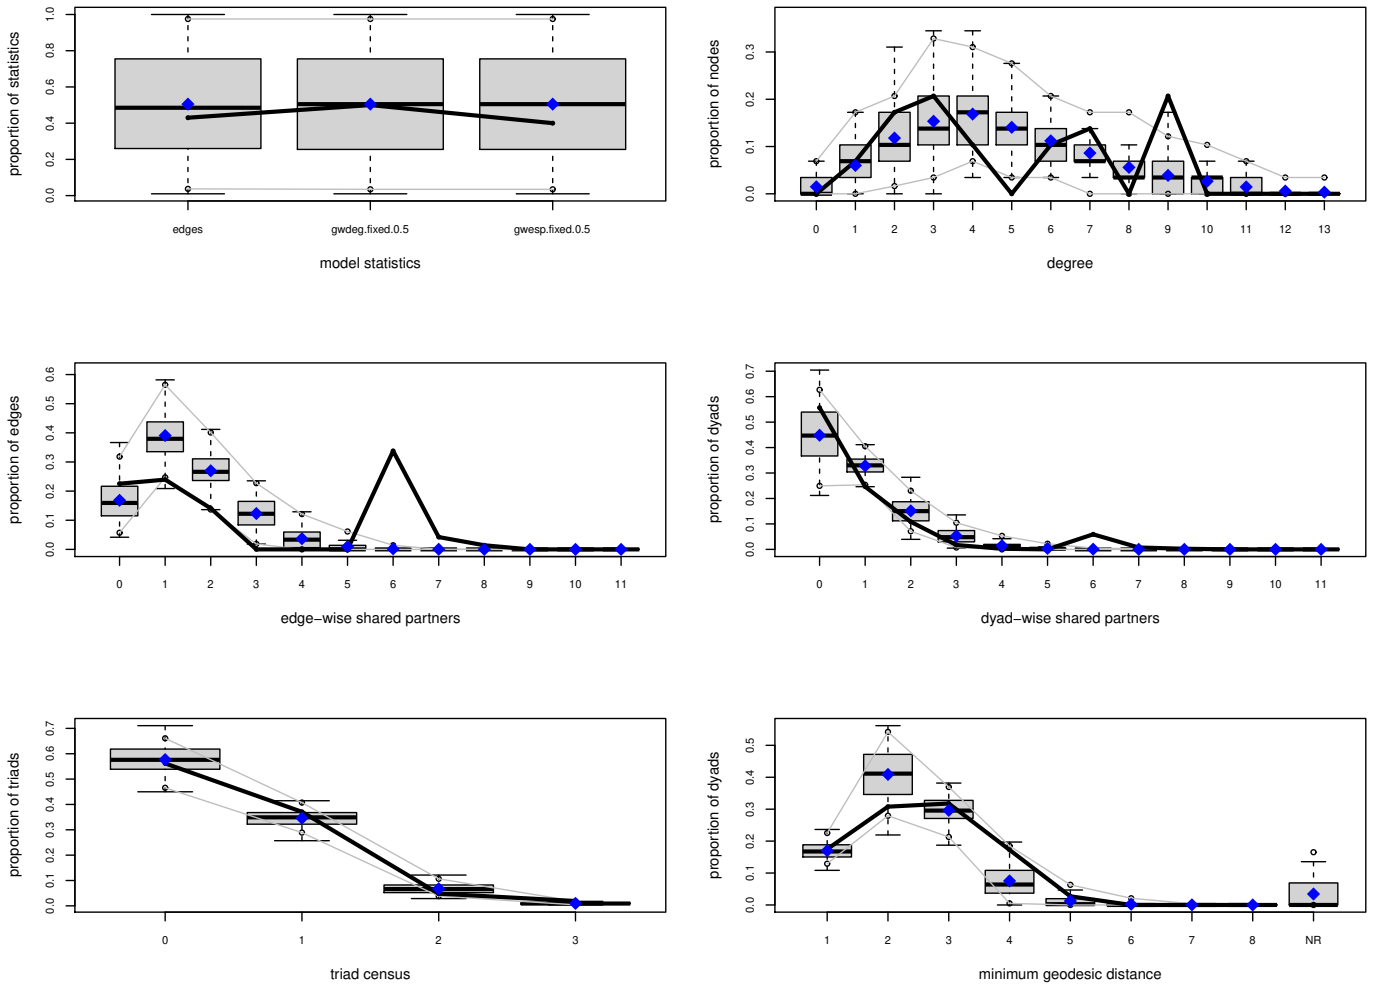


Figure C1: Statnet goodness-of-fit plot for the Dumbledore's Army ERGM in Table C1.

## Goodness-of-fit diagnostics

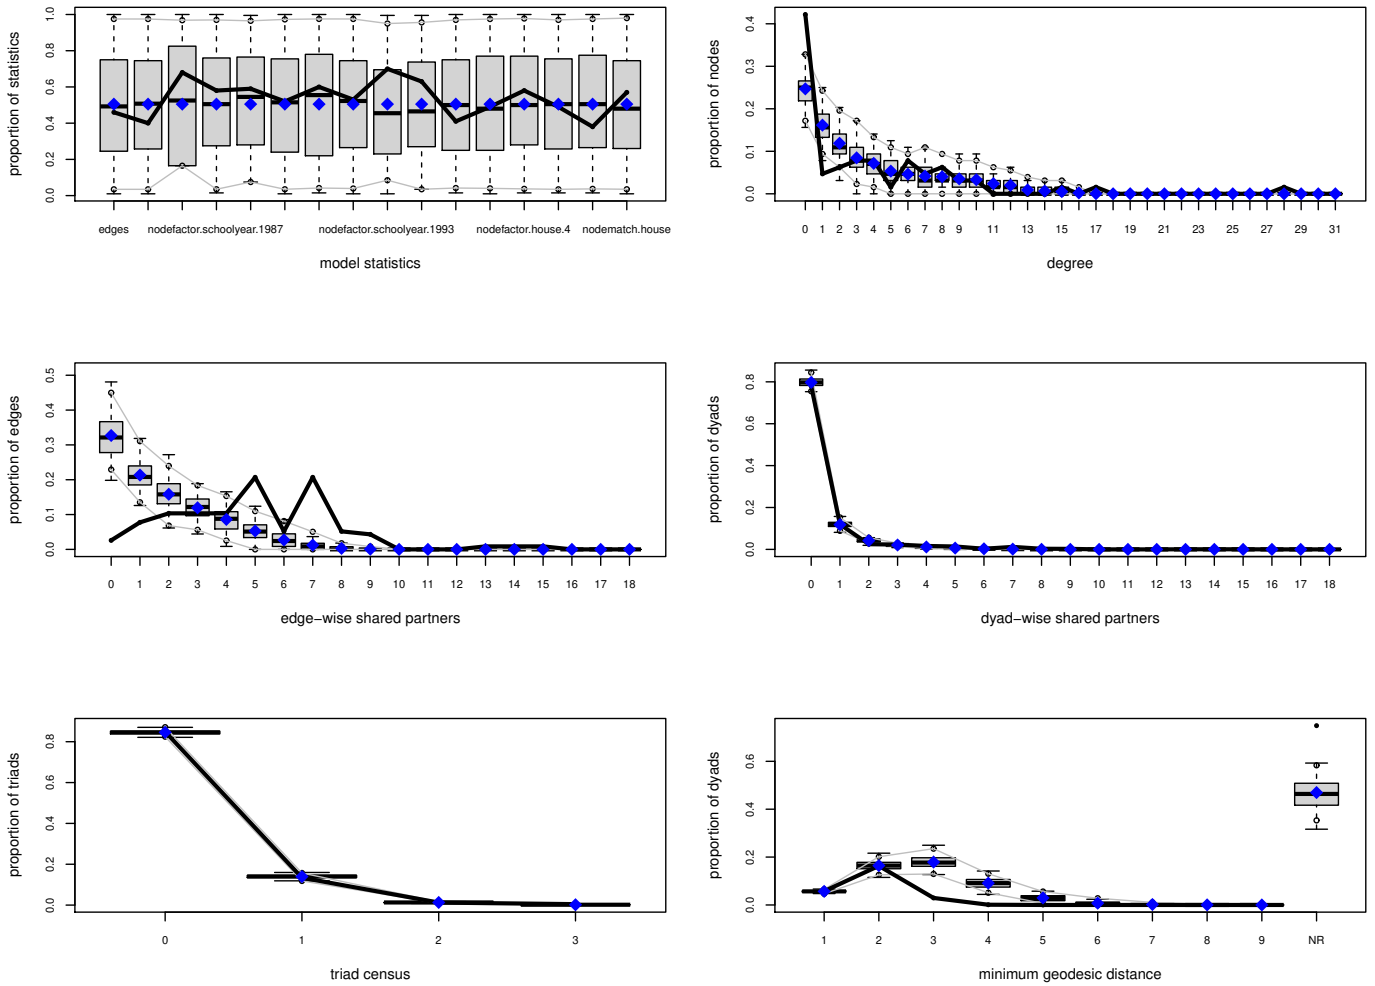


Figure C2: Statnet goodness-of-fit plot for the Harry Potter peer support network ERGM in Table C2.



## References

- E. Amaldi, C. Iuliano, T. Jurkiewicz, K. Mehlhorn, R. Rizzi, et al. Breaking the  $O(m^2n)$  barrier for minimum cycle bases. In *European Symposium on Algorithms (ESA)*, volume 5757 of *Lecture Notes in Computer Science*, pages 301–312. Springer, 2009.
- E. Amaldi, C. Iuliano, and R. Rizzi. Efficient deterministic algorithms for finding a minimum cycle basis in undirected graphs. In F. Eisenbrand and F. B. Shepherd, editors, *Integer Programming and Combinatorial Optimization: 14th International Conference, IPCO 2010, Lausanne, Switzerland, June 9-11, 2010. Proceedings 14*, volume 6080 of *Lecture Notes in Computer Science*, pages 397–410. Springer, 2010. doi: 10.1007/978-3-642-13036-6\_30.
- E. Amaldi, C. Iuliano, T. Jurkiewicz, K. Mehlhorn, and R. Rizzi. Improved minimum cycle bases algorithms by restriction to isometric cycles. <https://people.mpi-inf.mpg.de/%7Emehlhorn/ftp/BarrierCycleBasisJournal.pdf>, 2011. Accessed 10 January 2023.
- V. Amati, A. Lomi, and A. Mira. Social network modeling. *Annual Review of Statistics and its Application*, 5:343–369, 2018.
- J. A. Barnes. Social networks. Addison-Wesley Module in Anthropology, 26, 1972.
- G. Bossaert and N. Meidert. “We are only as strong as we are united, as weak as we are divided” a dynamic analysis of the peer support networks in the Harry Potter books. *Open Journal of Applied Sciences*, 3(2):174–185, 2013. doi: 10.4236/ojapps.2013.32024.
- C. T. Butts. network: a package for managing relational data in R. *Journal of Statistical Software*, 24(2), 2008. URL <http://www.jstatsoft.org/v24/i02/paper>.
- C. T. Butts. *network: Classes for Relational Data*. The Statnet Project (<http://statnet.org>), 2015. URL <http://CRAN.R-project.org/package=network>. R package version 1.18.0.
- F. Catrina, R. Khan, I. Moorman, M. Ostrovskii, and L. I. C. Vidyasagar. Quantitative characteristics of cycles and their relations with stretch and spanning tree congestion. *arXiv preprint arXiv:2104.07872*, 2021.
- P. Colomer-de Simón and M. Boguñá. Double percolation phase transition in clustered complex networks. *Physical Review X*, 4(4):041020, 2014.
- P. Colomer-de Simón, M. A. Serrano, M. G. Beiró, J. I. Alvarez-Hamelin, and M. Boguñá. Deciphering the global organization of clustering in real complex networks. *Scientific Reports*, 3:2517, 2013.
- G. Csárdi and T. Nepusz. The igraph software package for complex network research. *InterJournal, Complex Systems*: 1695, 2006. URL <https://igraph.org>.
- D. Cunningham, S. Everton, and P. Murphy. *Understanding dark networks: A strategic framework for the use of social network analysis*. Rowman & Littlefield, Lanham, MD, 2016.
- A. David, R. Kemp, L. Smith, and T. Fahy. Split minds: Multiple personality and schizophrenia. In P. W. Halligan and J. C. Marshall, editors, *Method in Madness: Case Studies in Cognitive Neuropsychiatry*, chapter 7, pages 122–146. Psychology Press, Hove, UK, 1996.
- E. S. Dias, D. Castonguay, H. Longo, and W. A. R. Jradi. Efficient enumeration of chordless cycles. *arXiv preprint arXiv:1309.1051v4*, 2014.
- S. Everton, T. Everton, A. Green, C. Hamblin, and R. Schroeder. Strong ties and where to find them: Or, why Neville (and Ginny and Seamus) and Bellatrix (and Lucius) might be more important than Harry and Tom. *Social Network Analysis and Mining*, 12:112, 2022. doi: 10.1007/s13278-022-00947-z.
- R. Ferreira, R. Grossi, R. Rizzi, G. Sacomoto, and M.-F. Sagot. Amortized  $\tilde{O}(|V|)$ -delay algorithm for listing chordless cycles in undirected graphs. In A. S. Schulz and D. Wagner, editors, *Algorithms - ESA 2014*, volume 8737 of *Lecture Notes in Computer Science*, pages 418–429. Springer, 2014. doi: 10.1007/978-3-662-44777-2\_35.

- E. Gabasova. The Star Wars social network. <http://evelinag.com/blog/2015/12-15-star-wars-social-network/>, December 2015. Accessed 7 December 2019.
- E. Gabasova. Star Wars social network, Jan. 2016. URL <https://doi.org/10.5281/zenodo.1411479>.
- M. Gashler. Waffles: A machine learning toolkit. *Journal of Machine Learning Research*, 12(69):2383–2387, 2011.
- M. Gashler and T. Martinez. Robust manifold learning with CycleCut. *Connection Science*, 24(1):57–69, 2012.
- M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, Morris, and Martina. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software*, 24(1):1–11, 2008. URL <http://www.jstatsoft.org/v24/i01>.
- M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, S. Bender-deMoll, and M. Morris. *statnet: Software Tools for the Statistical Analysis of Network Data*. The Statnet Project (<http://www.statnet.org>), 2016. URL <http://CRAN.R-project.org/package=statnet>. R package version 2019.6.
- M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, P. N. Krivitsky, and M. Morris. *ergm: Fit, Simulate and Diagnose Exponential-Family Models for Networks*. The Statnet Project (<http://www.statnet.org>), 2022. URL <http://CRAN.R-project.org/package=ergm>. R package version 4.3.2.
- A. J. Holanda, M. Matias, S. M. Ferreira, G. M. Benevides, and O. Kinouchi. Character networks and book genre classification. *International Journal of Modern Physics C*, 30(08):1950058, 2019.
- J. D. Horton. A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM Journal on Computing*, 16(2):358–366, 1987.
- R. M. Hummel, D. R. Hunter, and M. S. Handcock. Improving simulation-based algorithms for fitting ERGMs. *Journal of Computational and Graphical Statistics*, 21(4):920–939, 2012.
- D. R. Hunter, M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29, 2008. URL <https://www.jstatsoft.org/v024/i03>.
- W. A. Jradi, E. S. Dias, D. Castonguay, H. Longo, and H. A. do Nascimento. A GPU-based parallel algorithm for enumerating all chordless cycles in graphs. *arXiv preprint arXiv:1410.4876v2*, 2015.
- B. Kapferer. *Strategy and transaction in an African factory: African workers and Indian management in a Zambian town*. Manchester University Press, 1972.
- K. Klemm and P. F. Stadler. Statistics of cycles in large networks. *Physical Review E*, 73(2):025101, 2006.
- D. E. Knuth. *The Stanford GraphBase: a platform for combinatorial computing*. ACM Press, New York, 1993.
- J. Koskinen. Exponential random graph modelling. In P. Atkinson, S. Delamont, A. Cernat, J. Sakshaug, and R. Williams, editors, *SAGE Research Methods Foundations*. SAGE, London, 2020. doi: 10.4135/9781526421036888175.
- J. Koskinen and G. Daraganova. Exponential random graph model fundamentals. In D. Lusher, J. Koskinen, and G. Robins, editors, *Exponential Random Graph Models for Social Networks*, chapter 6, pages 49–76. Cambridge University Press, New York, 2013.
- P. N. Krivitsky, D. R. Hunter, M. Morris, and C. Klumb. ergm 4: New features for analyzing exponential-family random graph models. *Journal of Statistical Software*, 105(1):1–44, 2023. doi: 10.18637/jss.v105.i06.
- V. Labatut and X. Bost. Extraction and analysis of fictional character networks: A survey. *ACM Computing Surveys (CSUR)*, 52(5):1–40, 2019. doi: 10.1145/3344548.
- E. Lazega. *The collegial phenomenon: The social mechanisms of cooperation among peers in a corporate law partnership*. Oxford University Press, 2001.

- E. Lazega and P. E. Pattison. Multiplexity, generalized exchange and cooperation in organizations: a case study. *Social Networks*, 21(1):67–90, 1999.
- A. Leavitt and J. Clark. SNA in R workshop. <https://github.com/alexleavitt/SNAinRworkshop>, March 2014. Accessed 12 July 2019.
- Y. Li and L. Shi. Geodesic cycles in random graphs. *Discrete Mathematics*, 341(5):1275–1281, 2018.
- B. Lind. Lessons on exponential random graph modeling from Grey’s Anatomy hook-ups. <https://badhessian.org/2012/09/lessons-on-exponential-random-graph-modeling-from-greys-anatomy-hook-ups/>, September 2012. Accessed 12 July 2019.
- D. Lokshantov. Finding the longest isometric cycle in a graph. *Discrete Applied Mathematics*, 157(12):2670–2674, 2009.
- D. Lusher, J. Koskinen, and G. Robins, editors. *Exponential Random Graph Models for Social Networks*. Structural Analysis in the Social Sciences. Cambridge University Press, New York, 2013.
- D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson. The bottlenose dolphin community of Doubtful Sound features a large proportion of long-lasting associations. *Behavioral Ecology and Sociobiology*, 54(4):396–405, 2003.
- P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. *Computer Communication Review*, 36(4):135–146, 2006.
- J. L. Martin. The structure of node and edge generation in a delusional social network. *Journal of Social Structure*, 18(1):1–22, 2017. doi: 10.21307/joss-2018-005.
- J. L. Martin. Comment on geodesic cycle length distributions in delusional and other social networks. *Journal of Social Structure*, 21(1):77–93, 2020. doi: 10.21307/joss-2020-003.
- R. Mastrandrea, J. Fournet, and A. Barrat. Contact patterns in a high school: a comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLoS ONE*, 10(9):e0136497, 2015.
- R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- J. C. Mitchell. Social networks. *Annual Review of Anthropology*, 3(1):279–299, 1974.
- S. Negami and G.-H. Xu. Locally geodesic cycles in 2-self-centered graphs. *Discrete Mathematics*, 58(3):263–268, 1986.
- M. E. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(2):026118, 2001.
- C. Orsini, M. M. Dankulov, P. Colomer-de Simón, A. Jamakovic, P. Mahadevan, A. Vahdat, K. E. Bassler, Z. Toroczkai, M. Boguná, G. Caldarelli, et al. Quantifying randomness in real networks. *Nature Communications*, 6:8627, 2015.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016. URL <https://www.R-project.org/>.
- N. Sokhn, R. Baltensperger, L.-F. Bersier, J. Hennebert, and U. Ultes-Nitsche. Identification of chordless cycles in ecological networks. In K. Glass, R. Colbaugh, P. Ormerod, and J. Tsao, editors, *Complex Sciences: Second International Conference, COMPLEX 2012, Santa Fe, NM, USA, December 5-7, 2012, Revised Selected Papers 2*, volume 126 of *LNICST*, pages 316–324. Springer, 2013.
- A. Stivala. Geodesic cycle length distributions in delusional and other social networks. *Journal of Social Structure*, 21(1):35–76, 2020a. doi: 10.21307/joss-2020-002.
- A. Stivala. Reply to “Comment on geodesic cycle length distributions in delusional and other social networks”. *Journal of Social Structure*, 21(1):94–106, 2020b. doi: 10.21307/joss-2020-004.

- T. Uno and H. Satoh. An efficient algorithm for enumerating chordless cycles and chordless paths. In S. Džeroski, P. Panov, D. Kocev, and L. Todorovski, editors, *International Conference on Discovery Science*, volume 8777 of *LNAI*, pages 313–324. Springer, 2014.
- K. Wasa. Enumeration of enumeration algorithms. *arXiv preprint arXiv:1605.05102*, 2016.
- G. Weissman. Grey’s Anatomy network of sexual relations. <https://gweissman.github.io/post/grey-s-anatomy-network-of-sexual-relations/>, April 2019. Accessed 12 July 2019.
- E. W. Weisstein. Chordless cycle. From MathWorld—A Wolfram Web Resource. <https://mathworld.wolfram.com/ChordlessCycle.html>, 2023. Accessed 28 February 2023.
- H. Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. ISBN 978-3-319-24277-4. URL <https://ggplot2.tidyverse.org>.
- M. Wild. Generating all cycles, chordless cycles, and Hamiltonian cycles with the principle of exclusion. *Journal of Discrete Algorithms*, 6(1):93–102, 2008.
- W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33(4):452–473, 1977.