

The Gain and Cost of Multicast Routing Trees

Milena Janic and Piet Van Mieghem
Delft University of Technology
Faculty of EEMCS
2600 GA Delft, The Netherlands
{M.Janic, P.VanMieghem}@ewi.tudelft.nl

Abstract – *The last several years we witness the proliferation of multimedia applications on Internet. One of the unavoidable techniques to support this type of communication is multicasting. However, even a decade after its initial proposal, multicast is not widely deployed. One reason lies in the lack of a business model. If the gain and the cost of multicast could be predicted, network operators might be encouraged to deploy multicast on a larger scale. In this paper we propose several analytical expressions that could be used to estimate the gain and cost of network-layer multicast. We show that the theoretical model we propose matches simulation and Internet measurement results remarkably well.*

Keywords: Multicast, Shortest Path Tree, Steiner Minimum Tree, Uniform Recursive Tree.

1 Introduction

The number of multimedia applications on Internet, combining audio, video and data streams, is growing explosively. Multimedia applications, even when data compression is used, require in general a considerable amount of bandwidth, and they are often delay sensitive. Such applications include radio/television broadcast, desktop video/audio conferencing, shared white boards, tele-classing, file transfers to multiple locations, online gaming and animated simulations. IP multicast, offering a scalable point-to-multipoint delivery, is regarded as a promising network service for group multimedia applications.

Multiple user communication can be realized in several ways. The minimum cost approach would be to construct a single tree to distribute the traffic from all senders in the group, regardless of the sender's location, and to minimize the total weight of the tree. Hence, it optimizes the use of network resources. The problem of finding a minimum weight tree that spans all multicast users is known as the Steiner Tree problem [14]. However, the computational complexity of finding a Steiner Minimum Tree (SMT), proven to be *NP*-complete, together with its less stable dynamic behavior [20] prohibits the implementation of this algorithm for

multicast routing protocols on Internet. Instead, most of the current Internet protocols forward packets based on the (Reverse) Shortest Path. A Shortest Path Tree (SPT) is then composed from shortest paths between the source and the destination. The SPT algorithm does not necessarily result in a tree that economizes on network resources but it is easy to compute and it offers a minimum delay.

Even though the first deployment of multicast occurred in 1992, and in spite of the continuously rising demand for a ubiquitous multicast service, IP multicast is still experiencing slow wide-scale deployment. One of the reasons for this is the lack of a proper business model. The computational and administrative overhead of multicast group management increases the deployment cost compared to the cost of unicast. Clearly, the deployment of multicast can only be justified if the netto gain defined as savings minus costs is larger than the netto gain for unicast. Therefore, we believe the deployment of multicast on larger scale would be encouraged if the gain/cost of multicast could be successfully estimated.

Different researchers have introduced different definitions for the gain and the cost of multicast. One possible criterion to assess the gain of the multicast is the number of hops in the tree rooted at a particular source to m randomly chosen destinations. Alternatively, the cost of multicast can be defined as the sum of all the link weights in a tree connecting m uniformly chosen nodes. In this paper we propose analytic expressions that can be used to compute the gain and cost of the Shortest Path Tree, which we also compare with the Steiner Minimum Tree. Moreover, we investigate how well the simulation and Internet measurement data fit the proposed model.

The remainder of the paper is organized as follows: in the following Section we give a short overview of the related work on the multicast cost and the gain analysis, together with our previous work on this subject. Analytical results on the Uniform Recursive Tree (URT) are given in Subsection 3.1, followed by simulation results in Subsection 3.2. In the Section 3.3, the quality of the

URT as model for multicast trees is compared with Internet measurements. Finally, we conclude in Section 4.

2 Related Work on a Multicast Cost Analysis

Quantifying the gain/cost of multicast has been initiated by Chuang and Sirbu [6], and further investigated by Philips *et al.* [16] and Chalmers and Almeroth [4]. In [6], Chuang and Sirbu assumed that multicast packets are being delivered along a shortest-path tree from source to m destinations. As most of the current Internet protocols forward packets based on the (reverse) shortest path, the assumption of shortest-path tree delivery is quite realistic. Further, they assumed that m destinations are uniformly chosen out of N nodes, an assumption that is confirmed by Internet measurements as well [16] [4].

Let us denote by $H_N(m)$ the number of hops in the shortest path tree rooted at a particular source to m randomly chosen destinations. The average number of hops of a message to a uniformly chosen destination in the graph containing N nodes is $E[H_N] = E[H_N(1)]$. Chuang and Sirbu defined the normalized cost of a multicast tree, as the ratio $E[H_N(m)]/E[H_N]$. Via simulations on various network topologies, they suggested a power-law $E[H_N(m)]/E[H_N] \approx m^{0.8}$. However, Van Mieghem *et al.* [18] have shown that this *Chuang-Sirbu* law cannot hold for all m .

Finally, related work has been done on multicast pricing and billing. One of the earliest schemes has been proposed by Herzog *et al.* [12]. Einsiedler *et al.* [9] propose to charge the users for the resources they use. The weights of a link is determined by the maintenance costs, congestion on the link, or other factors. At each router in the tree where branching takes place (branching point), costs are determined by splitting the cost among subtrees. Each router stores the information on the number of branches and on link weights, and distributes this information along the branches to other branching points.

3 The Gain and the Cost of Multicast

3.1 Theory of the Uniform Recursive Tree (URT)

As mentioned above, Van Mieghem *et al.* [18] have shown that the *Chuang-Sirbu* law cannot hold for all m and they presented the general framework, valid for any number of nodes N and of users m and any topology. In [18] the exact expression for the multicast efficiency over unicast is derived in the random graph $G_p(N)$ [2], with independent exponential link weights w , and m multicast users uniformly chosen out of the N nodes. Earlier,

the same group of authors [19] has shown that the Shortest Path Tree in the complete graph is exactly, and in the class $G_p(N)$ asymptotically, a Uniform Recursive Tree [17]¹ (URT). For the URT, the average number of hops to m randomly chosen destinations, for every N and m is given by

$$E[H_N(m)] = \frac{mN}{N-m} \sum_{k=m+1}^N \frac{1}{k} \quad (1)$$

Van Mieghem *et al.* [18] have further shown that for the complete graph with exponentially distributed weights, and for m small with respect to N , the ratio $E[H_N(m)]/mE[H_N]$ increases about linearly with m on a log-log scale, explaining the empirical Chuang-Sirbu law. In addition to the average (1), the same authors derive in [13] the exact probability generating function and probability distribution $\Pr[H_N(m) = k]$ from which the variance follows as

$$\begin{aligned} \text{var}[H_N(m)] &= \frac{N-1+m}{N+1-m} E[H_N(m)] - \\ &\quad - \frac{(E[H_N(m)])^2}{(N+1-m)} - \\ &\quad - \frac{m^2 N^2}{(N-m)(N+1-m)} \sum_{k=m+1}^N \frac{1}{k^2} \end{aligned} \quad (2)$$

The interest of this result as shown in [13] is that for all $m = o(\sqrt{N})$, $\frac{H_N(m) - E[H_N(m)]}{\sqrt{\text{var}[H_N(m)]}}$ converges to a standard normal (Gaussian) random variable when $N \rightarrow \infty$.

Recently, Bollobas *et al.* [3] have shown that in the complete graph with N nodes and with exponentially distributed link weights with mean 1, the asymptotic weight of the Steiner tree spanning $m+1$ uniformly chosen nodes and m small with respect to N ,

$$E[W_{\text{Steiner}}(k)] = (1 + o(1)) \frac{m}{N} \log \frac{N}{m+1} \quad (3)$$

Triggered by this result, van der Hofstad *et al.* [13] have derived the average $E[W_N(m)]$ of the sum of the weights $W_N(m)$ in the Shortest Path Tree to m uniform multicast users in the random graph $G_p(N)$ with exponentially distributed link weights,

$$E[W_N(m)] = \sum_{j=1}^m \frac{1}{N-j} \sum_{k=j}^{N-1} \frac{1}{k} \leq \frac{\pi^2}{6} \quad (4)$$

It has been shown [17] that the ratio of the average of the number of nodes with degree k , denoted by D_k^N , over the total number of nodes in the URT obeys for large N

¹A URT is defined by the following growth rule: given a URT with N nodes, a URT with $N+1$ nodes is deduced by attaching the $N+1$ -th node uniformly (thus with probability $1/N$) to each of the N other nodes in the tree.

$$\frac{E[D_k^N]}{N} = \frac{1}{2^k} + O\left(\frac{\log^{k-1} N}{N^2}\right) \quad (5)$$

which is, for large N , close to $\Pr[\text{deg} = k]$, the probability that an arbitrary node has a degree k . Hence, the decay rate of the pdf of the node degrees in the URT equals $-\ln 2 = 0.693$.

The importance of these results lies in the fact that, if we can demonstrate that the URT characteristics match the Internet measurements for multicasting, then the cost/gain of multicast routing trees can be estimated.

3.2 Simulation Results and a Conjecture

The aim of the simulations described in this Section was to complement the analytical results derived in [13]. We confined ourselves to complete graphs with exponentially distributed link weights with mean 1. For each number of nodes N , 10^5 topologies were generated randomly. For each of these topologies, $m \in \{1, N-1\}$ nodes were uniformly chosen. The Shortest Path Tree (SPT) and the Steiner Minimum Tree (SMT) have been computed subsequently. The SPT is computed by using the Dijkstra algorithm, with $N \leq 100$. Depending on m , the Steiner tree [14] is generated using different algorithms. For $m = 2$, the SMT problem reduces to the computation of the shortest path between those two users. If $m = N$, the SMT is actually the (complete) minimum spanning tree, and is computed with the Prim algorithm. For $2 < m < N$, the SMT problem belongs to the class of hard *NP-complete* problems. Certain reductions [14] in the topology decrease the number of nodes and links and increase the speed of simulations in that reduced graph. In spite of the implemented reductions, the simulation process is nevertheless extremely time consuming for large N . Therefore, we restrict the simulations of *SMT* to graphs with $N \leq 20$. In each graph and for each m , the sum of the weights as well as the number of links in both the SPT and the SMT have been stored in 4 histograms. From these histograms, the probability density function of the sum of the link weights $\frac{d}{dw}(\Pr[W_N(m) \leq x]) = f_{W_N(m)}(x)$ in the SPT and the SMT, as well as the probability density function of the number of links $\Pr[H_N(m) = k]$ have been deduced. Since $\Pr[H_N(m) = k]$ is analytically known [13], these simulations are not shown, but served as a verification for the simulations.

Figure 1 gives the probability density function of the sum of the link weights $f_{W_N(m)}(x)$ in the SPT.

The average value $E[W_N(m)]$ of the sum of the link weights in the SPT and the SMT is plotted as a function of the multicast group size m , for the number of nodes $N = 20$ in Figure 2. Apart from the match between simulations and theory for the SPT, this figure reveals




Figure 1: The pdf of sum of the weights in the Shortest Path Tree for $N=100$

that $E[W_N(m)]$ for the SMT seems similar (apart from some scaling factor) to that of the corresponding SPT.




Figure 2: The average value of the sum of the link weights for SPT and SMT ($N=20$)

In Figure 3, simulation results of the variance of $W_N(m)$ in the SPT and the SMT are shown. So far, $\text{var}[W_N(m)]$ has not been derived analytically.

Although N is small (which allows us to show the entire m -range), Figure 4 indicates that the scaled random variable $\frac{W_N(m) - E[W_N(m)]}{\sqrt{\text{var}[W_N(m)]}}$ is close to a Gumbel type $e^{-e^{-x}}$, which suggests us to conjecture, for all m , that

$$\lim_{N \rightarrow \infty} \Pr \left[\left(\frac{W_N(m) - E[W_N(m)]}{\sqrt{\text{var}[W_N(m)]}} \right) \leq x \right] = e^{-e^{-\frac{x}{\sqrt{2}}}} \quad (6)$$

For the particular case of $m = 1$, we are able to prove this result (see Appendix).




Figure 3: The variance of the sum of the link weights for SPT and SMT ($N=20$)




Figure 4: The pdf of the scaled random variable $X_N(m) = \frac{W_N(m) - E[W_N(m)]}{\sqrt{\text{var}[W_N(m)]}}$.

3.3 Measurements on Internet

In 1999, Faloutsos *et al.* [10] studied three different Internet AS level instances. Their results indicated that the degree distribution follows power law characteristics. These findings have provoked astonishment among network researchers and stimulated the desire to explain this behavior. However, Chen *et al.* [5] have criticized the results published in [10] implying that they were obtained on incomplete AS graphs. They show that when data is obtained from more complete AS graphs, the degree distribution follows heavily skewed distributions (the values of degree vary over 3 to 4 orders of magnitude), but where only tails decay like a power law.

Faloutsos *et al.* [10] have also suggested the power-law node degree distribution for nodes in a router-level map of Internet. Ever since, several global router-level Internet mapping projects have been initiated, almost

all being based on *traceroute*, or some *traceroute-like* utility. Over the time the mapping techniques have advanced and improved, resulting in more and more accurate maps. However, due to several important reasons, as we discuss in [15], it is dangerous to apply the conclusions drawn on those imperfect maps to the actual Internet topology. Nevertheless, in [15] we have performed a thorough analysis of results obtained by other researchers on node degrees. We have concluded that most of the researchers have observed power-laws, with similar values for the slope coefficient α . Only the results based on CAIDA measurement data had the considerably larger slope coefficient. These measurements indicate that the random graph $G_p(N)$ with exponentially distributed weights and binomially distributed node degree distribution does not appear to be a good model for Internet topology.

Whereas the node degree distribution in Internet maps attracted much attention, only few results have been published on the characteristics of multicast routing trees so far. The first one has been provided by Chalmers and Almeroth [4], who have looked into the properties of the Internet multicast trees on Mbone. They have gathered multicast tree data for four live multicast sessions: the 43rd IETF meeting in December 1998 and the NASA shuttle launch in February 1999, each of them consisting of a separate audio and video channel. The path from each receiver to the source has been traced via *mtrace* (multicast *traceroute*) [11]. The *mtrace* data has been used for each *dataset* to reconstruct a multicast tree. Chalmers and Almeroth have developed the tool *mwalk*, that builds an activity graph of all possible trees over time. In [15, Figure 3(a).] we have plotted the node degree distribution for one realization of the tree, in 43rd IETF video *dataset*, when 129 receivers have been traced to belong to the group. After fitting their data on different scales, we noticed that the best fit of all (with the correlation coefficient² of 0.91) has been obtained for the linear fit on the log-lin scale, suggesting rather exponentially than polynomially distributed node degrees. Interestingly, the slope of the curve is approximately the same as that of the URT (see (5)).

The only other result on the multicast tree degree distribution so far has been provided by Dolev *et al.* [8]. In [8] Dolev *et al.* have investigated properties of multicast trees obtained from Internet measurement data. The data they used for their multicast analysis has been obtained via unicast *traceroute* measurement. They have used two *datasets*: first, on the underlying topology provided from *traceroute* measurements, they generated shortest path trees using the Dijkstra algorithm. The second *dataset* has been created based on

²The linear correlation coefficients measure the extent of linear relationship of two variables, and are given by $r = \frac{\text{cov}(y,x)}{\sqrt{\text{var}(y)\text{var}(x)}}$.

traceroute measurements of the paths between the root and the clients in the client population of `www.bell-labs.com`. In [8, Figure 6 and Figure 7] they have plotted the node degree distributions in both *datasets* on a log-log scale, and fitted with the linear function decaying with the rate -3.40 and -3.18 for the first and the second *dataset*, with the correlation coefficients of 0.9897 and 0.9829 , respectively. These findings seem to suggest a power-law structure of the node degree distribution in multicast routing trees, which contradicts that of [4].

With the purpose of understanding this discrepancy, in [15] we have investigated the node degree distribution in multicast trees constructed as unions of *traceroutes* obtained from both CAIDA's Skitter project [21] and RIPE [22]. We have demonstrated that the degree distribution in the multicast tree for small m does not follow a power law. For larger values of the number of destinations m the degree distribution seems to follow a power-law, when multicast trees are created as union of *traceroutes*. The result of Chalmers and Almeroth implies that if another method for constructing trees is used other than the union of *traceroute* paths, power-laws might not even be observed for larger values of m . Multicast routing trees have been obtained from CAIDA's Skitter project [21] *traceroute* data in the following way: first, we have randomly chosen $m = 50, 100, 500, 1000$ destinations (multicast users). For three monitor boxes (two of them situated in United States and one in Japan) the collection of paths from these three sources to randomly chosen destinations has been obtained via *traceroute*. In this way, we obtained for each source a set of 4 trees. For each of these 12 trees, the number of links in the tree has been calculated. In Figure 5 this number has been plotted, for each source, as a function of m . In addition, for $N = 135000$ (the approximate number of nodes in the Internet map derived from *traceroute* measurements, see [15]) and for various values of m (in the range $[50, 20000]$), the values of the functions $E[H_N(m)]$ and $E[H_N(m)] \pm 6\sigma_N(m)$ (where $\sigma_N(m) = \sqrt{\text{var}(H_N(m))}$) have been computed, and plotted in the same Figure. Since $\sigma_N(m)$ is much smaller than $g_N(m)$, the number of links in the URT is well approximated by the mean, $H_N(m) \approx E[H_N(m)]$ for large values of N .

Also in our previous work we have shown that our analytical findings were supported by Internet measurements (we refer to [18, Figure 7]). Hence, all of our measurement results so far indicate that the URT represents a reasonable, first order model for the multicast tree in Internet.

4 Conclusions

One of the reasons for the slow wide-scale deployment of IP multicast is the lack of a proper business model. In this paper we proposed analytical expressions (1) and




Figure 5: The average number of links (Caida measurements and theoretical value)

(3) for assessing the gain/cost of Shortest Path. We have compared the cost of Shortest path to Steiner Trees using the analytical expression (3) for estimating the gain/cost of Steiner Trees. From the extensive simulations we were lead to conjecture that $\Pr[W_N(m) \leq x] \sim \exp\left(-\exp\left(-\frac{\pi(x-E[W_N(m)])}{\sqrt{2\text{var}[W_N(m)]}}\right)\right)$ for large N .

Internet measurements seem to suggest that for the small number of users m , the URT model we used to derive the laws given above represents a reasonable model for multicast trees in Internet. The measurement results indicated that the average number of links in multicast trees lies in the range $E[H_N(m)] \pm 6\sigma_N(m)$ (where $\sigma_N(m) = \sqrt{\text{var}(H_N(m))}$) for all values of m . Since $\sigma_N(m)$ is much smaller than $g_N(m)$ for large N , the number of links in the URT is well approximated by the mean, $H_N(m) \approx E[H_N(m)]$.

5 Acknowledgements

We would like to thank A. Broido and KC Claffy (CAIDA) for the CAIDA measurement data.

References

- [1] M. Abramowitz and I.A. Stegun, *Handbook of Mathematical Functions*, Dover, 1968.
- [2] B. Bollobas, *Random Graphs*, Academic Press, 1985.
- [3] B. Bollobas, D. Gamarnik, O. Riordan and B. Sudako, "On the value of a random minimum weight Steiner tree", *Combinatorica* 24 (2004), 187-207.
- [4] R. Chalmers and K. Almeroth, "On the Topology of Multicast Trees", 153- 165, *IEEE Transactions on Networking*, April 2003.
- [5] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The origin of power laws in Internet topologies revisited", *Proceedings of IEEE INFOCOM*, 2002.
- [6] J. C.-I. Chuang and M. A. Sirbu, "Pricing multicast communications: A cost based approach", *Proceedings of INET*, 1998.

- [7] T.H. Cormen, C.E. Leiserson and R.L. Rivest, *Introduction to Algorithms*, The MIT Press, McGraw-Hill, 1995.
- [8] D. Dolev, O. Mokryn and Y. Shavitt, "On Multicast Trees: Structure and Size Estimation", Proceedings of IEEE INFOCOM, April 2003.
- [9] H. Einsiedler, P. Hurley, B. Stiller and T. Braun, "Charging Multicast Communications Based on A Tree Metric", Proceedings of Multicast Workshop, Braunschweig, Germany, 1999.
- [10] M. Faloutsos, P. Faloutsos, C. Faloutsos, "On Power-Law Relationships of the Internet Topology", Proceedings of SIGCOMM '99, pp. 251-262, August 1999.
- [11] W. Fenner and S. Casner, "A 'traceroute' facility for IP multicast", Tech. Rep, draft-ietf-idmr-traceroute-ipm-*txt, Internet Engineering Task Force (IETF), August 1998.
- [12] S. Herzog, S. Shenker, and D. Estrin, "Sharing the 'cost' of multicast trees: An axiomatic analysis", IEEE Transactions on Networking, 5(6):847-860, 1997.
- [13] R. van der Hofstad, G. Hooghiemstra and P. Van Mieghem, "Size and Weight of Shortest Path Trees with Exponential LinkWeights", submitted to *Combinatorica*.
- [14] F. Hwang, D. Richards, P. Winter, *The Steiner Tree Problem* (Annals of Discrete Mathematics, Vol 53), North Holland, pp. 103, 1992.
- [15] M. Janic and P. Van Mieghem, "Inferring the Internet Topology based on Traceroute Measurements", Delft University of Technology, Technical Report, 2004.
- [16] G. Phillips, S. Shenker, H. Tangmunarunkit, "Scaling of Multicast Trees: Comments on the Chuang-Sirbu scaling law", Proceedings of SIGCOMM'99, August 1999.
- [17] R. T. Smythe and H. M. Mahmoud, "A Survey of Recursive Trees", *Theor. Probability and Math Statist.*, No. 51, pp. 1-27, 1995.
- [18] P. Van Mieghem, G. Hooghiemstra and R. van der Hofstad, "On the Efficiency of Multicast", IEEE/ACM Transactions on Networking, vol. 9, No. 6, pp. 719-732, 2001.
- [19] P. Van Mieghem, G. Hooghiemstra and R. van der Hofstad, "A Scaling Law for the Hopcount in Internet", report 2000125 (<http://www.nas.ewi.tudelft.nl/people/Piet/teleconference.html>).
- [20] P. Van Mieghem and M. Janic, "Stability of a Multicast Tree", in Proceedings of IEEE INFOCOM, July 2002.
- [21] <http://www.caida.org>.
- [22] RIPE Test Traffic Measurements, <http://www.ripe.net/ripence/mem-services/ttm/>

A Proof of Conjecture (6) for $m = 1$

In [13], the probability generating function of the weight $W_N = W_N(1)$ of the shortest path is derived

$$\varphi_{W_N}(z) = E[e^{-zW_N}] = \frac{1}{N-1} \sum_{k=1}^{N-1} \prod_{n=1}^k \frac{n(N-n)}{z+n(N-n)}$$

The variance, computed from $\varphi''_{W_N}(0) - (\varphi'_{W_N}(0))^2$, is

$$\text{var}[W_N] = \frac{3}{N(N-1)} \sum_{n=1}^{N-1} \frac{1}{n^2} - \frac{\left(\sum_{n=1}^{N-1} \frac{1}{n}\right)^2}{(N-1)^2 N}$$

and for large N ,

$$\text{var}[W_N] = \frac{\pi^2}{2N^2} + O\left(\frac{\log^2 N}{N^3}\right)$$

The limit for $N \rightarrow \infty$ will be computed from which the distribution then follows by taking the inverse Laplace transform. With $y = \sqrt{\left(\frac{N}{2}\right)^2 + z}$, we have

$$\prod_{n=1}^k \frac{n(N-n)}{z+n(N-n)} = \frac{k!(N-1)!}{(N-k-1)!} \prod_{n=1}^k \frac{1}{\left(y + \frac{N}{2} - n\right)\left(y - \frac{N}{2} + n\right)}$$

The products can be written in terms of the Gamma function,

$$\varphi_{W_N}(z) = (N-2)! \frac{\Gamma\left(y - \frac{N}{2} + 1\right)}{\Gamma\left(y + \frac{N}{2}\right)} \sum_{k=1}^{N-1} \frac{\Gamma(k+1)}{\Gamma(N-k)} \frac{\Gamma\left(y + \frac{N}{2} - k\right)}{\Gamma\left(y - \frac{N}{2} + k + 1\right)}$$

Let the number of nodes be even $N = 2M$ such that $y = \sqrt{M^2 + z} \sim M + \frac{z}{2M}$ (provided $|z| < 2M$). The sum, denoted by S , can be rewritten

$$S = \sum_{j=-(M-1)}^{M-1} \frac{\Gamma(y+j)}{\Gamma(M+j)} \frac{\Gamma(M-j+1)}{\Gamma(y-j+1)}$$

such that

$$\begin{aligned} \varphi_{W_{2M}}(z) &= (2M-1)! \frac{\Gamma(y-M+1)}{\Gamma(y+M)} \\ &\times \frac{1}{2M-1} \sum_{j=-(M-1)}^{M-1} \frac{\Gamma(y+j)}{\Gamma(M+j)} \frac{\Gamma(M-j+1)}{\Gamma(y-j+1)} \end{aligned}$$

For large M ,

$$(2M-1)! \frac{\Gamma(y-M+1)}{\Gamma(y+M)} \sim (2M)^{-\frac{z}{2M}} \Gamma\left(\frac{z}{2M} + 1\right)$$

which makes us consider $z \rightarrow 2Mz$ since then, using [1, 6.1.47],

$$\begin{aligned} \varphi_{W_{2M}}(2Mz) &\sim (2M)^{-z} \Gamma(z+1) \\ &\times \frac{1}{2M-1} \sum_{j=-(M-1)}^{M-1} \left(1 + O\left(\frac{1}{M}\right)\right) \\ &\sim (2M)^{-z} \Gamma(z+1) \end{aligned}$$

Hence,

$$\lim_{N \rightarrow \infty} N^z \varphi_{W_N}(Nz) = \Gamma(z+1)$$

or equivalently,

$$\lim_{N \rightarrow \infty} E\left[e^{-(NW_N - \log N)z}\right] = \Gamma(z+1) \quad (7)$$

The inverse Laplace transform of $\Gamma(z+1)$ is a Gumbel distribution. Since $E[W_N] \sim \frac{\log N}{N}$ and $\sqrt{\text{var}[W_N]} \sim \frac{\pi}{\sqrt{2N}}$, we arrive at the asymptotic distribution (6) for the weight of the shortest path ($m = 1$).