The Skeleton of the Internet

Márton Csernai, András Gulyás, Gábor Rétvári,
Zalán Heszberger
Budapest University of Technology and Economics
Department of Telecommunications and Media Informatics
1117 Budapest, Magyar tudósok körútja 2. Hungary
Email: {csernai,gulyas,retvari,heszberger}@tmit.bme.hu

András Császár
TrafficLab, Ericsson Research
1117 Budapest, Irinyi út 4-20. Hungary
Email: Andras.Csaszar@ericsson.com

***Abstract*—Research works concerning AS (Autonomous Systems) level Internet topology measurements typically aim at obtaining near-complete maps of the AS structure. In this paper, we take a fundamentally different approach by inspecting several concurrently visible local views of the AS graph stored at individual BGP route servers. We find that each of these views exhibits the characteristic properties of complex graphs having power-law degree distribution, large clustering coefficient and the small world property. As a main contribution, the intersection of these views is investigated to identify the *skeleton* of the Internet consisting of edges seen by most of the ASes. Our measurements support the surprising claim that this skeleton is a scale-free complex network, having a giant connected component with a dense part in its heart forming the *critical AS level core*. We identify the edges in the skeleton as critical infrastructure, any changes of which induces an Internet-wide effect with BGP updates propagating to all ASes. Finally, we reinterpret the path inflation metric using the local view approach and show that local path inflation can be very diverse in different ASes.**

***Index Terms*—Internet, AS level topology, Network core, Measurements**

I. Introduction

Internet AS topology measurements have always generated considerable interest in the networking research community. Early results in this area significantly contributed to the triumph of the science of complex networks by providing evidence of heavy-tailed degree distribution and small-world phenomenon in the AS level Internet graph [1]. The efforts of obtaining more complete AS level maps fostered the development of diverse techniques for finding relationships between ASes, despite their reluctance to disclose routing information.

The most direct way of tracking down AS relationships is to investigate the BGP routing tables of publicly available "looking glass" and route servers. These servers store in their RIBs (Routing Information Base) the valid routing paths towards all IP prefixes as a sequence of AS-AS links. Since this information strongly depends on the network location of the route server, BGP based measurements usually aggregate the AS links extracted from numerous route servers during an extended period of time into a large common database [2]. Another class of measurements tries to override the limited availability of route records by using active probing from a large number of vantage points [3], [4], [5], [6]. Using the global AS topologies constructed from the combination of the above mentioned measurements, several interesting aspects of the AS level Internet structure have been uncovered. Besides the determination of the power law exponent of the degree distribution, significant research efforts have been made, e.g., to express path inflation metrics [7], [8], to infer AS business relationships [9], [10], [11], [12], and to the analysis of the highly connected Internet core [13], [14].

Unfortunately, accumulating AS links obtained from different vantage points at different time instances into a large *union* graph may restrict the application possibilities. This is because only a part of the union graph can be seen from any single point at any given time, and there is a substantial portion of paths that are not open to all ASes due to the strict business-oriented policy relationships governing the way ASes exchange routes, and hence, traffic between each other (see more on BGP policies in Section II).

In this paper, we take a fresh look at the Internet AS level topology. In particular, we address the problem of finding the subset of AS links that are visible from *any* arbitrary network location. For this analysis, BGP RIBs are taken as the authentic source of routing information at a given time and network location. By downloading these databases from various route servers at the same time and generating their *intersection* graph, we recover the concurrently visible AS links at a given instant (see Section III-A). We show that there exists a significant amount of "common knowledge" of around 27000 edges that can be found in a clear majority of the vantage points. This globally visible set of edges can be suggestively referred to as the *skeleton* of the Internet. Since the average number of AS links in our vantage points is around 45000, the skeleton contains approximately 60% of any local routing information. It is shown that, similarly to the global AS graph, the skeleton is also a scale-free complex network (Section III-C) having a giant connected component with a quite densely connected nucleus forming the *critical AS level Internet core* (Section III-D). Next, we reconsider the notion of path inflation (Section IV) that is usually calculated by comparing the length of the actual forwarding path to the shortest one extracted from the global AS level graph. We propose a local path inflation metric that does not rely on global AS maps and show that the local path inflation values have high variance among different vantage points (see Section IV). Finally, we conclude our work in Section V.

II. On the AS level graph of the Internet

In the AS level graph of the Internet the vertices are ASes that are connected via basically two types of edges. A provider edge stands for the customer-provider relationship between ASes, while peer links connect ASes that voluntarily exchange their traffic with each other in a settlement-free manner. According to the valley-free principle in BGP routing policy [9], an AS avoids propagating routes that use a customer AS to transit between two provider ASes. More formally, while exporting routing information to a provider or peer, an AS can export its own routes and its customer's routes, but usually does not announce its provider or peer routes, while can announce all of its routes to its customers.

Since not all the routes connecting two ASes propagate through the network, the complete AS graph cannot be seen from any single vantage point. This property renders the global topology graph to be a fictive theoretical object, which practical applications cannot explicitly rely on. Besides the significance of the global AS topology in, e.g., modeling frameworks or surveying aggregate network features, for several practical issues the local perspectives can be more appropriate. There is an ample amount of literature about the analysis of the global AS graph, whilst there are only few papers concentrating on the local consequences of the Internet routing policy [15], [16], [17], [18].

In the followings, we consider local views of the AS level topology originating from numerous public route servers. More specifically, a local view is the AS connectivity graph that can be extracted from the BGP RIB of a single route server. With the combination of these views we can identify a subset of edges —the so called *critical infrastructure*— as the set of the AS links that generate globally propagated BGP updates when changing. Furthermore, we can identify the core of this critical infrastructure as the *critical core*, which can be more meaningful as it is visible from arbitrary network location.

III. The skeleton of the Internet

In this section, we present the measurement studies concerning the different aspects of AS topologies seen from the local views of various ASes. Besides showing that the topology graphs exhibit the characteristic properties of complex graphs, we also reveal the critical infrastructure of the AS level Internet, and shed some light on important properties.

A. Measurement details

For the analysis of the AS Internet topology from different vantage points, our measurements were based on BGP routing databases downloaded from various public route servers [19]. The BGP RIB of a route server contains the corresponding AS topology map seen from the AS's location in which the route server resides in, and the AS peering relations representing the links in the topology graphs are extracted[1] from the *AS_PATH* field of the routing table entries [17]. Some high-level statistics of the considered ASes are summarized in Table I. Besides public AS numbers to identify the ASes, the table contains the location information on a country basis, the number of nodes and links of the topology graphs in the AS local views, the average path length (APL), the transitivity (or clustering coefficient) property [20], and finally the size of the RIBs (Routing Information Base) and a classification of the ASes according to [21].

[1]The source codes of our implementation can be downloaded from http://gonosz.tmit.bme.hu/topology_ meas/.

TABLE I: AS views data

ASN	Country	# of nodes	# of links	APL	Trans	RIB [KB]	AS class
AS553	Germany	33533	44875	4.541	0.1316	26,087	t2
AS812	Canada	33406	46281	4.449	0.1214	100,178	t2
AS852	Canada	33403	44199	4.578	0.1254	26,956	t2
AS2547	Hungary	33495	44553	4.560	0.1068	54,839	edu
AS3257	Germany	33393	44281	4.509	0.1127	26,019	t2
AS3549	USA	33562	45948	4.459	0.1092	271,537	t1
AS4323	USA	33437	46188	4.516	0.1226	148,155	t2
AS5413	UK	33403	45171	4.528	0.1388	115,804	t2
AS5713	South Africa	33425	46010	4.207	0.1650	102,601	nic
AS6539	Canada	33418	45789	4.559	0.1521	63,887	t2
AS6730	Switzerland	33441	45997	4.448	0.1357	76,230	t2
AS6939	USA	33706	47433	4.457	0.1613	607,924	t2
AS7018	USA	33382	48839	4.338	0.1208	488,493	t1
AS7474	Australia	33428	45104	4.457	0.1419	102,714	t2
AS8301	Gibraltar	33378	44476	4.516	0.1193	26,424	t2
AS11260	Canada	33405	46983	4.312	0.1431	36,977	t2
AS13645	USA	33459	44084	4.556	0.1426	27,801	t2
AS15290	Canada	33430	45636	4.514	0.1723	219,201	t2
AS20121	Brazil	33568	54108	4.158	0.2109	88,083	nic
AS21229	Hungary	33409	44361	4.514	0.1480	57,114	t2
AS22548	Brazil	33845	57254	4.125	0.2330	365,245	nic
Skeleton	NA	22808[a]	27303	5.781	0.0720	NA	NA

[a]vertices having one or more edges

This type of analysis has received several criticisms. It has been pointed out that a significant portion of AS-AS edges cannot be explored from BGP route servers [22], [17]. Additionally, some partially configured BGP routers (e.g. for the purpose of saving resources) may collect only a small portion of the obtainable path information using default entries instead. Such routers are to be avoided as they may invalidate measurement results. It is proven that many further AS peering links, hidden from public BGP routing databases, can be discovered using traceroute based methodologies [3], [4], [5]. Nevertheless, the measurements in this paper are based solely on BGP extracted topologies, since the analysis specifically focuses on topological data locally available at the given ASes.

Downloading of the routing tables from the different sources was done concurrently ensuring that a real snapshot is taken of the collective BGP knowledge. The small variance in the number of ASes seen from each view demonstrates that properly configured BGP speaking routers obtain routing information for essentially the same set of Internet domains at any given time, regardless of their AS class or location.

B. Degree distributions and correlations

The degree distribution of the graphs extracted from the downloaded BGP datasets is depicted in Figure 1. We observe that each local view has a power-law tail with essentially the same exponent around 2.2. From Table I, one can infer that the clustering coefficients are actually the same for all local views and the small world property also holds. This means that from each local perspective the AS level Internet graph is perceived as a complex graph, with all its characteristic properties.

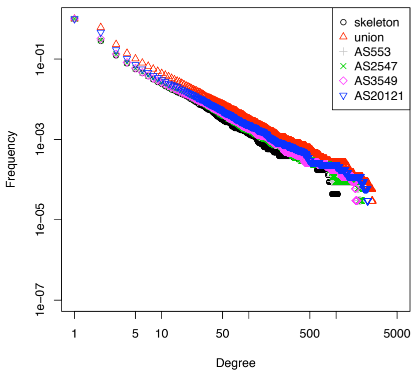

Fig. 1: Degree distributions of the intersection graph, union graph and the different local views of the AS topology.

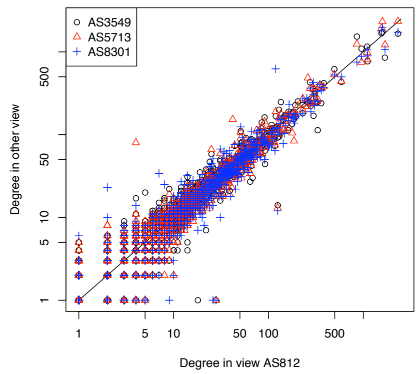

Fig. 2: AS degree correlation between different local views.

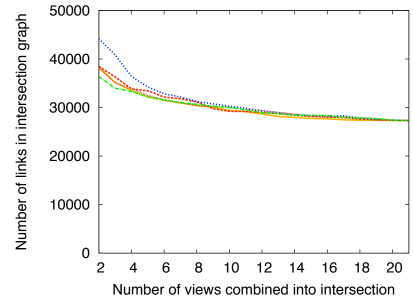

Fig. 3: Number of links in intersection graph vs. number of views combined in the intersection graph. The curves with different colors show different sequencing of the ASes.

Despite the similar behavior of the degree distributions, the exact degree of a given AS can vary in the local views depending on the applied BGP policies. For example, to guarantee the valley-free property of Internet routes peer relationships are only propagated downstream to customer ASes, rendering them invisible from the provider's view. In order to explore the degree diversity of the nodes in different local views, we calculated the degree correlations of the views of each AS pairs in our dataset. As seen in Figure 2, the degree correlation diagram places an AS in the x and y axis according to its degree in the two views. To make the plot more general, a randomly chosen AS is paired with three other ASes of diverse geographical locations and AS classes[2].

Figure 2 reveals the strong degree correlation between the local perspectives. This means that if an AS is a hub having a large number of AS relationships, then it is seen as a hub from arbitrary local perspective. Nevertheless, the exact degree can vary significantly especially in case of lower degree ASes. For example the triangle symbol at $(4, 90)$ represents that an AS has a degree of 4 as visible form $AS812$ while its degree is 90 from the point of $AS5713$.

C. The critical infrastructure

In this section, we investigate the intersection of the downloaded local views. The intersection of two views include only those links (along with end nodes) that are present in both. Intersecting the resulting graph with the remaining local views one by one gives the the common part shared by all local views. Figure 3 shows the edges remaining in the k-th step of this process for some randomly chosen iteration sequence. It turns out that the exact order in which we take the intersection of the views does not impact the characteristics of the plot substantially. The resulting set of common edges can be considered as a *skeleton* that is "dressed" by additional AS links in the particular local views. We deem this skeleton to be a critical infrastructure of the Internet, as it consists of the AS links contained in at least one AS path of essentially any AS in the Internet, whose failure, flapping, or any other status change propagates throughout the entire Internet thus significantly contributing to BGP churn.

[2]We generated such plots for all possible AS pairs and got statistically identical results.

Perhaps the most important property of this skeleton is that it also is a complex network, with power law degree distribution (Figure 1) and high clustering (Table I), which is coherent with the strong degree correlation found between the local perspectives. Another interesting property is that the skeleton contains a surprisingly large number of AS links (about 60%). Still, the skeleton is not a connected graph. Figure 4 shows the size distribution of the connected components. We can see that there is a giant connected component, consisting of approximately 20000 ASes, and there are numerous smaller components with increasing frequency as size reduces. Figure 5 shows these smaller components, with separate nodes with zero degree and 1 degree pairs of nodes omitted for the sake of visibility. One can see that these smaller components are all trees with a central node.

In some respect, the existence of so many AS links in the skeleton plus the presence of the giant connected component are somewhat unexpected findings, considering that BGP in no ways fosters the disclosure of AS-AS connections to *all* ASes. And indeed, there are many reasons for which the skeleton may fall into separate components, like extensive prefix deaggregation for traffic engineering purposes, site multi-homing with selective downstream distribution of provider links, prefix aggregation for routing table compaction, etc. Despite of this, there still seems to be a huge body of shared knowledge in the Internet routing system.

Finally, we present the geographical visualization of the skeleton in Figure 6. Similarly to the geographical representation of the Internet AS level topology in [4] the ASes are placed on a circular graph where the radial distance is

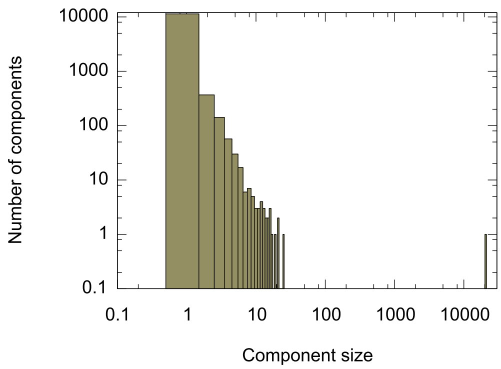

Fig. 4: Distribution of components in intersection graph according to their sizes

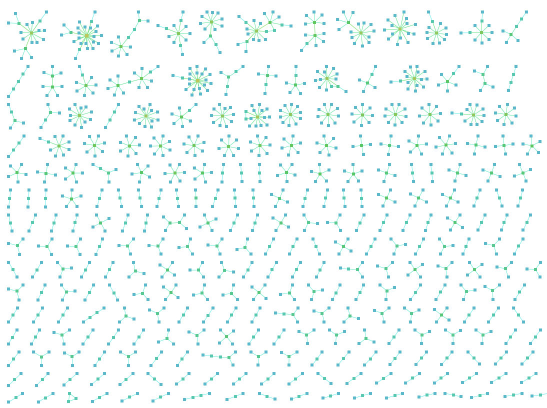

Fig. 5: The structure of the smaller components in the skeleton.

proportional with the AS degree and the azimuth corresponds to the meridian angles of the geographic location.

D. The critical core of the Internet

From the survey presented in the previous sections, we recognize that there is a skeleton of AS links that are present in all local views of the AS level topology. Here, we analyze the core of the skeleton according to the widely known k-core algorithm [23], that has been used many times for AS level core analysis [24], [13]. Comparing to the previously investigated Internet cores based on the near-complete AS topologies, we believe that the k-core is more meaningful since it is considered as the core from any point of the network. According to our datasets, 36 ASes constitute the 5-core of the skeleton graph. It contains 6 of the 12 generally considered tier-1 ASes, 27 tier-2 ASes and 3 company ASes according to the various AS Taxonomy datasets [21], [25]. In this core there are large and small ASes as well, the number of their advertised prefixes greatly varies from 1 to a few thousands. The geographical distribution of the critical core generated with GeoPlot [4] can be seen in Figure 7. If we compare the k-core of the skeleton with that of the union graph, we find that it contains more than twice as many ASes. Interestingly, not all

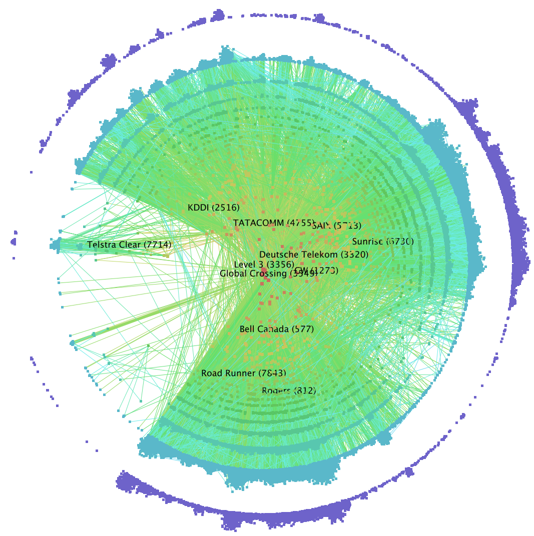

Fig. 6: The skeleton of the Internet plotted on a geographically mapped plane.

the ASes seen in the skeleton k-core are contained in the union k-core, which reveals that loosely connected ASes can also play strategic roles in the Internet infrastructure. Nevertheless there is a subset of 16 ASes, which is contained in both cores. This subset contains the 6 previously mentioned tier-1 and also other tier-2 ASes of varying sizes and geographical locations, which identifies the most popular Internet domains.

IV. Path inflation revisited

The notion of path inflation captures the fact that AS routes used for actual packet forwarding in the Internet takes longer than the shortest AS path available. The inflation is caused by the routing policies restricting inter-domain links to be used for specific (eg. non-local) traffic. In the traditional approach [7], [8] shortest paths between AS pairs are determined by using the (near-)global AS topology map [22] assembled by collecting data from many AS routing databases or by other means of topology measurements. By considering only a subgraph locally available from an AS and recalculating shortest paths accordingly, one may assume that, the path inflation fades away. In contrast, our interesting finding is that such a redefined local path inflation metric is still seems worthwhile to investigate. Extensive measurements show that local path inflation can be very diverse in different ASes. Furthermore, we can calculate the inflation effect of restricting the knowledge of the AS topology map to local data. In Figure 8 the average global and local path inflation is plotted for all investigated AS local views.

V. Conclusions

In the paper local views of AS graphs were investigated. By finding the intersection of the graphs the *skeleton* of the Internet is presented. It is shown that similarly to the local views, the skeleton also shows the characteristic properties of

Fig. 7: The critical core of the Internet plotted on a political world map.

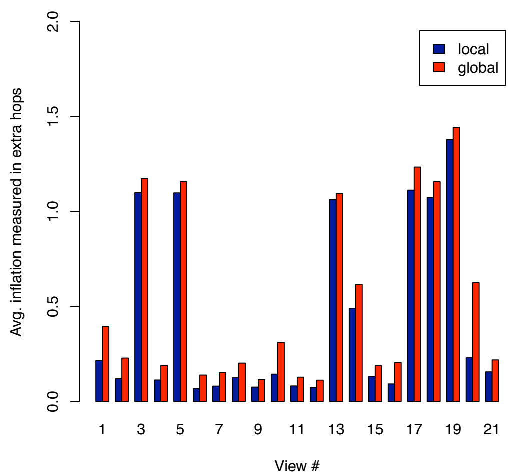

Fig. 8: Difference between the measured average path inflation based on local views and the global AS level graph.

complex graphs. The intersection AS graph, also referred to as the critical infrastructure, was discovered to contain the bulk of the locally available AS links with a centrally forming giant connected component. We also matched the skeleton to the global AS graph and made some remarkable observations. Finally, the properties of a local version of the path inflation metric is investigated and compared to the original.

References

[1] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proceedings of the ACM SIGCOMM*, 1999.

[2] University of Oregon Route Views Project. [Online]. Available: http://www.routeviews.org

[3] Y. Shavitt and E. Shir, "DIMES: Let the internet measure itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, p. 74, 2005.

[4] CAIDA Project. [Online]. Available: http://www.caida.org/

[5] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, "iplane: An information plane for distributed services," in *Proceedings of the OSDI*, 2006, pp. 367–380.

[6] N. T. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *Proceedings of the SIGCOMM*, 2002, pp. 133–145.

[7] N. Spring, R. Mahajan, and T. Anderson, "Quantifying the causes of path inflation," in *In Proceedings of ACM SIGCOMM*, 2003.

[8] L. Gao and F. Wang, "The extent of as path inflation by routing policies," in *In Proceedings of IEEE GLOBECOM*, 2002.

[9] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions on Networking*, vol. 9, pp. 733–745, 2000.

[10] X. Dimitropoulos, D. Krioukov, B. Huffaker, K. Claffy, and G. Riley, "Inferring as relationships: Dead end or lively beginning?" *Lecture Notes in Computer Science*, vol. 3503, pp. 113–125, 2005.

[11] H. Chang, S. Jamin, and W. Willinger, "Inferring AS-level Internet topology from router-level path traces," in *Proc. Workshop on Scalability and Traffic Control in IP Networks, SPIE ITCOM Conference*, 2001.

[12] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "On inferring AS-level connectivity from BGP routing tables," in *Proceedings of the IEEE Infocom*, 2002.

[13] G. Zhang, G. Zhang, Q. Yang, S. Cheng, and T. Zhou, "Evolution of the internet and its cores," *New Journal of Physics 10 123027*, 2008.

[14] J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani, "k-core decomposition: a tool for the analysis of large scale internet graphs," *CoRR*, vol. abs/cs/0511007, 2005.

[15] J. G.-V. José M. Barceló, Juan I. Nieto-Hipólito, "Study of internet autonomous system interconnectivity from bgp routing tables," *Computer Networks*, vol. 45, pp. 333–344, 2004.

[16] L. Subramanian, S. Agarwal, J. Rexford, and Y. H. Katz, "Characterizing the internet hierarchy from multiple vantage points," in *In Proceedings of IEEE INFOCOM*, 2002.

[17] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "Towards capturing representative AS-level Internet topologies," in *Computer Networks Vol. 44*, 2004, pp. 737–755.

[18] K. Chen, C. Hu, W. Zhang, Y. Chen, and B. Liu, "On the eyeshots of bgp vantage points," in *In Proceedings of IEEE GLOBECOM*, 2009.

[19] Public route server list. [Online]. Available: http://traceroute.org/

[20] D. Watts and S. Strogatz, "Collective dynamics of small-world networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.

[21] X. Dimitropoulos, D. Krioukov, G. Riley, and K. Claffy, "Revealing the Autonomous System Taxonomy: The Machine Learning Approach," in *Passive and Active Measurements Workshop (PAM), Adelaide, Australia*, 2006.

[22] R. Cohen and D. Raz, "The internet dark matter–on the missing links in the as connectivity map," in *Proceedings of the IEEE INFOCOM*, 2006.

[23] V. Batagelj and M. Zaveršnik, "Generalized cores," *Arxiv preprint cs/0202039*, 2002.

[24] M. Gaertler and M. Patrignani, "Dynamic analysis of the autonomous system graph," in *IPS 2004, International Workshop on Inter-domain Performance and Simulation, Budapest, Hungary*, 2004, pp. 13–24.

[25] Robtex homepage. [Online]. Available: http://www.robtex.com/