

Beyond Centrality - Classifying Topological Significance Using Backup Efficiency and Alternative Paths

Yuval Shavitt and Yaron Singer

Tel Aviv University, Tel Aviv, Israel
yaronsin@post.tau.ac.il

Abstract. In networks characterized by broad degree distribution, such as the Internet AS graph, node significance is often associated with its degree or with centrality metrics which relate to its reachability and shortest paths passing through it. Such measures do not consider availability of efficient backup of the node and thus often fail to capture its contribution to the functionality and resilience of the network operation. In this paper we suggest the Quality of Backup (QoB) and Alternative Path Centrality (APC) measures as complementary methods which enable analysis of node significance in a manner which considers backup. We examine the theoretical significance of these measures and use them to classify nodes in the Internet AS graph while applying the BGP valley-free routing restrictions. We show that both node degree and node centrality are not necessarily evidence of its significance. In particular, some medium degree nodes with medium centrality measure prove to be crucial for efficient routing in the Internet AS graph.

Keywords: Internet topology, network analysis, Internet AS graph.

1 Introduction

The topological study of networks appears in a wide spectrum of research areas such as physics [3], biology [13], and computer science [10]. In research of the Internet, node significance classification has received attention in past studies [18,4,5] and was treated in two different contexts: study of the Internet resiliency against attacks and failures [3,12,14,7] and identification of the Internet core nodes as well as significance categorization [18,5]. Both study threads were conducted at the Internet AS level graph.

Several attempts have been made in the past to characterize the core of the Internet AS graph. In [18] the most connected node was used as the natural starting point for defining the Internet's core. Other ASes were also classified to four shells and tendrils that hang from the shells, where ASes in shells with a small index are considered more important than ones in higher indices. Further work has dealt with classification of nodes into few shells with decreasing importance [8,17]. In recent study [5] k -shell graph decomposition was used to classify nodes by importance to roughly 40 layers of hierarchical significance. The

k -shell classification, based on the node's connectivity, identified over 80 ASes as the Internet core, some of which with medium degrees. Almost exclusively, an attempt to rank ASes by metrics other than node degree was done by CAIDA [1], where the 'cone' of the node was used to determine its importance, namely the number of direct and indirect customers of the AS.

As network functionality is often measured by connectivity and vertex distances in the graph used as its model, measures which credit vertices connected to a relatively large number of vertices at relatively short distances are often used as significance indicators [9,3]. However, inadequate consideration of backup by such measures often overshadows significance in context of its contribution to functionality and resilience of the network. Existence of backup raises question regarding a node's significance since failure of a node with backup does not effect connectivity nor does it increase path lengths in the network and therefore the effect of failure in such instances is minimal. Furthermore, existence of backup denies exclusivity of the information passing through the node in the network. Since nodes can have backups of various qualities, measures of backup efficiency and topological significance which considers backup are crucial for analysis of network functionality.

In this paper, we suggest two complementary measures which capture a node's contribution to the network's functionality: the Quality of Backup (QoB) and the Alternative Path Centrality (APC). The QoB measures backup quality of a vertex regardless of its centrality or effect on the functionality of the network, enables comparison of backup efficiency between vertices in the graph as well as between vertices from different graphs, and can thus serve as a universal measure for backup. The APC measures functionality which considers both backup quality and centrality of vertices in graphs and therefore enables analysis of nodes' significance in a wider context in comparison to other topological measures.

Our starting point for the node significance classification problem is examination on levels of theoretical abstraction, and then evaluation of our results on the Internet AS graph. Since failure of a node on the AS level is possible [7] though highly unlikely, applying APC and QoB on the Internet AS graph allows a unique insight to the Internet as oppose to quantifying effects of failures. On the AS level, centrality which considers backup reveals significance in context of potential information which exclusively passes through a node, and its backup quantifies the dependency of its customers on its transit services. In our study we use APC to identify the most significant nodes in the Internet AS graph, and show that these are not necessarily members of the Internet core. In accordance with properties of APC, it is not surprising that the largest ASes in the core, such as UUNET and Sprint, also have very high APC values due to the large number of customer ASes. However, small networks with poor backup like the French research network RENATER, and the GEANT and Abilene academic backbones which have degrees as low as 51 (RENATER) and low centrality values, have very high APC values as well.

The rest of this paper is organized as follows. The next section discusses the concept of backup in networks and introduces the QoB as a measure of universal

backup efficiency. Section 3 provides detail of the APC construction and discusses its properties. In Section 4, we discuss modifications of our methods in order to maintain relevance in the Internet AS graph model. Section 5 holds our analysis of the Internet AS graph using the modified measures in comparison to previous works. Finally, we summarize and discuss our work in section 6.

2 Quantifying Backup Efficiency

In our attempt to quantify backup in networks we observed that the quality of backup of a given vertex in the graph is determined by the number of direct children covered by a set of backup vertices, and the efficiency of reaching this set by the set of direct parents. For a vertex $v \in V$, we define the set of children C_v , set of parents P_v , and the backup set B_v , as follows:

$$\begin{aligned} C_v &= \{u \in V | (v, u) \in E\}, \\ P_v &= \{u \in V | (u, v) \in E\}, \\ B_v &= \{w \in V | \exists u \in C_v : (w, u) \in E\}. \end{aligned}$$

Clearly, in instances of undirected graphs $C_v \equiv P_v$, and the discussion which follows remains relevant for these instances as well.

For $u, w \in V$, we use $\delta(u, w)$ to denote the shortest path distance between u and w in G . The shortest distance, $\delta(u, w)$, can be calculated by any set of rules, e.g., based on additional annotations on the graph edges, and is not limited to minimum hop. By convention, if u cannot reach v through any path in G , then $\delta(u, v) = \infty$. We also use $\delta_v(u, w)$ to represent the distance of the shortest path which bypasses v from u to w .

Let $G = (V, E)$ be a directed or undirected graph where V is the set of vertices and E is the set of edges. For $v \in V$, let P_v be the set of v 's direct parents and let C_v be the set of v 's direct children. The Quality of Backup of v , denoted $\rho(v)$ is:

$$\rho(v) = \frac{\sum_{u \in P_v} \sum_{w \in C_v} (\max\{\delta_v(u, w) - 1, 1\})^{-1}}{|P_v| \cdot |C_v|}$$

The rationale behind this measure is the following. To measure backup efficiency of a given vertex, it is enough to examine the cost of re-routing paths from its set of parents to its set of direct children. Note that $\max\{\delta_v(u, w) - 1, 1\} = \delta_v(u, w) - \delta(u, w) + 1$ for all pairs $\langle u, w \rangle$, where $u \in P_v$ and $w \in C_v$. For $v \in V$, it is easy to see that $\rho(v) = 1 \iff \forall \langle x, u \rangle \in P_v \times C_v \exists w \in B_v : (x, w) \in E \wedge (w, u) \in E$. Also, note that $\rho(v) = 0 \iff \delta_v(x, u) = \infty \forall \langle x, u \rangle \in P_v \times C_v$. Thus, $\rho : V \rightarrow [0, 1]$, and returns 1 for vertices with perfect backups and 0 for vertices with no backup. Formal implementation of QoB on unweighted directed graphs is presented in the figure below. Here, BFS_v denotes the BFS algorithm which bypasses a vertex v , and $\vec{\delta}_v(u)$, denotes the vector of shortest path distances from u to all the vertices in the graph, which bypass v .

```

QoB( $v, G$ )
 $\rho \leftarrow 0$ 
for all  $u \in P_v$  do
     $\bar{\delta}_v(u) \leftarrow \text{BFS}_v(u, G)$ 
    for all  $w \in C_v$  do
         $\rho \leftarrow \rho + (\max\{\delta_v(u, w) - 1, 1\})^{-1}$ 
 $\rho \leftarrow \frac{\rho}{|P_v| \cdot |C_v|}$ 
return  $\rho$ 

```

The following theorem shows the QoB measure indeed enables local measurement of a vertex’s backup in the graph.

Theorem 1. *For $G = (V, E)$, for a vertex $v \in V$ with $P_v \neq \emptyset$ and $C_v \neq \emptyset$, $\rho(v)$ monotonically increases with respect to rise in backup efficiency.*

Proof. For $u \in P_v$ and $w \in C_v$, assume that $\delta_v(u, w) = c$ in G , where $1 < c \leq \infty$. Construct G' by adding some edge $e \notin E$, such that $\delta'_v(u, w) = c' < c$, where $\delta'_v(u, w)$ represents the distance between u and w bypassing v in G' . Therefore, $\frac{1}{\delta'_v(u, w)} > \frac{1}{\delta_v(u, w)}$, and $(\max\{\delta'_v(u, w) - 1, 1\})^{-1} \geq (\max\{\delta_v(u, w) - 1, 1\})^{-1}$ (where equality holds only when $c = 2$). It thus easily follows that $\rho'(v) > \rho(v)$, where $\rho'(v)$ is the QoB measure of v in G' . □

3 Alternative Path Centrality

The above section discusses backup efficiency of a vertex regardless of centrality considerations. In an attempt to quantify significance, note that centrality of a node (its ability to reach a relatively large number of nodes efficiently) also plays a vital role in analysis: a node which has relatively efficient backup can be crucial to the network’s functionality due to its high centrality, while a node with poor backup and low centrality can have little effect on functionality in the network. The Alternative Path Centrality (APC) measure presented in this section enables quantifying topological contribution of a node to the functionality of the network as it considers both centrality and backup efficiency.

Given a graph $G = (V, E)$ as above and $u \in V$, the topological centrality measure used here, denoted χ , where $\chi : V \rightarrow \mathbb{R}$ is:

$$\chi(u) = \sum_{w \in V \setminus \{u\}} \frac{1}{\delta(u, w)}$$

Clearly, $0 \leq \chi(u) \leq |V| - 1 \quad \forall u \in V$.

For a vertex $u \in V$, the value of $\chi(u)$ depends on the number of vertices connected to u and their distances from it; χ monotonically increases with respect to both centrality and connectivity of the vertex. Thus, in relation to other vertices in the graph, high χ values are obtained for a vertex which is connected to a large

number of vertices at short distances. Symmetrically, a vertex connected to a small number of vertices at large distances yields low χ values. These properties make the χ function a favorite candidate for measuring vertices' centrality in the network. In [11] for a network G , the average of χ values was used to define the *efficiency* of the network. Similar topological measures have also been used in [3] and in [9] to study functionality in complex networks.

For $G = (V, E)$, The APC value of $v \in V$, denoted $\varphi(v)$ is:

$$\varphi(v) = \sum_{u \in V \setminus \{v\}} \chi(u) - \sum_{u \in V \setminus \{v\}} \chi_v(u)$$

Where χ_v denotes centrality values calculated in the graph using alternative paths which bypass v .

The rationale behind APC is simple. In instances where network functionality is determined by shortest paths and connectivity, the significance of a node v to the network's functionality can be measured by its effect on these criteria. Computing the difference between vertices' topological centrality using v , and topological centrality bypassing v , enables witnessing v 's exclusive contribution to the network's functionality.

The algorithm presented below is a simple implementation of APC using the Breadth First Search (BFS) algorithm for unweighted directed graphs.

```

APC( $v, G$ )
 $\varphi \leftarrow 0$ 
for all  $u \in V \setminus \{v\}$  do
     $\bar{\delta}(u) \leftarrow \text{BFS}(u, G)$ 
     $\bar{\delta}_v(u) \leftarrow \text{BFS}_v(u, G)$ 
     $\chi_\Delta \leftarrow 0$ 
    for all  $w \in V \setminus \{v, u\}$  do
         $\chi_\Delta \leftarrow \frac{1}{\bar{\delta}(u, w)} - \frac{1}{\bar{\delta}_v(u, w)}$ 
return  $\varphi$ 

```

Using the BFS algorithm, the overall computational complexity of APC is $O(|V| \cdot |E|)$. For weighted graphs, one can substitute the BFS algorithm with a single-source shortest path algorithm for non-negative weighted graphs, such as Dijkstra's algorithm [6] and achieve $(|V| \cdot (|V| \cdot \log |V| + |E|))$ running time.

We conclude our discussion of the APC properties with the following theorem which shows that APC properly considers both centrality and backup of a vertex in the graph.

Theorem 2. *For $G = (V, E)$ and $v \in V, C_v \neq \emptyset$, $\varphi(v)$ monotonically increases with respect to rise in topological centrality and decrement in backup quality.*

Proof. To prove $\varphi(v)$ monotonically increases with respect to rise in centrality, let $\chi(v) < |V| - 1$, and $w \in V$ be a vertex for which $1 < \delta(v, w) \leq \infty$. Let $e \notin E$ be some edge for which $\delta'(v, w) < \delta(v, w)$, where $\delta'(v, w)$ denotes the shortest

path distance in $G' = (V, E \cup \{e\})$, and e does not create new alternative paths to w in G' (otherwise backup efficiency increases). We show that $\varphi(v) < \varphi'(v)$, where $\varphi'(v)$ denotes the APC value of v in G' . For all $x \in V$, which reach w through v , $\delta'(x, w) < \delta(x, w)$ and $\delta'_v(x, w) = \delta_v(x, w)$. For all such vertices, x , we have:

$$\frac{1}{\delta'(x, w)} - \frac{1}{\delta'_v(x, w)} > \frac{1}{\delta(x, w)} - \frac{1}{\delta_v(x, w)}$$

and it therefore follows that $\varphi(v) < \varphi'(v)$.

To prove monotonic increase with respect to decrement in backup quality, assume that some edge $e \notin E$ has been added to G , such that $\rho(v)$ increases. We again denote $G' = (V, E \cup \{e\})$, and use similar notation as above. We therefore assume $\rho'(v) > \rho(v)$. Specifically, there is some pair $\langle u, w \rangle \in P_v \times C_v$ such that $\delta'_v(u, w) < \delta_v(u, w)$. For this pair we have:

$$\frac{1}{\delta'(u, w)} - \frac{1}{\delta'_v(u, w)} < \frac{1}{\delta(u, w)} - \frac{1}{\delta_v(u, w)}$$

It trivially follows that $\varphi'(u) < \varphi(u)$, and concludes proof of the theorem. \square

4 Adaptation of APC and QoB for the Directed AS Graph

To apply QoB and APC on the Internet, we have adjusted these measures to conform to the model of the AS graph and specifically to the routing restriction which it imposes. We begin with a brief description of the AS graph model.

The Internet AS Graph

The Internet today consists of tens of thousands of networks, each with its own administrative management, called Autonomous Systems (ASes). Each such AS uses an interior routing protocol (such as OSPF, RIP) inside its managed network, and communicates with neighboring ASes using an exterior routing protocol, called BGP. The graph which models inter-connection between ASes in the Internet is referred to as the Internet AS graph. Since the ASes in the Internet are bound by commercial agreements, restrictions are imposed on the paths which may be explored. The commercial agreements between the ASes are characterized by customer-provider, provider-customer and peer-to-peer relations. A customer pays its provider for transit services, thus the provider transmits all packets to and from its customers. The customer, however, will not transit packets for its provider. Specifically, a customer will not transit packets between two of its providers, or between its provider and its peers. Peers are two ASes that agree to provide transit information between their respective customers.

In pioneering work, Lixin Gao [8] has deduced that a legal AS path may either be an *up hill* path, followed by a *down hill* path, or an *up hill* path, followed by a peering link, followed by a *down hill* path. An *up hill* path is a sequential set, possibly empty, of customer-provider links, and a *down hill* path is a sequential

set, possibly empty, of provider-customer links. Therefore a legal route between ASes can be described as a *valley free* path. A peering link can be traversed only once in each such path, and if it exists in the path it marks the turning point for a *down hill* path.

The ASQoB and ASAPC Measures

Since transitivity is not immediate in the AS graph, the QoB requires two cardinal adjustments to maintain relevance. Consider the AS graph $G = (V, E)$, and some $v \in V$, for which we wish to obtain $\rho(v)$ in G . Let $u \in P_v$ and $w \in C_v$. The first adjustment is to consider the pair $\langle u, w \rangle \iff u$ can reach w through v using a legal AS path. Since the BFS algorithm does not consider the *up*, *down*, and *peer* labels, *valley free* paths are not exclusively discovered, and it cannot be used to measure minimum-hop distances in the AS graph. For this, we use the ASBFS algorithm [16] which discovers *valley free* shortest paths from a source vertex in the unweighted AS graph in linear time.

In order to provide motivation for the second adjustment required, we present the following example. Consider the graph illustrated in Fig. 1. In quantifying the QoB of $v \in V$ Suppose a vertex $u \in P_v$ has reached a vertex $w \in C_v$ through an *up hill* path through v , though by using the alternative path through the vertex $b \in B_v$, u now reaches w through a *down hill* path. All vertices in C_w which are reached through an *up hill* path (x in this example), are now unreachable to u as this creates an illegal AS path. Therefore, to factor this into the QoB measure in the AS graph, we use the following strategy. For all vertices $w \in C_v$ we scan for vertices $x \in C_w$ which are reachable from v through legal AS paths, and consider the pairs $\langle u, x \rangle \in P_v \times C_w$ as well. The ASQoB algorithm is described in the figure below. We denote by R_{uv} the set of reachable children of v from u in accordance to policy based routing in the AS graph.

```

ASQoB( $v, G$ )
 $\rho \leftarrow 0$ 
for all  $u \in P_v$  do
     $\bar{\delta}_v(u) \leftarrow \text{ASBFS}_v(u, G)$ 
    for all  $w \in R_{uv}$  do
         $\rho \leftarrow \rho + (\max\{\bar{\delta}_v(u, w) - 1, 1\})^{-1}$ 
        for all  $x \in R_{vw}$  do
             $\rho \leftarrow \rho + (\max\{\bar{\delta}_v(v, x) - 1, 1\})^{-1}$ 
 $\rho \leftarrow \frac{\rho}{\sum_{u \in P_v} \sum_{w \in R_{uv}} |R_{uv}| + |R_{vw}|}$ 
return  $\rho$ 

```

Drawing its strength from the properties of the QoB measure, the ASQoB remains faithful to the principles of measuring backup efficiency in the AS graph. For $v \in V$, as reachable children are scanned in two levels, we are guaranteed that $\rho(v) = 1 \iff v$ has a perfect backup which does not disqualify legal AS paths.




Fig. 1. Illustration of an instance in an AS graph where a direct child can be reached through a backup vertex, though its paths cannot be used. Direction of an edge implies it is an *up* edge, and for each *up* edge a *down* edge in the opposite direction exists (not portrayed). Here, *b* serves as a backup for *v*. In accordance to the *valley free* restrictions, *u* can reach *w*, though cannot reach *x* through *b*.

Substituting the BFS algorithm with its analogous for the AS graph, ASBFS, applying APC on the AS graph is immediate. The calculation of a shortest path, δ , is done while considering the *valley free* routing and all the properties discussed in section 3 hold.

5 Analyzing the Directed AS Graph

We used the combined data from the DIMES [15] and RouteViews [2] projects for week 11 of 2006. The AS graph is comprised of 20,103 ASes and 57,272 AS links. We approximate the AS relationship by comparing the *k*-core index [5] of two ASes and taking the one with the highest *k*-core index as the provider of the other. If the *k*-core indices of two ASes are equal, the ASes are treated as peers. While we are aware that our approximation involves some inaccuracies, there is no known error free algorithm for this task. Since the majority of the interesting ASes are within the range of AS numbers 1-22,000, we present results of these 11,407 ASes along with results of ASes with degree higher than 40 of the rest of AS graph.

We first show that while centrality is closely related to the node degree in the AS graph, our APC criteria captures significance which is not necessarily associated with high degree. Fig. 2 shows the centrality values of AS nodes averaged by their degree on a log-log scale. There is almost a monotonic increase in centrality for nodes of degree above 300, and the close relationship between centrality and degree is evident. On the other hand, Fig. 3 shows there is a clear monotonic (and fairly linear in the log-log scale) increase in the average APC value from degree 3 up to around 40, and above this value the number of nodes with the same degree is below 10. Therefore any one ‘outlayer’, namely a node with extreme high or low APC values, can change the average significantly.

To display the relationship between high centrality and high APC we plot the degree and APC values of the nodes with the highest centrality (Fig. 5) and the




Fig. 2. Average centrality as a function of its degree




Fig. 3. Average APC as a function of its degree

degree and centrality of the nodes with the highest APC values (Fig. 4). The five ASes with the highest degree, 701 (UUNET), 7018 (AT&T), 1239 (Sprint), 3356 (Level3), and 174 (Cogent), are also the five ASes with the highest centrality. These are the largest tier-1 providers. In contrast, only UUNET is in the top ten APC list, mainly due to its high number of peer ASes; Sprint and Cogent also have high APC values. These three tier-1 providers support many stub ASes but have relatively low backup measure (0.7–0.75) which explain their high APC values. Level3, which has high centrality, has low APC value because it has a rather high QoB around 0.82. This means that although Level3 (3356) plays a central role in Internet routing, it may be replaced through alternative routes and thus is not as important as the previous three nodes. The next nodes with high centrality are 3549 (GBLX), 2914 (Verio), 7132 (SBC), 6461 (Abovenet), and 12956 (Telefonica). These are all tier-1 providers or major providers in Europe.

For the nodes with the highest APC values the picture is different: while UUNET (701) has the fourth largest APC value, many of the high locations in the list are captured by medium sized ASes with poor (and sometimes extremely poor) backup. Through study of the QoB distribution in the AS graph we have learned that there is a large concentration around 1, which is a testament of perfect backup. The median QoB value is 0.9799, and a large majority of the nodes have QoB values above 0.95. The nodes ranked first, third, and eighth in the top APC list are educational networks: GEANT (20965) in Europe, ENA (11686) in the USA, and RENATER (2200) in France (Abilene the US research network was ranked eleventh). The other group of nodes is of medium size providers, France Telecom (3215), YIPES (6517), Ukraine Telecom (6849), and ServerCentral (23352), each appears to have high APC values due to a different reason. France Telecom, YIPES and UKR Telecom have extremely low QoB, while ServerCentral connects remote locations that may not have efficient alternative paths. Statistics of nodes with highest APC values are displayed in table 1.

Fig. 6 shows the distribution of the APC values in the AS graph (note the truncation of the first column). The APC distribution is shown to have a long but narrow tail with only a few nodes with very high APC values, these nodes are




Fig. 4. The degree and centrality of the nodes with the highest APC values




Fig. 5. The degree and APC of the nodes with the highest centrality values

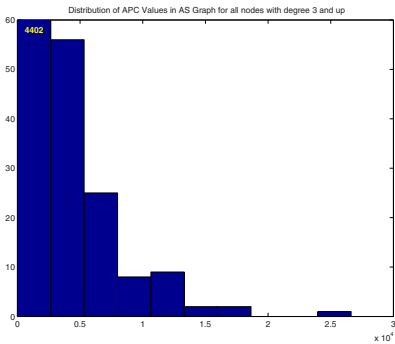


Fig. 6. A histogram of the APC values for nodes of degree greater than 2. The first bin holds 4402 ASes and was truncated.




Fig. 7. A histogram of the backup values for nodes of degree greater than 2

scattered almost over the entire degree range, starting with nodes with degree of just above 50 (see Fig. 4 and Table 1).

The QoB distribution shown in Fig. 7 has a large concentration around 1, which is a testament of perfect backup. The median value is 0.9799, and as the histogram shows a large majority of the nodes have QoB values above 0.95.

To discuss our results in comparison to other measures of node significance, we refer to table 2 which shows the top ten nodes in the CAIDA ranking [1] based on the number of customers a node has. The list is dominated by high degree nodes; the two medium degree nodes in the list have also rather high APC values; in general all the nodes have relatively high APC values and eight of them are in the top 38 APC list. All the nodes in the list have poor QoB values, possibly due to relatively large stub ASes connecting to them. It is vivid that the centrality of the nodes in the CAIDA list is much larger than on our

Table 1. statistics of AS nodes with highest APC values

AS No.	degree	cent.	QoB	APC
20965	74	1190	0.79	26628
10910	205	385	0.59	16298
11686	187	3389	0.92	16042
701	2616	7956	0.72	14276
3215	115	422	0.80	13493
6517	175	474	0.83	12851
6849	186	472	0.56	12765
2200	51	347	0.50	12549
12859	79	1017	0.94	12396
23352	71	2113	0.94	12065

Table 2. statistics of AS nodes with highest CAIDA significance rankings

AS No.	degree	cent.	QoB	APC
3356	1784	7690	0.82	7559
209	1272	5381	0.72	6113
7018	2354	7992	0.74	11448
1239	2020	8022	0.74	10604
701	2616	7956	0.72	14276
3561	708	5762	0.79	2579
174	1483	7144	0.76	8797
703	216	1441	0.86	10539
19262	188	905	0.75	10763
702	680	5672	0.77	2101

APC list. While all the nodes identified as important in the CAIDA list have high APC values, the opposite analogy does not apply. Several of the nodes in our top 10 list are ranked below 200 in the CAIDA list.

6 Conclusion

We have shed light on the contribution of backup efficiency for the node significance classification problem. Given our theoretical analysis, we believe this contribution has merit in classification of network nodes in other fields outside the data networking domain.

We are aware that our results are not accurate for several reasons. First, as we stated in the main text, our AS relationship approximation is not accurate. Second, although we used the most detailed Internet map available through the DIMES project, the graph itself is still missing many links which can effect the calculation of all the measures, as well as the AS relationship deduction.

In the future we intend to broaden this research to study the effect of node failure on the Point of Presence (PoP) level as well as study relationship of sets of nodes in the AS graph in the context of backup and functionality. On the theoretical level, we intend to study the robustness of the APC and QoB measures to error in measurements, as well as further formal analysis of their properties.

References

1. CAIDA AS ranking. <http://as-rank.caida.org/>.
2. University of oregon RouteViews project. <http://www.antc.uoregon.edu/route-views/>.
3. Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, (406):378–382, 2000.
4. Sagy Bar, Mira Gonen, and Avishai Wool. An incremental super-linear preferential internet topology model. In *PAM '04*, Antibes Juan-les-Pins, France, April 2004.

5. Shai Carmi, Shlomo Havlin, Scott Kirkpatrick, Yuval Shavitt, and Eran Shir. Medusa: New model of Internet topology using k -shell decomposition. Technical report, arXiv, January 2006.
6. Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. 2001.
7. Danny Dolev, Sugih Jamin, Osnat Mokryn, and Yuval Shavitt. Internet resiliency to attacks and failures under BGP policy routing. *Computer Networks*, 50(16):3183–3196, November 2006.
8. Lixin Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, December 2001.
9. H. Jeong, S. P. Mason, A. L. Barabási, and Z. N. Oltvai. Lethality and centrality in protein networks. *Nature*, 411:41–42, 2001.
10. Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. In *SODA*, pages 668–677, 1998.
11. Vito Latora and Massimo Marchiori. Efficient behavior of small-world networks. *Physical Review Letters*, 87(19):198701, 2001.
12. Damien Magoni. Tearing down the internet. *IEEE Journal on Selected Areas in Communications*, 21(6):949–960, August 2003.
13. Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, D. Chklovski, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298:824–827, 2002.
14. Seung-Taek Park, Alexy Khrabrov, David M. Pennock, Steve Lawrence, C. Lee Giles, and Lyle H. Ungar. Static and dynamic analysis of the internet’s susceptibility to faults and attacks. In *IEEE INFOCOM 2003*, San-Francisco, CA, USA, April 2003.
15. Yuval Shavitt and Eran Shir. DIMES: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5), October 2005.
16. Yuval Shavitt and Yaron Singer. A linear time shortest paths algorithm for the Internet AS graph. *Tel Aviv University Technical Report*, (EE102), 2007.
17. Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *IEEE INFOCOM 2002*, New-York, NY, USA, April 2002.
18. L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos. A simple conceptual model for the Internet topology. In *Global Internet*, November 2001.