WICER: A Weighted Inter-Cluster Edge Ranking for Clustered Graphs

Divya Padmanabhan
Dept. of Computer Science
University of Minnesota
Minneapolis MN 55455
USA
dpadman@cs.umn.edu

Prasanna Desikan
Dept. of Computer Science
University of Minnesota
Minneapolis MN 55455
USA
desikan@cs.umn.edu

Jaideep Srivastava
Dept. of Computer Science
University of Minnesota
Minneapolis MN 55455
USA
srivasta@cs.umn.edu

Kashif Riaz
Dept. of Computer Science
University of Minnesota
Minneapolis MN 55455
USA
riaz@cs.umn.edu

Abstract

Several algorithms based on link analysis have been developed to measure the importance of nodes on a graph such as pages on the World Wide Web. PageRank and HITS are the most popular ranking algorithms to rank the nodes of any directed graph. But, both these algorithms assign equal importance to all the edges and nodes, ignoring the semantically rich information from nodes and edges. Therefore, in the case of a graph containing natural clusters, these algorithms do not differentiate between inter-cluster edges and intra-cluster edges. Based on this parameter, we propose a Weighted Inter-Cluster Edge Ranking for clustered graphs that weighs edges (based on whether it is an inter-cluster or an intra-cluster edge) and nodes (based on the number of clusters it connects). We introduce a parameter 'α' which can be adjusted depending on the bias desired in a clustered graph. Our experiments were two fold. We implemented our algorithm to relationship set representing legal entities and documents and the results indicate the significance of the weighted edge approach. We also generated biased and random walks to quantitatively study the performance.

1. Introduction

The World Wide Web has grown rapidly in size over the last decade and has proven to be a vast distributed source of semi structured information. Various Web Mining techniques have been developed and can be classified into Web Content Mining, Web Usage Mining and Web Structure Mining based on the data being analyzed. The success of Google [1] has invoked a lot of research focus on Web Structure Mining algorithms that typically involve link structure analysis [2] to improve ranking of the web search query results. PageRank Algorithm [1], [7] and Hypertext Induced Topic Selection (HITS) algorithm [6]are two of the most popular algorithms that have been developed for ranking pages on the web. A variety of modifications and improvements to these approaches have been developed in recent years [5][8]

In general, the PageRank Algorithm and HITS Algorithm can be applied to rank the nodes in any graph with directed edges. But, in scenarios wherein natural clusters exist in the graph, due to an added importance of nodes belonging to specific clusters, these algorithms in their naïve form do not capture semantic information of clusters to produce an efficient ranking of the nodes. This paper focuses on developing a measure called Weighted Inter-Cluster Edge Rank to rank the nodes of a clustered graph. The proposed algorithm includes a parameter α denoting the inter-cluster edge weight, to weigh edges between different clusters higher than edges within the same cluster. The second weighing factor, we introduced was to weigh the nodes based on the number of different kinds of edges that connect to it. The weights on the edges and nodes are used in determining the ranks of the nodes in the graph.

The rest of the paper is organized as follows. Section 2 describes reasons as to why PageRank and HITS algorithms are not sufficient to rank the nodes in a clustered graph. A brief overview of the PageRank algorithm has been presented in Section 3 prior to describing the proposed algorithm. The Weighted Inter-Cluster Edge Rank algorithm has been presented in Section 4. This is followed by the description of experiments performed and the observed results in Section 5. Section 6 summarizes the conclusions of the proposed algorithm.

2. Motivation

As mentioned earlier, HITS and PageRank are the two primary algorithms that have been developed for ranking the nodes in a graph that are based entirely on the topology of the graph. Kleinberg's HITS algorithm [6] is used to identify web pages that serve as "authorities" on a topic or as "hubs" that point to good "authority" pages. PageRank algorithm [1] [7] computes the rank of a node based on the ranks of the nodes that link to it. A more detailed explanation of the PageRank algorithm has been presented in Section 3.

Ranking techniques such as HITS and PageRank give uniform importance to all the nodes and edges of the graph. However, there exist domains that have richer semantic information describing the type of a node or a type of an edge. In such domains the nodes of the graph can be clustered and the edges categorized based on the semantic information, and this additional information provides a need for a new ranking scheme that distinguishes the type of nodes or edges. Existing algorithms do not take this information into consideration when computing the ranks. For example PageRank algorithm computes the rank of a node as a sum of ranks of backlinks independent of whether the inlink is from a node within the same cluster or from a different cluster. Neglecting such information might not truly represent the extent of importance of a node. Consider a network consisting of nodes representing airports and an edge representing a direct flight between the two airports. The nodes can be classified based on the country in which the airport is located. In order to rank the airports based on importance, it is necessary to rank the airports which have incoming international flights higher than the airports that have only domestic incoming flights. Also, the airport having incoming flights from a large number of countries should be ranked high. Existing algorithms do not include such edge and node properties in their computations.

In order to overcome these limitations for a clustered graph application, this paper suggests a Weighted Inter-Cluster Edge Rank (WICER) an extension of the PageRank, which takes into consideration both the node clustering and edge properties. WICER algorithm assigns specific weights to inter-cluster edges and intra-cluster edges and the rank of a document is computed as a weighed sum of ranks of the backlinks. The parameter inter-cluster edge weight (α) can be assigned an appropriate value based on the extent of bias required for inter-cluster edges in the graph for a specific application.

3. Page Rank

The PageRank algorithm [1] has been intuitively justified to model a random surfer in which a user clicks on links at random and the rank of a page signifies the probability of a user arriving at that page. A user can arrive at a page either by clicking on links or by randomly jumping to a page. The algorithm includes a parameter *d* which represents the probability of a user continuing to click on links and *(1-d)* as the probability that the user jumps to a random page.

PageRank of a page is determined using the random surfer model described above. The PageRank of a page can be computed as follows:

$$PR(A) = (1 - d)/N + d * \sum_{j \in S} PR(j)/C(j)$$

Where:
PR (j) is the PageRank of page j
S is the set of nodes that have an inlink to page A
C (j) is the out degree of page j
d is the dampening factor that is set to a value between 0 and 1. It is usually set to 0.85 for the web graph.
N is the number of nodes in the graph.

Addition of *d* also addresses the issue of a node being a *rank source*, i.e. a node having zero inlinks will be assigned at least a minimum PageRank value of d. In order to overcome the problem of a *rank sink*, i.e. a node having zero outlinks, it is assumed that there exists an outlink to all other nodes in the graph and there is an increased probability of starting at a random page.

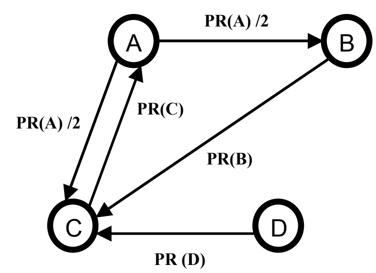

One iterative computation would be:

PR(A) = (1-0.85)/4 + 0.85 * (PR(C))
PR(B) = (1-0.85)/4 + 0.85 * (PR(A)/2)
PR(C) = (1-0.85)/4 + 0.85 * (PR(A)/2) + PR(B) + PR(D))
PR(D) = (1-0.85)/4

Fig 1: Example Computation of PageRank Algorithm

The PageRank of all the nodes in the graph are computed using an iterative algorithm. Each node is assigned an initial value and the PageRank of all the nodes are then calculated in several iterations based on the equations determined by the PageRank Algorithm. In the PageRank Algorithm, a page uniformly distributes its rank to each outlink and in turn when computing the PageRank of a page, the rank of each inlink is weighed equally. Therefore, the PageRank Algorithm is not sufficient to rank the nodes of a clustered graph in which the inter-cluster edges and intra-cluster edges are given varying importance.

4. Weighted Inter-Cluster Edge Rank Algorithm (WICER)

The Weighted Inter-Cluster Edge Rank (WICER) Algorithm can be modeled using a Biased Surfer Model. This model is based on the idea that if a user browsing for information in a specific domain comes across a link that points to a document outside the domain, will be more intrigued to click on the external link that leads to another domain rather than continuing within the same. Therefore, there is a higher chance that the user will arrive at a document that contains links from multiple domains than a document that has links from a single domain. An example for this would be in a scenario such as of an expert lawyer searching for law documents related to a specific "case" would be very interested in a document that is being referenced by multiple types of cases rather than a document being referenced by similar cases. Therefore it is logical to assign such a document a higher rank amongst the search results. Also it is most likely that the lawyer being an expert would already be familiar with documents of similar cases.

The WICER Algorithm has been developed for ranking the nodes of a clustered graph. The basic ideas of the algorithm are: (a) a node that has incoming inter-cluster edges should be ranked higher than a node that has incoming intra-cluster edges, (b) the rank of a particular node is weighted by the number of different clusters from which there exists an incoming edge to this node and (c) each cluster is weighted based on its importance.

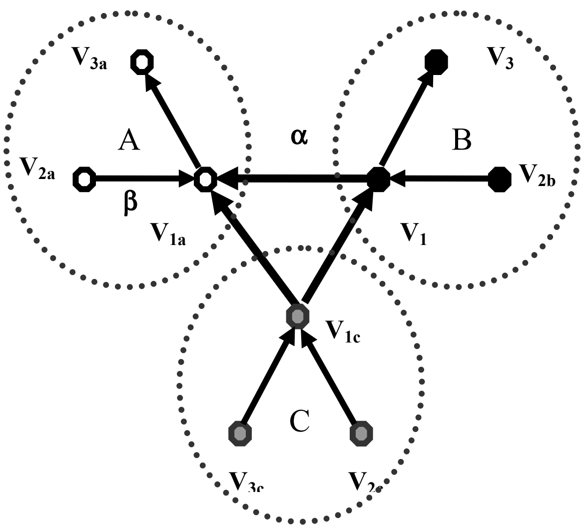

Fig 2: Illustration of WICER on Clustered Graph

The Weighted Inter-Cluster Edge Rank is given by,

$$R(V_{ic}) = \frac{(1-d)}{N} + d * \left(1 + \frac{C}{N_c}\right) * \left(\sum_{j=1..C} W(j) * \sum \frac{R(V_{kj})}{OD(V_{kj})} * w_{jc}\right)$$

Where,
R (V_{ic}) is the rank of vertex V_i of cluster c.
N is the number of nodes in the graph.
N_c is the number of clusters in the graph.
d is the damping factor to handle rank sinks.
W(j) is the weight of cluster j.
C is the set of clusters that have an edge to cluster c.
S_j is the set of vertices in cluster j having inlinks to vertex V_i of cluster c.
w_{jc} is the weight of the edge from cluster j to c.

$$w_{jc} = \alpha, \text{ if } j \neq c$$
$$w_{jc} = \beta, \text{ if } j = c$$

α is the inter-cluster edge weight
β is the intra-cluster edge weight

In the process of determining the rank of each node, when computing the sum of the ranks of the backlinks, the WICER algorithm weighs the rank of a backlink of an inter-cluster edge higher than the rank of a backlink of an intra-cluster edge. The factor of C/N_c represents the fraction of clusters from which there exists an edge to the present node being processed. This fraction in the equation is responsible for assigning a higher rank to a node that has incoming edges from a larger number of varying clusters. In Fig 2, the node V_{1a} has incoming edges from 3 clusters, V_{1b} has incoming edges from 2 clusters and V_{1c} has incoming edges from just 1 cluster. Since the node V_{1a} is responsible for connecting all three clusters, it is considered an important node by this algorithm. In order to account for the varying importance of different clusters, the clusters can be assigned specific weights and the ranks of the nodes in a particular cluster are multiplied by the corresponding weight.

The parameters α and β which are the inter-cluster and intra-cluster edge weights correspondingly, can be assigned values based on the application domain and their semantic significance in the network being considered. In most applications, α is higher than β indicating that the inter-cluster edges are more important than the intra-cluster edges. But, in some cases it might be semantically more meaningful to give higher weights to intra-cluster edges i.e. $\alpha < \beta$. In our experiments, we have only used the parameter α. The value of β is set to 1.0.

Another modification to the above equation is possible when the edges can be classified based on a specific semantic properties of edges independent of the node classification. This implies that for each node there could be multiple types of incoming and outgoing edges. For each node, the rank of the node can be weighed with a factor E/N_E , where E is the number of types of incoming edges and N_E is the total types of edges in the whole graph. A node that has more types of incoming edges will be weighed higher. Therefore, in the above equation the factor of C/N_C can be replaced with E/N_E, since semantically C/N_C represents the edge property that has been derived from node clustering.
The modified equation is given by:

$$R(V_{ic}) = \frac{(1-d)}{N} + d * \left(1 + \frac{E}{N_E}\right) * \left(\sum_{j=1..C} W(j) * \sum \frac{R(V_{kj})}{OD(V_{kj})} * w_{jc}\right)$$

5. Experiments and Results

The WICER algorithm has been implemented on a real data set to rank documents in a database that represents relationships between those documents.

The database contains several million rows of such relationships. The documents represent legal material, and the profile of the attorneys or judges as represented in the legal directories like West Legal Directory and Martindale Hubble. The document instances are stored in a typical distributed environment according to their legal category. There exists multiple relationship types between documents, i.e., each document can be involved in relationships of different types.

```
Function WICER( in G : Directed Graph with N nodes
                 Nc : Number of clusters, α : Inter-Cluster Edge Weight,  β : Intra-Cluster Edge Weight
                 W : Weights of Nc clusters)
return R[1..N] : Rank values of the nodes
prevR[1..N] : Temporary storage of rank values
for ( i = 1 to N)
   R[i] = 1/N;
RSink : Set of nodes those have zero OutDegree
while (diff > ε)
  for ( i = 1 to N)
          c : cluster to which Node i belongs
          Compute the rank of Node i
          for (j = 1 to C)
                  R(Vic)= R(Vic) + W(j)*∑kϵSj prevR(Vkj)*wjc/OutDegree(Vkj)
    R(Vic)=(1-d)/N + d *R(Vic);
          // Handle Dangling nodes with OutDegree = 0
          HandleRankSinks (RSink);
  end for
  diff = | R - prevR |
end while
```

Fig 3: Pseudo Code for WICER Algorithm

The database does not represent the exhaustive collection of legal documents. Currently it primarily consists of relationships where the profile of attorneys or judges is referred by other legal documents. The other relationship types that are present include litigation related material and documents related by precedence. In all, there were 2851826 documents with 11761584 relationships among them. The graph was constructed with documents as nodes and a relationship between two documents as a directed edge between the corresponding nodes in the graph. The nodes in the graph can be clustered based on the type the category of documents. Also, the edges are labeled depending on the relationship type they represent. This graph therefore has a node label and edge label associated with it. The key fields in each row of the database are: Document Id of the source of the relationship, Document Id of the target of the relationship and the relationship type between them. The category to which a document belongs is derived from additional information based on the relationship types it is involved in. Each node in the graph consists of a node Id, the cluster to which the node belongs, a vector holding the Ids of nodes that have an inlink to it, out degree of the node and the number of relationship types the node is involved in. The vector of inlinks also holds information as to whether an inlink is an inter-cluster edge or an intra-cluster edge. In the current implementation, the damping factor d is taken as 0.85, the inter-cluster (α) and intra-cluster (β) edge weights are assumed to be 1.2 and 1 respectively. The WICER ranks of the nodes are iteratively computed until convergence. The generic PageRank algorithm was implemented as well for the same data set with a damping factor of 0.85.

A limitation in the data set being used is that it consists of mostly profiler relationships and therefore it is only possible to analyze the impact of the ranking techniques particularly to those documents involved in these relationships.

Table 1 contains the top 10 documents ranked by WICER and their corresponding PageRank. Table 2 contains the top 10 documents ranked by PageRank and their corresponding ranks as computed by WICER. The document Ids have been modified for confidentiality purposes. The set of documents as ranked by PageRank and WICER were presented to domain experts to compare and analyze the results of the two methods. It was observed that most of the top documents in both ranking belong to the relationship type in which the target document is the profile of the attorney or the judge. This ranking of documents that refer to profiles is useful in an online service that needs to present all the legal documents in which attorney names have been mentioned e.g. law reviews, dockets, and cases etc.

Thus, the profile document that has a higher number of documents referring to it from different category databases should be ranked higher. The generic PageRank does not rank the documents according to this very important feature since there are situations where the algorithm has ranked the profile document (4th ranked) that has 57 profile references (inter-cluster edges) lower than a profile document (1st ranked) that has 42 profile references. This has been consistent across the results of generic PageRank algorithm.

Table 1: Documents Ranked by WICER

Document Id	WICER	PageRank
Ixxx468	1	63664
Ixxx44f	2	41597
Ixxx0db	3	63666
Ixxx469	4	70945
Ixxx452	5	63669
Ixxxffa	6	63674
Ixxx463	7	63675
Ixxx4da	8	63677
Ixxx43c	9	63680
Ixxx462	10	39630

Table 2: Documents Ranked by PageRank

Document Id	PageRank	WICER
Ixxx455	1	2792907
Ixxx9ee	2	2808245
Ixxx44d	3	2792908
Ixxx416	4	2792909
Ixxx43d	5	2792910
Ixxx4d1	6	2792912
Ixxx4d8	7	2792913
Ixxx60d	8	2792914
Ixxxf70	9	2792915
Ixxx53f	10	2804808

In contrast, the WICER algorithm consistently ranks document that have higher profile references above the document have lower profile references. The WICER algorithm performs better in this scenario because all profiler references are inter-cluster edges since legal documents belonging to different category databases are all referring to profile documents present in the same category database. The algorithm therefore succeeds in assigning a higher weight to an edge between a case law and a profile than an edge between two case law documents. Therefore, the profile document that has a higher number of profile references implying a higher number of inter-cluster edges is ranked higher by WICER.

In the second set of experiments we analyze and compare the performance of a biased walk with a random walk on a clustered graph. We measure the performance in terms of the coverage of clusters in each walk. Coverage is defined as the number of times a surfer enters a different cluster during the walk. The idea being, the more number of clusters (or topics) covered by the surfer, the wider and better is the information gained.

In a biased walk, at each node, in the process of choosing an outgoing link, inter-cluster edges, if present are given α times higher probability to be chosen over intra-cluster edges. Therefore, a value of 1.0 for α represents a random walk on the graph. The biased walk translates directly to our algorithm with the α set to an assigned biased value and β set to 1.0. Since in the process of determining the rank of each node, the ranks of inter-cluster backlinks are weighed α times higher than the ranks of intra-cluster backlinks implying that an inter-cluster edge is α times more important than an intra-cluster edge.

There are primarily four parameters that influence the coverage in a biased walk. These are: (a) the length of the walk, (b) the ratio of inter-cluster edges to the intra-cluster edges present in the graph referred to as Edge Ratio, (c) the ratio of the number of intra-cluster edges present in each cluster (i.e. ratio of cluster sizes) referred to as Cluster Ratio, and (d) the inter-cluster edge weight α. The edge ratio is represented as the ratio of the number of inter-cluster edges to the number of intra-cluster edges in cluster 1 and cluster ratio is represented as the ratio of number of edges in cluster 2 to the number of edges in cluster 1. In our experiment, we have maintained a constant length of 10000 and the size of cluster 1 as 1000 for all generated walks. For this given walk length, we have studied the influence of the other three parameters on the coverage. The constructed graph consists of only two clusters, but can be extended into a general case involving more number of clusters. However, to study the effect of these parameters on coverage, it is enough to consider the simplest case involving two clusters.

First, we studied the effect of the inter-cluster edge weight α on coverage. For this purpose, biased walks were generated with a fixed walk length and α values varying from 0.5 – 10.5 and the Coverage for each α has been plotted in Fig 4. It can be observed from the above graph that with all other parameters maintained a constant; the no. of inter-cluster edges traversed is directly proportional to α. This is true because with an increase in bias to inter-cluster edges, the walk traverses a higher number of inter-cluster edges.

Fig 5 represents the variation in Coverage with different Edge Ratios and Fig 6 represents the variation in Coverage with different Cluster Ratios with the other parameter, Edge Ratio maintained a constant equal to 0.5. It can be observed from Fig 7 that an increase in the edge ratio results in an increase in coverage. But, when the number of inter-cluster edges becomes significantly high, the chances of choosing an inter-cluster edge in both random and biased walks remain more or less the same.

Similar observations can be made in Fig 6 in which coverage decreases with increase in the cluster ratio until the clusters are of comparable sizes and then coverage stabilizes. On comparing the two graphs, it can be observed that for each value of α, the Coverage is significantly different depending on the values for the edge ratio and the cluster ratio.

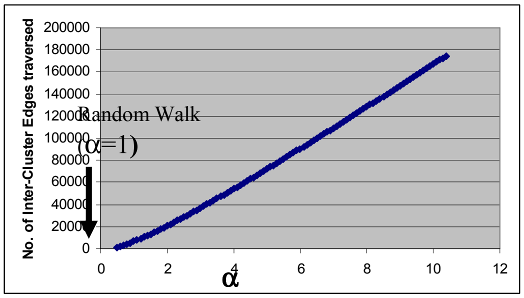

Fig 4: Effect of α on Coverage

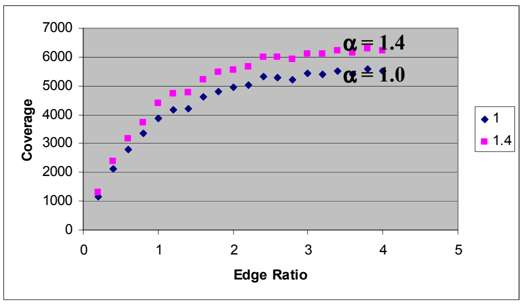

Fig 5: Variation in Coverage with Edge Ratio for different values of α

Fig 7 plots the coverage for both α = 1.4 (Biased walk) and α = 1 (Random walk) for varying values of the edge ratio and cluster ratio. The graph consists of two plots for each value of the edge ratio, the lower plot is for a random walk and the upper plot is for a biased walk. It can be observed that the biased walk consistently performs better that a random walk and the improvement in performance increases with increase in the edge ratio.

From the above experiment, it is seen that the performance of the biased walk experiment is better than random walk and also that the performance is dependant on the ratio of inter-cluster edges to intra-cluster edges present and the ratio of cluster sizes present in the graph. Since a biased walk is the underlying model of the Weighted Inter-Cluster Edge Rank Algorithm, the observations from the above experiment holds true for the performance of WICER algorithm as compared to PageRank.

6. Related Work

Various link based techniques have been developed for improving the results of a web search query. Two of the most important algorithms, PageRank and HITS have already been described before in Sections 2 and 3.

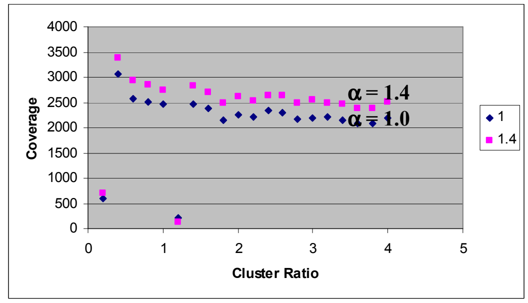

Fig 6: Variation in Coverage with Cluster Ratio for different values of α

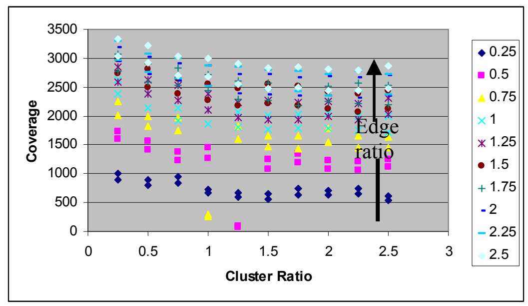

Fig 7: Variation in Coverage with Cluster Ratio for different values of Edge Ratio (0.25 – 2.5)

Many different approaches of the above algorithms have been proposed for the purpose of improving the computational efficiency, personalization or for specific applications. Topic Sensitive PageRank [4] by Haveliwala describes computing a set of PageRank vectors, each biased with a particular topic and therefore, each page has a set of scores, one for each topic. The ranking of the query results is done based on the scores of a page on the specific topics that the query belongs. To efficiently compute PageRank, the BlockRank algorithm proposed in [5] uses the nested block structure of the web to compute the local PageRank vector for each block and using these vectors to determine an approximate initial vector for computing the global PageRank.

Extensions of PageRank and HITS algorithms have also been proposed to address the limitation of these algorithms assigning equal weights to all edges in the graph. Xing and Ghorbani [8] present a Weighted PageRank algorithm in which instead of evenly distributing the rank of a page among its outlinks, each outlink page gets a value that is proportional to its number of inlinks and outlinks. The algorithm assigns a larger rank to a more popular page based on its outdegree and indegree. Incremental approaches to compute PageRank have also been developed [3]. However, the issue of ranking in a graph with labeled nodes and edges had not been dealt in depth. And this paper focuses on addressing measures to rank nodes that have labels and are naturally clustered with different types of edges between them.

7. Conclusions and Future Work

The commonly used hyper-link analysis algorithms such as PageRank and HITS can be applied to rank the nodes of any unlabeled directed graph. This paper presents a weighted edge ranking algorithm for a clustered graph, where nodes and/or edges have labels. The Weighted Inter-Cluster Edge Rank algorithm introduces a parameter α in order to assign higher weights to inter-cluster edges as compared to the intra-cluster edges while computing the rank for a node. It also weighs a node based on the number of different clusters from which there exists an incoming edge to this node. The proposed metric therefore succeeds in weighing the nodes that connect more clusters higher and ranks the documents belonging to the relationship database better as compared to the generic PageRank. Also, interesting observations have been made from the biased walk experiment about dependencies that exist between the number of inter-cluster edges between clusters, relative cluster sizes and the assigned inter-cluster edge weight.

The WICER algorithm can also be applied in domains such as social networks, a sub graph of the web etc., in which there exist natural clusters in the graph and it is semantically meaningful to rank the nodes of the graph.

At present, the WICER algorithm includes a weight on each cluster, but this has been assumed to be 1 for all the experiments. As future work, this parameter and its effectiveness need to be studied. The weight on intra-cluster edges β has been assumed to be 1. The influence of this parameter on the ranking and the interaction between α and β has to be analyzed. Though we have proposed a straightforward approach to rank documents, the richness of information from a graph that comprises of labeled nodes and edges is still to be fully exploited. Efficient determination of communities and other knowledge measures is a challenge that needs to be explored.

8. Acknowledgements

This work was supported by AHPCRC contract number DAAD19-01-2-0014. The content of the work does not necessarily reflect the position or policy of the government and no official endorsement should be inferred. Access to computing facilities was provided by the AHPCRC and the Minnesota Supercomputing Institute. We also thank the Data Analysis and Management Group at the Department of Computer Science for their valuable feedback.

9. References

[1] S. Brin and L. Page. "The anatomy of a large-scale hypertextual Web search engine", *Computer Networks and ISDN Systems*, 1998.

[2] P. Desikan, J. Srivastava, V. Kumar, P.-N. Tan, "Hyperlink Analysis – Techniques & Applications", *Army High Performance Computing Center Technical Report*, 2002.

[3] P. Desikan ,Nishith Pathak, J. Srivastava and V. Kumar "Incremental PageRank Computation on evolving graphs", Poster Paper at *Fourteenth International World Wide Web Conference* on May 10-14, 2005, in Chiba, Japan

[4] T.Haveliwala, "Topic-sensitive pagerank", in *Proceedings of the Eleventh International World Wide Web Conference, Honolulu, Hawaii,* 2002

[5] S.Kamvar, T.Haveliwala, C.Manning, and G.Golub, "Exploiting the block structure of the web for computing pagerank", Department of Stanford University, March 2003.

[6] J. M. Kleinberg. "Authoritative sources in a hyperlinked environment", *Journal of the ACM*, September 1999.

[7] L.Page, S.Brin, R.Motwani, and T.Winograd, "The PageRank Citation Ranking: Bringing Order to the Web", *Stanford Digital Technologies Project*, 1998.

[8] W.Xing and A.Ghorbani, "Weighted PageRank Algorithm". *Proc. of the 2rd Annual Conference on Communication Networks & Services Research (CNSR 2004),* Fredericton, Canada, pp. 305-314.