Generalized Random Surfer-Pair Models

Sai Kiran Narayanaswami
Robert Bosch Centre for Data Science and AI,
Indian Institute of Technology Madras
Chennai, India
saikirann94@gmail.com

Balaraman Ravindran
Dept. of Computer Science and Engineering
Robert Bosch Centre for Data Science and AI,
Indian Institute of Technology Madras
Chennai, India
ravi@cse.iitm.ac.in

Venkatesh Ramaiyan
Dept. of Electrical Engineering
Robert Bosch Centre for Data Science and AI,
Indian Institute of Technology Madras
Chennai, India
rvenkat@ee.iitm.ac.in

Abstract—SimRank is a widely studied link-based similarity measure that is known for its simple, yet powerful philosophy that two nodes are similar if they are referenced by similar nodes. While this philosophy has been the basis of several improvements, there is another useful, albeit less frequently discussed interpretation for SimRank known as the Random Surfer-Pair Model. In this work, we show that other well known measures related to SimRank can also be reinterpreted using Random Surfer-Pair Models, and establish a mathematically sound, general and unifying framework for several link-based similarity measures. This also serves to provide new insights into their functioning and allows for using these measures in a Monte Carlo framework, which provides several computational benefits. Next, we describe how the framework can be used as a versatile tool for developing measures according to given design requirements. As an illustration of this utility, we develop a new measure by combining the benefits of two existing measures under this framework, and empirically demonstrate that it results in a better performing measure.

I. Introduction

Advances in computing and information systems have enabled the collection of data on a large scale. A sizable portion is available in network form, with some well known examples being social networks and the World Wide Web. Naturally, the need arises to effectively utilize this data in the relevant domains. One area where network data has enormous potential is in recommender systems.

With the advent of electronic commerce, data is often available for products that customers have purchased, which can be used for collaborative filtering to provide recommendations for further purchases. As human knowledge grows, leading to rapidly expanding bodies of literature, there is an increasing requirement for effective recommender systems to aid researchers. Thus, utilizing bibliographic network data such as citation and co-authorship networks to provide recommendations is of rising importance.

In general, some form of similarity assessment would be a key component of recommender systems. We consider one class of similarity measures that work with only the link structure of the network known as *structural* (or *link-based*) similarity measures. Among the first such well known measures are Co-Citation [14] and its counterpart Bibliographic coupling [8] that measure respectively the frequency with which nodes refer to two given nodes , and the frequency with which nodes

are referenced together by two given nodes. Amsler [1] is a combination of the former two measures. SimRank [7] with its intuitive graph-theoretic model was a breakthrough which formed the basis of several subsequent measures. Among the notable ones are P-Rank [17], PSimRank [3] and SimRank* [16].

At the focus of this work is the probabilistic interpretation of SimRank known as the Random Surfer-Pair (RSP) Model [7]. We develop a generalization of this interpretation that applies to several measures and present a unified view of these measures. Casting them under the Generalized Random Surfer-Pair (GRSP) model also provides new insights into their functioning: with P-Rank, the GRSP model brings to light an otherwise indiscernible peculiarity, and for SimRank* provides an elegant and intuitive justification. We show how the framework also provides the tools necessary for performing Monte Carlo computation of these measures and discuss the computational benefits that result. Next, we outline the ways in which the GRSP model can be used to design measures, and the considerations involved in its usage for developing new measures. We then apply the framework to develop a hybrid measure, PSimRank* to combine the benefits of two existing measures, PSimRank and SimRank*. Experiments are performed to demonstrate empirically how this combination improves upon both of the measures on which it is based in terms of retrieval efficiency on a large (more than 2 million nodes) citation network from the Arnetminer dataset.

II. Background

In this section, we present SimRank and other relevant link-based similarity measures, and then the RSP interpretation for SimRank. Throughout, we consider networks with directed, simple graphs. We denote the graph by G, its vertex set by V, and its edge set by E. For any vertex a, $I(a)$ denotes its in-neighbors and $O(a)$ its out-neighbors. We index into these sets as $I_i(a)$ and $O_i(a)$.

A. SimRank and related measures

The famous SimRank philosophy that two nodes are similar if they are referenced by similar nodes is formalized recur-

sively as:

$$s\left(a,b\right) = \frac{C}{|I(a)||I(b)|} \sum_{i=1}^{|I(a)|} \sum_{j=1}^{|I(b)|} s\left(I_i(a), I_j(b)\right) \qquad (1)$$

where $a \neq b$ and $C < 1$ is a fixed, positive parameter. In the rare cases where either of the nodes a and b have no in-neighbors, the similarity is considered to be zero. Maximal self similarity applies, with $s\left(x,x\right) = 1 \; \forall x \in V$. These two base cases also apply for other measures discussed subsequently and nodes a and b on the left hand side are assumed to be distinct except as specified otherwise.

An out-link variant of SimRank, rvs-SimRank is also described in [5], which has the following form:

$$s\left(a,b\right) = \frac{C}{|O(a)||O(b)|} \sum_{i=1}^{|O(a)|} \sum_{j=1}^{|O(b)|} s\left(O_i(a), O_j(b)\right) \qquad (2)$$

Some deficiencies that have been identified in SimRank are as follows [5] :

- The In Links Consideration Problem : SimRank is unavailable (i.e set to zero) when either node has no in-neighbors, even though there may be evidence of similarity in the out-neighbors.
- The Pairwise Normalization Problem : This is the counter-intuitive effect that the SimRank score of a pair of nodes can *decrease* as there are more and more nodes referring to both of them. Consider two nodes u and v having several (say k) nodes that refer to both of them. Now, if these nodes are unrelated to each other (i.e they have zero similarity), the SimRank score between u and v is found to be $\frac{C}{k}$, which *decreases* with k. Thus, even though there are more witnesses to the similarity of u and v, their SimRank score decreases.
- The Level-wise computation problem : It can be shown that SimRank is unavailable for node pairs that don't have any paths of *equal* length to a common node. That is, it discards any evidence of similarity provided by path pairs of unequal length to a common node.

We now present some measures that were developed to address these issues.

1) P-Rank: P-Rank ([17]) was proposed to take into account out-links as well in computing similarity. It has the following recursive form :

$$s\left(a,b\right) = \lambda \times \frac{C}{|I(a)||I(b)|} \sum_{i=1}^{|I(a)|} \sum_{j=1}^{|I(b)|} s\left(I_i(a), I_j(b)\right)$$
$$+ (1-\lambda) \times \frac{C}{|O(a)||O(b)|} \sum_{i=1}^{|O(a)|} \sum_{j=1}^{|O(b)|} s\left(O_i(a), O_j(b)\right) \qquad (3)$$

It essentially adds an additional clause to the SimRank philosophy: *two entities are also similar if they reference similar entities*.

The base cases are similar to those of SimRank, except that only the term that corresponds to in (out) neighbors gets zeroed out if one or both of (a,b) doesn't have in (out) neighbors. This ensures that P-Rank is not unavailable for node pairs without in-links as is the case with SimRank, as long as they have out links.

2) PSimRank: The recursive form of PSimRank ([3]) is as follows :

$$s\left(a,b\right) = \frac{C|I\left(a\right) \cap I\left(b\right)|}{|I\left(a\right) \cup I\left(b\right)|} \cdot 1$$
$$+ \frac{C}{|I\left(a\right) \cup I\left(b\right)||I\left(b\right)|} \sum_{\substack{a' \in I(a) \setminus I(b) \\ b' \in I(b)}} s\left(a', b'\right) \qquad (4)$$
$$+ \frac{C}{|I\left(a\right) \cup I\left(b\right)||I\left(a\right)|} \sum_{\substack{b' \in I(b) \setminus I(a) \\ a' \in I(a)}} s\left(a', b'\right)$$

PSimRank solves the pairwise normalization problem by assigning greater importance to common in-neighbors. It is evident that unlike SimRank where each term $s\left(a,b\right)$ has the same weight, the weights have been redistributed so that terms of the form $s\left(x,x\right)$ (which are always 1, and make up the constant term) are given more weight than all other terms.

3) SimRank*: SimRank* ([16]) was proposed to solve the level-wise computation problem of SimRank and has the following recursive form:

$$s\left(a,b\right) = \frac{C}{2|I(a)|} \sum_{i=1}^{|I(a)|} s\left(I_i(a), b\right) + \frac{C}{2|I(b)|} \sum_{i=1}^{|I(b)|} s\left(a, I_i(b)\right) \qquad (5)$$

The measure is derived in [16] by actually enumerating all pairs of paths of (possibly) unequal length from a and b to a common node and computing the weighted sum of an exponentially decayed score associated with each path. Later on, we present a much simpler explanation as to how it works under the GRSP interpretation.

B. The RSP Model for SimRank

The Random Surfer-Pair interpretation for Simrank is based on a random experiment involving two random walks (or surfers) starting at the given nodes a and b, and traversing the graph *backwards* until they meet. That is, at the end of each step, each walk transitions to a randomly chosen in-neighbor. If either of the current nodes have no in-neighbors, the experiment is stopped.

Definition. Let $L(a,b)$ be the random variable that gives the number of steps taken until the surfers meet starting from a and b respectively. The *expected f-meeting distance* between a and b is defined for a given function f as $\mathbb{E}\left[f\left(L\left(a,b\right)\right)\right]$.

The f-meeting distance can be viewed as a score resulting from each instance of the experiment, and is itself a random variable. It turns out that for a specific choice of f, the expected score is nothing but the SimRank of (a,b). This equivalence of the Random Surfer-Pair formulation and the recursive form in equation 1 are proved in [7] :

Theorem. *SimRank as defined by Equation (1) is the same as the expected f-meeting distance between a and b for $f(t) = C^t$, where $0 < C < 1$.*

If the experiment is stopped because of unavailability of neighbors, $L(a, b)$ is considered to be infinite, thus making the f-meeting distance zero for that run of the experiment. Also, the base case of maximal self-similarity follows naturally because if $a = b$, the surfers deterministically meet at time $t = 0$, giving a score of 1 always.

In this interpretation, two nodes are similar if they are close to some source(s) of similarity. As we will discuss, various measures differ in what they consider to be sources of similarity and how much weight is assigned to them.

III. RELATED WORK

An RSP interpretation has already been made for PSim-Rank, and it was in fact the way in which the measure was developed [3]. PSimRank is defined the same way as the RSP model for SimRank, except that the transition probabilities for the surfers are modified so that they are more likely to meet at a common in-neighbor in the next step based on the number of such common in-neighbors at their current positions.

The improvements to Simrank are justified using this RSP model, and the recursive form in Eqn. 4 is also derived from it. However, unlike this work, the RSP model is not discussed as a general framework, and no general connections between the recursive forms and the RSP model are established.

We will arrive at the same RSP model later when we apply the GRSP model to PSimRank in Section VI-A.

IV. GENERALIZING THE RSP MODEL

For defining the GRSP model, we treat the Random Surfer-Pair experiment as a single random walk, but on a compound state space that consists of vertex pairs from $V \times V$ [7] to indicate the positions of both surfers, and also a "stopped" state, which represents unavailability. We use the letter h to denote a typical state from this space, which we denote by $\mathcal{S}$.

The transition probabilities for this random walk are denoted as $p(h' \,|\, h)$, the probability of transitioning to h' from h. The stopped state is an absorbing state, that is once the state is reached, it is impossible to leave. We can collect these probabilities into a matrix $\mathbf{P}$.

The idea is that different measures can be realized for different choices of transition probabilities, formalized by the following definition of the GRSP model :

Definition 1. For a particular matrix of transition probabilities $\mathbf{P}$, consider the following combined random walk experiment over the compound state space $\mathcal{S}$ starting from (a, b) at time $t = 0$:

- If the current state of the walk is h, the walk moves to the next state with probability $p(h' \,|\, h) \,\forall\, h' \in \mathcal{S}$ as specified by $\mathbf{P}$.
- The walk ends when a state of the form (x, x) is reached, for some $x \in V$.

Let $L(a, b)$ be the random variable that gives the number of steps taken until the walk ends. The *expected f-meeting distance* between a and b is defined for this combined walk for the function $f(t) = C^t$ as $\mathbb{E}\left[f\left(L\left(a, b\right)\right)\right]$ with $0 < C < 1$. This is a function of (a, b) which we call the *similarity measure induced by* $\mathbf{P}$ under the Generalized Random Surfer-Pair model.

One can think of other choices for f to encode different extents of decay of similarity, such as $f(t) = \frac{C}{t^2}$. However, the specific choice of $f(t) = C^t$ is what leads to the useful properties we discuss below.

The termination condition is equivalent to the random surfers meeting at some node for the first time. If the walk goes into the stopped state, it stays there forever, and does not reach a state of the form (x, x). This gives an infinite number of steps, thus leading to a score of zero. Again, the base case of maximal self-similarity applies here as well.

V. EQUIVALENCE TO RECURSIVE FORM

An interesting observation to be made is that the coefficients of any $s(a', b')$ in the recursive formulation of SimRank (Equation 1) are the same as the transition probabilities for the surfers in the Random Surfer-Pair Model going from (a, b) to (a', b'). This leads one to believe there could be a similar relationship for any transition probabilities $\mathbf{P}$. Indeed, this is true and the results are formally presented in the remainder of this section.

For a given transition probability matrix $\mathbf{P}$, consider the following set of recursive equations defined for all node pairs (a, b) :

$$s(a, b) = C \sum_{(a', b') \in \mathcal{R}((a, b))} p\left(\left(a', b'\right) \,\Big|\, (a, b)\right) s(a', b') \quad (6)$$

Here, $\mathcal{R}\left((a, b)\right) \subseteq V \times V$ is a region of support (which we will also refer to as *support set*) for (a, b) under $\mathbf{P}$, that is where the transition probability $p\left(\left(a', b'\right) \,\Big|\, (a, b)\right)$ is non-zero. Note that this *does not* include the stopped state, which means the sum of the coefficients appearing in the above equation need not be 1 (of course, they have to be less than 1).

The same base case of $s(a, a) = 1 \,\forall\, a \in V$ is used. If $\mathcal{R}\left((a, b)\right) = \{\phi\}$, $s(a, b)$ is taken to be zero unless $a = b$. These equations define what is called the *recursive similarity measure induced by* $\mathbf{P}$.

An iterative form is also defined:

$$s_{k+1}(a, b) = C \sum_{(a', b') \in \mathcal{R}((a, b))} p\left(\left(a', b'\right) \,\Big|\, (a, b)\right) s_k(a', b')$$
$$(7)$$

The following results establish the mathematical soundness of the GRSP model.

Theorem 1. *The following results hold true for the iterative form :*

- ***Monotonicity and boundedness :***

$$0 \leq s_k(a, b) \leq s_{k+1}(a, b) \leq 1 \qquad \forall\, (a, b) \in V \times V$$

- **Convergence to limit** : *The sequence $s_k(a,b)$ converges to a limit (obviously between 0 and 1 by the previous part) for all $(a,b) \in V \times V$.*

Proof. The monotonicity and boundedness are proved by induction. The inductive hypothesis is that

$$0 \leq s_{k-1}(a,b) \leq s_k(a,b) \leq 1 \qquad \forall\, (a,b) \in V \times V$$

The base case of this for $k = 0$ is trivial since $s_0(x,x) = 1$ and $s_0(a,b) = 0 \; \forall\, a \neq b$, and so is the case with $a = b$. The inductive step is as follows :

Monotonicity : We have

$$s_{k+1}(a,b) - s_k(a,b) = C \times$$
$$\sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) \left[s_k(a',b') - s_{k-1}(a',b')\right]$$

But $s_k(a',b') - s_{k-1}(a',b') \geq 0$ by the inductive hypothesis, and $p((a',b') \,|\, (a,b)) \geq 0$ since it is a probability, thus proving the monotonicity.

Boundedness : From the inductive hypothesis, $0 \leq s_k(a,b) \leq 1$. Therefore,

$$
\begin{aligned}
s_{k+1}(a,b) &= C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) s_k(a',b') \\
&\leq C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) \cdot 1 \\
&\leq C \leq 1
\end{aligned}
$$

Where we have used the fact that $\sum_{(a',b') \in \mathcal{R}((a,b))} p((a',b') \,|\, (a,b)) \leq 1$, since it is a sum of transition probabilities out of (a,b) (possibly less than one because of the stopped state). Similarly, it can be shown that $s_{k+1}(a,b) \geq 0$.

Convergence : Since $s_k(a,b)$ is bounded and non-decreasing, by the Completeness Axiom of Calculus, $s_k(a,b)$ converges to a limit $\forall\, (a,b) \in V \times V$, which we denote by $g(a,b)$. Of course, this limit must be between 0 and 1 as the sequence itself is bounded in that range. $\square$

From the above, the following result follows :

Theorem 2. *There exists a unique solution to the system of equations defined by equation 6.*

Proof. Let two solutions to Equation 6 be s_1 and s_2. Define their difference

$$\delta(a,b) = s_1(a,b) - s_2(a,b)$$

Let M be the maximum absolute value of δ, that is $\max_{(a,b)} |\delta(a,b)|$. Let this maximum value be achieved for (a,b), that is $M = |\delta(a,b)|$. If $a = b$, then clearly $M = 0$ as both s_1 and s_2 must satisfy the maximal self-similarity base case. Otherwise, we have

$$M = \left| C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) \left[s_1(a',b') - \right. \right.$$
$$\left. \left. s_2(a',b')\right] \right| \leq C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) \left|\delta(a',b')\right|$$
$$\leq C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \,\middle|\, (a,b)\right) M \leq CM$$

Here, we have used the fact that since (a,b) maximizes $|\delta(\cdot,\cdot)|$, $\left|\delta(a',b')\right| \leq M$ and again the fact that $\sum_{(a',b') \in \mathcal{R}((a,b))} p((a',b') \,|\, (a,b)) \leq 1$.

Now, since M is an absolute value, and $M \leq CM$ with $0 < C < 1$, we must have $M = 0$. This proves that s_1 and s_2 are always the same. Thus, there exists a unique solution to the recursive form of Equation 6, and that solution can be obtained as the limit of the iterative form. $\square$

Which leads to our central result :

Theorem 3. *The similarity measure induced by $\mathbf{P}$ according to definition 1 is the same as the recursive similarity measure induced by $\mathbf{P}$ (Equation 6).*

Proof. We first need to show that the expected f-meeting distances, for which we overload the notation $s(a,b)$ satisfy the recursive form. Let $\mathcal{W}_{a,b}$ be the set of all compound walks from (a,b) to a state of the form (x,x). Let $l(w)$ denote the length of such a walk w, and $p(w)$ the total probability, which is the product of the probabilities of the individual transitions. Then, by definition of the expected f-meeting distance,

$$s(a,b) = \sum_{w \in \mathcal{W}_{a,b}} p(w)\, C^{l(w)} \tag{8}$$

Now, consider the set of all such compound walks from one step ahead, that is, $\bigcup_{(a',b') \in \mathcal{R}((a,b))} \mathcal{W}_{a',b'}$. Note that all the individual sets $\mathcal{W}_{a',b'}$ are disjoint. Now, clearly this collection of walks differs from $\mathcal{W}_{a,b}$ only in the inclusion of the first transition to some (a',b'). Therefore, a bijection exists between this set and $\mathcal{W}_{a,b}$, and so $\mathcal{W}_{a,b}$ can be enumerated in terms of the new collection.

This means that for every member w of $\mathcal{W}_{a,b}$, there is some unique (a',b') and some unique $w' \in \mathcal{W}_{a',b'}$. Thus, it is possible to group the terms of the summation in Equation 8 by (a',b'). Now, the corresponding w' will have one step fewer, so $l(w') + 1 = l(w)$, and it omits the transition probability for the step from (a,b) to (a',b'), so

$$p(w) = p\left((a',b') \,\middle|\, (a,b)\right) p\left(w'\right)$$

Thus, Equation 8 is rewritten as:

$$s(a, b) =$$

$$\sum_{(a',b') \in \mathcal{R}((a,b))} \sum_{w' \in \mathcal{W}_{a',b'}} p\left((a',b') \big| (a,b)\right) p\left(w'\right) C^{l\left(w'\right)+1}$$

$$= C \sum_{(a',b') \in \mathcal{R}((a,b))} p\left((a',b') \big| (a,b)\right) \sum_{w' \in \mathcal{W}_{a',b'}} p\left(w'\right) C^{l\left(w'\right)}$$

But, by definition, $s(a', b') = \sum_{w' \in \mathcal{W}_{a',b'}} p\left(w'\right) C^{l\left(w'\right)}$.

This completes the proof that the expected f-meeting distances satisfy the recursive form of Equation 6. By the uniqueness result of Theorem 2, it follows that this is the same as the solution that can be arrived as a limit of the iterative form, thus establishing the equivalence of the Generalized Random Surfer-Pair model and its recursive form. □

This is the result necessary to convert an existing recursive form into a Random Surfer-Pair Model. All that needs to be done is to read off the non-zero coefficients into the appropriate places into the matrix, or equivalently get the support set and the corresponding transition probabilities as a function of (a, b). One thing to note here is that the probabilities corresponding to *actual* node pair destinations from (a, b) need not sum to 1, because it could go into the stopped state as well.

It would of course be more illustrative to get a concise description of the matrix. For SimRank for instance, $\mathcal{R}((a,b))$ is simply all neighbor pairs of a and b with the transition probabilities being uniform over this set. If either node has no in-neighbors, it transitions to the stopped state with probability 1 (i.e unavailable).

Our formulation allows for transitions from one compound state to any arbitrary state, and any number of destination states with non-zero transition probability. However, in existing measures, there are only a few possible transitions from any given state (a, b) compared to the total number of possible states, that too involving the neighbor pairs of a and b (on practical graphs which are not densely connected). This means that $\mathbf{P}$ is sparse, so even though there are $|V|^2$ states, only a few of them are involved in transitions from any given state, and it is no more complicated than the existing recursive formulations. However, we note that the above results continue to hold for any $\mathbf{P}$ regardless of sparsity.

VI. APPLICATION TO OTHER MEASURES

In this section, we demonstrate the usage of the GRSP model by applying it to PSimRank (Equation 4), P-Rank (Equation 3) and SimRank* (Equation 5) as discussed in the previous section. Throughout this exercise, any "unallocatable" probability due to e.g unavailability of in-neighbors or out-neighbors is given to the stopped state.

A. PSimRank

Support Set: $\mathcal{R}((a,b))$ is divided into 3 disjoint subsets: $\{(x,x),\, x \in I(a) \cap I(b)\}$, $(I(a) \setminus I(b)) \times I(b)$, $I(a) \times$
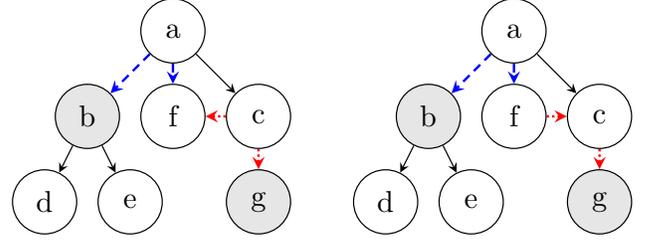

Fig. 1. Left: An instance of the GRSP experiment for P-Rank with the surfers starting at b (dashed path) and g (dotted path) and meeting at f. Right: In this graph, the same path-pair is invalid because the paths take opposite directions in the second step.

$(I(b) \setminus I(a))$. Note that parts of $I(a) \times I(b)$ are missing from this set unlike SimRank.

Transition Probabilities: The 3 subsets mentioned above are given total probabilities of $\frac{|I(a) \cap I(b)|}{|I(a) \cup I(b)|}$ (note that this is the the the Jaccard coefficients of in-neighbor sets), $\frac{|I(a) \setminus I(b)|}{|I(a) \cup I(b)|}$ and $\frac{|I(b) \setminus I(a)|}{|I(a) \cup I(b)|}$ respectively, and probabilities for individual elements are uniform over each set. This is exactly the RSP model presented in [3].

B. P-Rank

Support Set: $\mathcal{R}((a,b))$ now has two disjoint parts: $I(a) \times I(b)$ and $O(a) \times O(b)$.

Transition Probabilities: The two parts of $\mathcal{R}((a,b))$ are given total probabilities of λ and $1-\lambda$, and just like PSimRank, probabilities for individual elements are uniform over each set. These probabilities can be interpreted as follows: a coin with probability λ is tossed, and based on its result, *both* surfers move backward or forward and choose from applicable edges uniformly. This scheme is illustrated in Figure 1. Transitions where the surfers take different directions are not allowed. Thus, a theoretical deficiency exists that causes it to discard path-pairs containing such transitions as evidence for similarity. This is not at all evident from its recursive form.

C. SimRank*

Support Set: For SimRank*, $\mathcal{R}((a,b)) = (\{a\} \times I(b)) \cup (I(a) \times \{b\})$. Note that the Cartesian products involve singleton sets.

Transition Probabilities: The two parts of $\mathcal{R}((a,b))$ are given total probabilities of $\frac{1}{2}$ each, and again, individual probabilities are uniform over their respective subsets. These transitions are the same as tossing a fair coin, and based on the outcome, stepping *one* surfer to a uniformly chosen in-neighbor. An example of this is shown in Figure 2.

The notable feature here is that only one of the surfers is allowed to move at each step. The choice as to which surfer moves is made uniformly. From this, it becomes clear how SimRank* manages to consider path-pairs of unequal length. The surfers need not have made an equal number of jumps to meet at some node. This is a much more simpler and intuitive
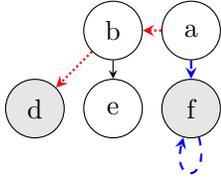
Fig. 2. An instance of the GRSP experiment for SimRank* with the surfers starting at d (dotted path) and f (dashed path) and meeting at a, with their paths having different effective lengths. The loop shows that surfer did not move that step.

explanation than the original analytic derivation in [16] by enumerating all such path pairs.

VII. MONTE CARLO COMPUTATION

A key benefit of the GRSP model is that it enables the use of Monte Carlo methods for the entire class of measures that fall under this model. This includes SimRank* and P-Rank, for which Monte Carlo methods have not been used so far. It is quite straightforward to apply: simulate the surfers' transitions starting from a given pair of nodes for some number of times, and return the average score as the similarity.

In practice, the surfers would have to be truncated after some number of steps L_{max}, and only a limited number of samples N_S can be drawn in the interest of fast querying, but decent guarantees for accuracy in practice are shown in [3]. Radius based pruning can be done to reduce the amount of nodes that need to be considered for top-k similarity queries, i.e restricting the search to nodes at a particular distance (or radius) from the query node.

The various benefits of using Monte Carlo computation are as follows:

1) *Complexity:* Typically, there is an easy ($O(1)$) way to generate a transition from any given state. In SimRank* for instance, all that needs to be done is to toss a coin, and advance one surfer to a randomly chosen in-neighbor. Thus, the complexity of a single similarity computation is just $O(N_S L_{max})$, where these quantities are much smaller than the size of the network. In comparison, solving the recursive equations takes at least $O(|V|^2)$ even when extensive optimizations can be made [16].

2) *Memory efficiency:* The memory requirement is also low since individual similarities can be computed without having to compute and store all pairs of similarities ($O(N^2)$ memory needed) as is necessary for the recursive forms. More precisely, it is $O(1)$ for each similarity being queried. This is particularly important when dealing with large networks (including citation networks). In fact, even storing the iterates while solving the recursive forms becomes infeasible for even medium sized networks (of the order of 10^6 nodes) by current standards.

3) *Adaptability to changes in network structure:* The network can be updated without having to recompute all similarities, which is useful in settings where nodes are dynamically added or removed. This is often the case, such as with social and citation networks, which change on a regular basis. Solving the recursive forms after each change would be highly impractical or even impossible.

4) *Parallelizability:* Possibly the biggest advantage is the extensive parallelizability; every instance of the RSP experiment can be run separately and concurrently. Further, under the RSP model, different similarities can be computed in parallel.

VIII. DESIGNING MEASURES UNDER THE GRSP MODEL

Theorem 3, along with our other results, characterizes a class of "sensible" measures that are non-negative, satisfy the base case of maximal self similarity being 1, and are bounded. It describes how such measures in the recursive form weight common sources of similarity in their computation in their GRSP model. We have also seen how existing measures can be better understood under the GRSP model, and the computational benefits that result. These observations highlight the utility of the GRSP model as a tool for developing measures according to given design requirements, rather than working with the recursive forms directly.

A. Modes of Design

The only mathematical requirement is a valid transition matrix $\mathbf{P}$, i.e $\mathbf{P}$ needs to be doubly stochastic. Thus, designing measures is a matter of allocating transition probabilities out of each compound state. This could be done from first principles such as those underlying the measures we have discussed so far, or existing measures could be involved, i.e $\mathbf{P}$ could be derived from some other transition matrices. For example, one might wish to combine existing measures in order to create a measure with the desirable properties present in each of the existing measures, an exercise we undertake in Section VIII-C.

One natural way to combine measures is to take a convex combination of the individual transition matrices, resulting in a transition matrix of the form

$$\mathbf{P} = \sum_{i=1}^{N} \lambda_i \mathbf{P}_i$$

with $\mathbf{P}_i$, $i = 1 \ldots N$ being transition matrices for some given measures, and λ_i are non-negative weights such that $\sum_i \lambda_i = 1$. In fact, it is easy to see that P-Rank itself is such a combination. The two measures involved are SimRank and Rvs-SimRank, as can be verified by using Theorem 3 to get the transition matrices from Equations 1, 2 and 3. Indeed, the rationale behind P-Rank is to account for both in-links as well as out-links, which are properties of these two measures that are easily combined as described above.

B. Domain Knowledge

The GRSP framework also enables the use of domain knowledge in the design process. We describe two possible ways in which this can be accomplished:

- Knowledge about the network structure: Different parts of the network could have different structures, and a mixture of behaviors of various measures could be necessary to accurately capture similarity. One possibility is to construct $\mathbf{P}$ based on independent transitions for each component of the node pair following different measures based on which part of the network the node belongs to, leading to probabilities of the form

$$p\left(\left(a^{'},b^{'}\right)\Big|(a,b)\right) = p_1\left(a^{'}\Big|a\right)p_2\left(b^{'}\Big|b\right)$$

where the RHS terms are obtained by appropriately marginalizing from two given transition matrices $\mathbf{P}_1$ and $\mathbf{P}_2$.

- Node and Edge attributes and other data: these can be used to give more weight to meaningful transitions that can be identified from the additional information. Similarly, this could help identify and prune out links that exist in the network but are not indicative of similarity in any way. This could prove particularly beneficial when utilizing features extracted from text data. For instance, [2] proposes a method to detect the intensity of references in scientific articles based on their textual content, that is how important a reference is to an article. This inferred attribute could be used to assign lower probabilities to tangential references when computing similarity as papers connected through such references can be very dissimilar.

C. PSimRank* : The Best of Both Worlds

In this section, we describe a measure designed based on two existing measures using the GRSP framework. Later in section IX, the efficacy of this measure is empirically evaluated.

Previously, we have described how PSimRank solves the Pairwise Normalization problem and SimRank* solves the Level-wise Computation problem, and what these entail in the Random Surfer-Pair domain. With PSimRank*, we attempt to combine these two benefits under the GRSP framework, resulting in a better measure because of solving both the problems.

The combination is straightforward; to make the surfers meet at a common in-neighbor with probability equal to the Jaccard coefficient just like in PSimRank, but the remainder of the time behave like SimRank*, moving only one at a time.

Support Set: $\mathcal{R}\left((a,b)\right) = (\{a\} \times I\left(b\right)) \cup (I\left(a\right) \times \{b\}) \cup \{(x,x), x \in I\left(a\right) \cap I\left(b\right)\}$. The first two subsets are for behavior like SimRank*, and the third subset is from PSimRank with both surfers stepping to a common in-neighbor.

Transition Probabilities: The third subset is allocated probability $\frac{|I(a) \cap I(b)|}{|I(a) \cup I(b)|}$, and the remaining is divided equally among the other two. The same scheme as before is used for individual probabilities.

The effect of this design is that node pairs which have a large number of common neighbors (i.e a high value for the Jaccard coefficient $\frac{|I(a) \cap I(b)|}{|I(a) \cup I(b)|}$) will greatly increase the tendency of the measure to adopt PSimRank-like behavior and jump to a common neighbor, thus solving the Pairwise

Measure	Arnetminer	Citeseer	Cora
SimRank	0.73	0.71	0.66
P-Rank(λ=0.4)	0.76	0.73	0.70
SimRank*	0.80	0.67	0.62
PSimRank	0.80	0.68	0.57
PSimRank*	0.81	0.69	0.63

Normalization problem for these nodes. For other node pairs with fewer common neighbors, this problem is not as severe, and the measure will focus more on the Level-wise Computation problem by adopting SimRank*-like behavior. This way, PSimRank* is expected to mitigate both problems overall.

IX. EXPERIMENTS

We compare the performance of PSimRank* against existing measures on a real-world dataset. For P-Rank, a sweep over the λ parameter is performed in increments of 0.1 from 0 to 1 and the best performing value of λ is used for comparison. SimRank and rvs-SimRank are the edge cases of this sweep with $\lambda = 1$ and $\lambda = 0$ respectively.

We use the Arnetminer dataset ([15]), which is a citation network of 2,244,021 papers and 4,354,534 citations extracted from DBLP[1]. A portion of these papers have been manually annotated and given labels corresponding to 10 different topics (clusters). The evaluation consists of running a top-k similarity query on some randomly chosen labeled nodes, and finding the Mean Average Precision (MAP) ([10]) for the answer set having the same label as the query. The rationale behind this is that papers in the same topic as the query are likely to be similar.

For all measures used here, pruning was done to radius 4 (in the undirected graph). The Random Surfer simulations were performed 200 times per query, and truncated after at most 15 steps. Top-100 queries were run on 50 randomly chosen labeled nodes that had at least 5 citations and 5 references to ensure that the measures wouldn't become unavailable. Since not all the nodes are labeled, only the nodes in the answer set that have a label are considered for calculating the MAP scores. Further, 50 such trials are performed and the averaged MAP scores are reported in Table I.

It is observed that PSimRank* outperforms all the other measures, improving on both PSimRank and SimRank* on which it is based. In networks much smaller than Arnetminer such as preprocessed versions [13] of the Cora [11] and CiteSeer [4] datasets which have only a few thousand nodes each, we found that PSimRank* as well as SimRank* and PSimRank performed worse than P-Rank (Table I). However, even in these cases, PSimRank* outperformed its predecessors. Thus, combining the benefits of PSimRank and SimRank*

[1] http://dblp.uni-trier.de/

under the GRSP interpretation is indeed effective, improving on both the measures.

X. Conclusions

The GRSP model serves as a unifying framework for a class of similarity measures based on the SimRank philosophy and subsumes several seemingly disparate measures. Any properties that are discovered for this framework would also apply to these measures. Admittedly, it is not all-encompassing; it is not evident how it can be applied to MatchSim [9] which uses weights obtained from a Maximum-Matching based scheme (which are not constant), and CoSimRank [12] which is based on Personalized PageRank [6].

Reinterpreting existing measures (P-Rank and SimRank*) under this framework has provided interesting insights about their functioning. The benefits of Monte Carlo computation are also brought to this class of measures. The development of PSimRank* under this framework, improving on the measures from which it was derived, highlights the potential of the GRSP Model to aid in the designing measures tailored to various applications and domains. One exciting avenue for future work is to use this framework to incorporate knowledge generated by Machine Learning methods, such as document embeddings generated by the state of the art Natural Language Processing methods. Hopefully, this work has opened up possibilities for theoretical dissection and development of more effective measures.

References

[1] Robert A Amsler. Applications of citation-based automatic classification, 1972.

[2] Tanmoy Chakraborty and Ramasuri Narayanam. All fingers are not equal: Intensity of references in scientific articles. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1348–1358, Austin, Texas, November 2016. Association for Computational Linguistics.

[3] Dániel Fogaras and Balázs Rácz. Scaling link-based similarity search. In *Proceedings of the 14th International Conference on World Wide Web*, WWW '05, pages 641–650, New York, NY, USA, 2005. ACM.

[4] C. Lee Giles, Kurt D. Bollacker, and Steve Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the Third ACM Conference on Digital Libraries*, DL '98, pages 89–98, New York, NY, USA, 1998. ACM.

[5] Masoud Reyhani Hamedani and Sang-Wook Kim. Simrank and its variants in academic literature data: Measures and evaluation. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, SAC '16, pages 1102–1107, New York, NY, USA, 2016. ACM.

[6] Taher H. Haveliwala. Topic-sensitive pagerank. In *Proceedings of the 11th International Conference on World Wide Web*, WWW '02, pages 517–526, New York, NY, USA, 2002. ACM.

[7] Glen Jeh and Jennifer Widom. Simrank: A measure of structural-context similarity. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, pages 538–543, New York, NY, USA, 2002. ACM.

[8] M. M. Kessler. Bibliographic coupling between scientific papers. *American Documentation*, 14(1):10–25, 1963.

[9] Zhenjiang Lin, Michael R. Lyu, and Irwin King. Matchsim: A novel similarity measure based on maximum neighborhood matching. *Knowl. Inf. Syst.*, 32(1):141–166, July 2012.

[10] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.

[11] Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3(2):127–163, Jul 2000.

[12] Sascha Rothe and Hinrich Schütze. Cosimrank: A flexible & efficient graph-theoretic similarity measure. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1392–1402, Baltimore, Maryland, June 2014. Association for Computational Linguistics.

[13] Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.

[14] Henry Small. Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science*, 24(4):265–269, 1973.

[15] Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. Arnetminer: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 990–998, New York, NY, USA, 2008. ACM.

[16] Weiren Yu, Xuemin Lin, Wenjie Zhang, Lijun Chang, and Jian Pei. More is simpler: Effectively and efficiently assessing node-pair similarities based on hyperlinks. *Proc. VLDB Endow.*, 7(1):13–24, September 2013.

[17] Peixiang Zhao, Jiawei Han, and Yizhou Sun. P-rank: A comprehensive structural similarity measure over information networks. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pages 553–562, New York, NY, USA, 2009. ACM.