



Resource location based on precomputed partial random walks in dynamic networks



Víctor M. López Millán^{a,*}, Vicent Cholvi^{b,1}, Antonio Fernández Anta^{d,2}, Luis López^{c,3}

^a Universidad CEU San Pablo, Spain

^b Universitat Jaume I, Spain

^c Universidad Rey Juan Carlos, Spain

^d Institute IMDEA Networks, Spain

ARTICLE INFO

Article history:

Received 23 November 2015

Revised 29 February 2016

Accepted 9 April 2016

Available online 13 April 2016

Keywords:

Resource location

Dynamic networks

Random walks

Complex networks

ABSTRACT

The problem of finding a resource residing in a network node (the *resource location problem*) is a challenge in complex networks due to aspects as network size, unknown network topology, and network dynamics. The problem is especially difficult if no requirements on the resource placement strategy or the network structure are to be imposed, assuming of course that keeping centralized resource information is not feasible or appropriate. Under these conditions, random algorithms are useful to search the network. A possible strategy for static networks, proposed in previous work, uses short random walks precomputed at each network node as partial walks to construct longer random walks with associated resource information. In this work, we adapt the previous mechanisms to dynamic networks, where resource instances may appear in, and disappear from, network nodes, and the nodes themselves may leave and join the network, resembling realistic scenarios. We analyze the resulting resource location mechanisms, providing expressions that accurately predict average search lengths, which are validated using simulation experiments. Reduction of average search lengths compared to simple random walk searches are found to be very large, even in the face of high network volatility. We also study the cost of the mechanisms, focusing on the overhead implied by the periodic recomputation of partial walks to refresh the information on resources, concluding that the proposed mechanisms behave efficiently and robustly in dynamic networks.

© 2016 Published by Elsevier B.V.

1. Introduction

Random walks are network routing mechanisms which have been extensively studied and used in a wide range of applications: physics, mathematics, population dynamics, bioinformatics, etc. [11,18,24]. Roughly speaking, they choose, at each point of the

route, the next node uniformly at random among the neighbors of the current node.

Among the advantages of random walks when applied to communication networks is the fact they need only local information, avoiding the bandwidth overhead necessary in other routing mechanisms to communicate with other nodes. This is especially useful when there is no knowledge on the structure of the whole network, or when the network structure changes frequently. For these reasons, random walks have been proposed as a base mechanism for multiple network applications, including network sampling [9,16], network resource location [1,10,28,33], network construction [5,15,19–21], and network characterization [8,29,32].

The emergence of the peer-to-peer (P2P) architecture model has been proven useful in many applications in recent years. While structured P2P systems (e.g., Chord [31], CAN [27] –Content-Addressable Network–, Kademlia [23], etc.) provide efficient search mechanisms, they introduce a significant management overhead. In turn, unstructured systems have little management overhead and, consequently, have been considered in several scenarios (e.g., Gnutella [5,20], CAP [14] –Cluster-based Architecture for P2P–,

* Corresponding author. Tel.: +34 91 3726424 (ext. 4911); fax: +34 91 3724049.

E-mail address: vmlopez.eps@ceu.es (V.M.L. Millán).

¹ Supported in part by the Spanish Ministerio de Economía y Competitividad under grant TIN2011-28347-C02-01.

² Supported in part by Ministerio de Economía y Competitividad grant TEC2014-55713-R, Regional Government of Madrid (CM) grant Cloud4BigData (S2013/ICE-2894, co-funded by FSE & FEDER), NSF of China grant 61520106005, and European Commission H2020 grants ReCred and NOTRE.

³ Supported in part by the European Commission under projects NUBOMEDIA (FP7-ICT GA 610576), FI-WARE (FP7-ICT GA 285248) and FI-CORE (FP7-ICT GA 632893). This work is also supported by the Spanish Ministerio de Economía y Competitividad under grant REACTIVEMEDIA (TIN2013-41819-R) and by Comunidad de Madrid under grant CLOUD4BIGDATA (P2013/ICE-2894) cofounded by FSE & FEDER.

etc.). For such systems, searching techniques based on *flooding*, *supernodes* and *random walks* have been used. However, it is known that flooding mechanisms do not scale well [13], and supernode systems are vulnerable to supernodes failures (technical problems, attacks, censorship, etc.). Therefore, random walks have been used to search for resources held in the nodes of a network (e.g., [5,22]), a problem usually known as *resource location*. The problem consists of finding a node that holds a given resource, the *target node*, starting at some *source node*. The source node is checked for the resource: if it is not found locally, the search hops to a random neighbor, checking that node for the resource. The search proceeds through the network in this way, until the target node is reached.

Nevertheless, by using random walks, some nodes may be (unnecessarily) visited more than once, while other nodes may remain unvisited for a long time. Avoiding this problem is the main objective of our study.

1.1. The Dynamic resource location problem

In this work, we are concerned with the resource location problem in networks with dynamic behavior regarding both resources and nodes.

In particular, we consider scenarios in which resources are randomly placed in the nodes across the network. Then, on the one hand, we consider scenarios in which the instances of the resources may appear and disappear from a time instant to another, maybe at different nodes. On the other hand, we also consider scenarios in which the network nodes themselves may also leave and join the network.

In these scenarios, all the nodes of the network may launch independent searches for different resources (e.g., files) at any time, without the help of a centralized server, and we are interested in measuring the average performance of searches between *any* pair of nodes.

Our assumptions regarding dynamicity cover a wide range of scenarios. For instance, in P2P networks nodes represent users, which may leave and join the network quite often. Also, resources represent the shared files, which may appear and disappear from time to time.

1.2. Contributions

In this work, we use the technique of concatenating *partial random walks* (PWs) to generalize the resource location mechanisms introduced in [17] for static networks to the case of dynamic networks. In particular, this paper provides new analytical models that predict the behavior of the resource location mechanisms in scenarios with dynamic resources and in scenarios with dynamic nodes, along with new simulation experiments to validate the analytical results. In addition, a new analysis of the cost of the mechanisms in these scenarios is provided. We consider the two versions of the mechanisms proposed in [17] and adapt them to operate in the dynamic scenarios. In the first version, which we refer as *choose-first PW-RW*, the search mechanism first chooses one of the PWs at random and then checks its associated information for the desired resource. In the second version, which we refer as *check-first PW-RW*, the search mechanism first checks the associated resource information of all the PWs of the node, and then randomly chooses among the PWs with a positive result. It is clear that there are other choices regarding the search mechanisms that seem reasonable. However, the ones considered in our study follow very different approaches (one chooses first, and the other checks first). Therefore, that will allow us to check the strength of our approach in very different circumstances.

Then, we have studied their performance, considering the following aspects:

- *Dynamic resources*: We have developed an analytical mean-field model for both mechanisms when resources are dynamic. Expressions are given for the corresponding *expected search length* (i.e., the expected number of hops taken to find the resource, averaged over all source nodes, target nodes, and network topologies) of each mechanism. These expressions provide predictions as a function of several parameters of the model, such as the network structure (size and degree distribution), the resource dynamics, and those of the mechanisms operation. The predictions of the models are validated by simulation experiments in three types of randomly built networks: regular, Erdős-Rényi, and scale-free. These experiments are also used to compare the performance of both mechanisms, and to investigate the influence of the resource dynamics. We have compared the performance of the proposed search mechanisms with respect to random walk searches. For the *choose-first PW-RW* mechanism we have found a reduction in the average search length with respect to simple random walk ranging from around 57% to 88%. For the *check-first PW-RW* mechanism such a reduction is even bigger, achieving reductions above 90%.
- *Dynamic nodes*: We have also considered the case where network nodes may leave and join the network, and have provided both analytical and experimental results. We have found a reduction in the average search length with respect to simple random walks above 94% (using the *check-first PW-RW* mechanism).
- *Cost*: Finally, we have analyzed the cost of the PW-RW mechanisms, defined as the number of messages, taking into account the cost of searches themselves and the cost of precomputing the PWs in each recomputation interval. We have provided analytical expressions for the relation between the cost and the length of the recomputation interval, as well as for the interval length that minimizes this cost. We have found that the impact of the precomputation of PWs on the cost is not significant in a wide range of lengths of the precomputation interval depending on the dynamic behavior of the network and on the dynamic behavior of searches.

1.3. Related work

Da Fontoura Costa and Travieso [6] study the network coverage of three types of random walks: traditional, preferential to untracked edges, and preferential to unvisited nodes. Also, Yang [33] studies the search performance of five random walk variations: no-back (NB), no-triangle-loop (NTL), no-quadrangle-loop (NQL), self-avoiding (SA) and high-degree-preferential self-avoiding (PSA). Self-avoiding walks (SAW) are those that try not to visit nodes that have already been visited. Several variations of this idea have been studied, differing in the probability of revisiting a node. Some examples are: strict SAW, true or myopic SAW, and weakly SAW [4,30]. In [7], Das Sarma et al. propose a distributed algorithm to obtain a random walk of a specified length ℓ in a number of rounds¹ proportional to $\sqrt{\ell}$.

López Millán et al [17] propose a mechanism for resource location based on building random walks connecting together partial walks (PW) previously computed at each network node. However, the mechanisms in [17] are only valid when both nodes and resources have a static behavior, contrary to the approach we follow in this paper.

The remainder of this paper is arranged as follows. Section 2 and 3 respectively present the *choose-first PW-RW* and *check-first PW-RW* mechanisms in scenarios with dynamic

¹ A *round* is a unit of discrete time in which every node is allowed to send a message to one of its neighbors. According to this definition, a simple random walk of length ℓ would then take ℓ rounds to be computed.

resources. Section 4 adapts the previous mechanisms to scenarios with dynamic nodes. In Section 6, the cost of the mechanisms is evaluated in dynamic scenarios. Finally, Section 7 concludes this paper and provides some future work lines.

2. Choose-first PW-RW with dynamic resources

Consider a communication network (e.g., the Internet, a wireless ad-hoc network, etc.) that provides full connectivity to the end system entities (e.g., computers, smartphones, etc.) connected to it. Next, consider a subset of N end system entities which establish logical neighboring relations for some purpose (a P2P file sharing system, a social network application, etc.). These end system entities (as the *nodes*) and their neighboring relations (as the *links*) form an *overlay network* on top of the *underlying network*.² The resource location mechanisms described in this paper apply to such an overlay network (referred to simply as the *network*), and we will focus on searches for resources in the end system entities (referred to as the *nodes*).

Each of the N nodes holds a set of resources. We focus on a given resource of interest, of which initially there is a number of instances randomly placed in many distinct network nodes. Our resource location problem is defined as finding one of the nodes that hold the resource (the one we encounter first, called the *target node*), starting by a certain node (the *source node*). We make no assumption on the underlying communication network, focusing on the number of nodes of the overlay visited to find a resource, as a measure of search performance.

In our analysis, for each search, we assume that the source node is uniformly chosen at random among all nodes in the network. Likewise, we consider that the instances of the resource have been randomly distributed throughout the network. The probability that a given node holds an instance of the resource is denoted by p_{res} . The expected number of instances of the resource for a given network is denoted by $\bar{R} = N \cdot p_{res}$.

Resources have a dynamic behavior. If we compared two snapshots of the network, one taken at time $t=0$ and another taken at time $t=T$, we would observe that some of the instances have disappeared while other new ones have appeared (at different nodes, in general). More concretely, an instance present in a node disappears with probability d . Conversely, an instance not initially present in a node appears in that node with probability a . We will use d as an input parameter to characterize resource dynamics. For a value of d , we will set a so that expectation of the number of resources (\bar{R}) remains unchanged. This way, our results will isolate the impact of resource dynamics on the search mechanism due only to the deterioration of information, discarding the effect of a possible increase or decrease on the expected number of resource instances. Fig. 1 provides an illustrative example of the dynamic behavior of the resources.

The search mechanism. A search performs a walk from the source node to the target node according to the mechanism that is defined below. The search mechanism proposed in this paper, referred to as PW-RW, exploits the idea of efficiently building *total random walks* from *partial random walks* available at each node of the network. It comprises two stages:

1. *Partial random walks construction:* In an initial stage at time $t=0$, every node i in the network precomputes a set W_i of w random walks before the searches take place, with the initial




Fig. 1. Example of the dynamic behavior of resources. The network is shown at $t=0$ (left hand side), when the PWs are precomputed, and at time $t=T$ (right hand side), the end of the PW recomputation interval. The resource instances have changed during the interval.

distribution of resource instances in the network. Each of these partial walks (PW) has length s , starting at i and finishing at a node reached after s hops. Using the PW-RW mechanisms, the PWs computed in this stage are simple random walks (i.e., the next node to be visited is chosen uniformly at random among the neighbors of the current node).

During the computation of each PW in W_i , node i registers the resources held by the s first nodes in the PW (from i to the one before the last node). The last node of the PW is excluded, being included in the PWs departing from it. In particular, for each PW computed by i , this node keeps the set of the identifiers of the resources held by the nodes in that PW. In this set, there is no indication of the particular node or nodes holding the resource. The registered information will be used by the searches in stage 2, to decide whether to traverse that PW (if the resource looked for is found in the set of identifiers), or to jump over it (if the resource is not found).

2. *The searches:* During the interval $0 < t \leq T$, after the PWs are constructed, searches are performed in the network. We will consider the system at $t=T$, in which, as stated above, the dynamic behavior of resource instances is characterized by d . Therefore, results obtained will reflect the performance of the search mechanism in a worst case scenario, since searches executed in $t < T$ will see a probability that an instance disappears less than or equal to d . There is no relation between the interval T , measured in time units, and the hops of a search, other than the assumption that T is much longer than the duration of a typical search, as discussed later in this section in paragraph *Resource dynamics*.

A search can be qualitatively described as a sequence of jumps over PWs, interleaved with some occasional unsuccessful PW traversals, and finished by the successful traversal of a PW until the target node is visited. Unsuccessful PW traversals are caused by outdated resource information associated to that PW, i.e., the resource was in the PW at $t=0$ but it has disappeared at the time of the search. The last PW traversal will be incomplete in general, in the sense that its length will be less than or equal to s , since the search stops when the resource is found.

We measure the length of searches in *hops*. Some of these hops are *jumps* (over PWs), and other are *steps* (traversing PWs). We distinguish between *unnecessary steps* (in unsuccessful PW traversals), and *final steps* (in the last, successful, PW). The definition of the search mechanism and the associated concepts are illustrated by the example in Fig. 2, in which PWs of length $s=6$ are used.

At this point, we emphasize the difference between the *search* just defined and the *total walk* that supports it. Indeed, the total walk consist of the concatenation of *partial walks* as defined above. Searches are therefore shorter in length than their

² Note that neighbors in the overlay are not in general neighbors in the underlying network. Note also that the underlying communication network provides connectivity between any pair of end system entities, even if they are not neighbors in the overlay.




Fig. 2. An example search: total walk, partial walks, jumps and steps.

corresponding total walks because of the number of steps saved in jumps over PWs in which we know that the resource is not located, although these saving may be reduced by the unnecessary steps due to outdated information within PWs.

More formally, we describe how searches are performed as follows. Let a search start at a node A . A PW in W_A is chosen uniformly at random. Its associated resource information collected in stage 1 is then queried (any instance of) the desired resource.

- If the query result is negative (i.e., the resource is not in the set of identifiers associated with that PW), the search *jumps* to node B , the last node of that PW. Note that the current node and the node to which the search *jumps* are not neighbors in the overlay network in general. Jumps therefore make use of the underlying communication network. The process is then repeated at B and the search keeps jumping in this way while the results of the queries are negative.
- If the query result is positive (i.e., the resource is in the set of identifiers), the search *traverses* that PW looking for the resource. It starts checking if the current node has the desired resource. If it does not, the search takes a *step* to the next node of the PW, checking again if it has the resource. The search proceeds through the PW in this way until the resource is found or the PW is finished.
 - If the resource is found, the search stops.
 - Otherwise (i.e., the search reaches the last node of the PW without having found the resource in the previous nodes), it means that the information collected in stage 1 for that PW and the resource of interest is no longer valid. The search considers that the result is negative, and the search process is repeated at the last node of the PW.

Resource dynamics. Regarding resource dynamics, we realize that searches are executed based on information collected at $t=0$ that may be outdated during the interval $0 < t \leq T$, when the queries are performed. The duration of the interval (T) is assumed to be much longer than the duration of a typical search, since we are interested in using the PWs computed at $t=0$ for as many searches as possible, with an acceptable degradation of information associated with those PWs.³ Therefore, most of the searches take place within the interval, and only a few start in one interval and finish in the following interval. The part of those searches in the following interval benefits from the newly computed PWs of this interval. Our analysis will consider that all searches take place within the interval, thus reflecting a worst case situation. Four cases arise when the information associated with a PW is queried for the resource, as is shown in Table 1 and illustrated by Fig. 3.

Table 1
Resource dynamics: query results.

case	Resource present in PW		Query result
	$t=0$	$t=T$	
a)	no	no	True Negative (TN)
b)	yes	yes	True Positive (TP)
c)	no	yes	False Negative (FN)
d)	yes	no	False Positive (FP)

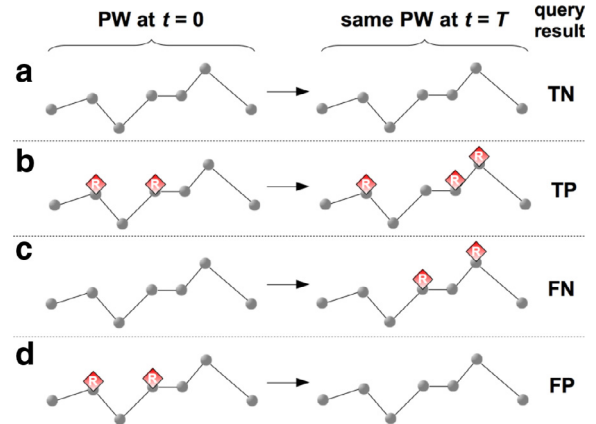


Fig. 3. Resource dynamics: examples of results when a PW precomputed at $t=0$ is queried at $t=T$.

A *True Negative* or TN (case a) occurs when no instances were present in the PW at $t=0$, and the same holds at $t=T$. A *True Positive* or TP (case b) occurs when one or more instances are present in the PW at $t=0$ and one or more instances (not necessarily the same ones) are present at $t=T$. The impact of resource dynamics on the performance of the search mechanism comes from the *False Negatives* (FN) and the *False Positives* (FP). A FN (case c) occurs when there were no instances in the PW at $t=0$, but at least one instance is present at $t=T$. It makes the search jump over that PW, ignoring the new instance(s). An FP (case d) occurs when there were one or more instances in that PW at $t=0$ but all of them are gone at $t=T$ and there are no new instances at $t=T$. It makes the search traverse a whole PW fruitlessly, since no instances are currently in that PW. Note that the case when all instances disappear from the PW, but some other instance(s) appears in that PW is included in the TP case.

At this point, we note that the performance of the search mechanism can be affected by the degradation of the information collected in the PWs. For instance, if all the instances of a given resource disappear and new instances appear in nodes that happen to be in PWs which did not have the resource at $t=0$, the resource would not be found, since searches would jump over those PWs due to the (false) negative result of the queries. Therefore, in order to keep their accuracy, PWs may need to be “refreshed”. Once the interval is finished, PWs are recomputed at $t=T$ to acquire fresh resource information, and searches from $t=T$ to $t=2T$ will use the new PWs. In other words, stages 1 and 2 described above are repeated with a period T . In Section 6, we assess the impact of this PW recomputation on the overall searches cost.

The search metrics. In this work, we are interested in the number of hops to find a resource, which is defined as the *search length*. This search length is a random variable that takes different values when independent searches are performed. The *search length distribution* is defined as the probability distribution of the search length random variable. So, the *expected search length*, derived from the mentioned distribution, taken over all networks, all

³ A discussion of the optimal duration of the interval is provided in Section 6.

source nodes, and target nodes, is an interesting performance measure of the searching mechanism in a given network.

In the next section we provide an analytical model that gives an estimation of the expected search length of the *choose-first PW-RW* mechanism. We use a *mean-field* model and analysis. In this approach, instead of computing the search length distribution for each network, source node and target node, and then taking expectations, the analysis itself handles average values. Whereas this approach may not accurately reflect some border cases (like some specific network topologies), it captures rather accurately the general characteristics and metrics of the mechanisms. This fact has been also confirmed by using simulations.

2.1. Analysis

Since we have defined the expected search length for *any* pair of source and target nodes, the expected length of the search from the source node is the same as the expected search length from any subsequent node in the search. This allows us to write a recursive expression for the expected search length. Consider the search is at some node i , one of whose PWs is queried for the resource. The expected search length from i can be written in terms of the expected search length from the next node j that will be queried for the resource, *plus* the number of hops taken from node i to node j . If the result of the query at i was negative (either false or true), the search will jump from i to j , adding 1 hop to the search. If the result was a false positive, the search will traverse the entire PW without finding the resource, adding s hops to the search. Finally, if the result was a true positive, the search will traverse the PW until the resource is found, adding a number of hops whose expected value we call \bar{F} . These cases occur with probabilities that we call, respectively: P_n , P_{tp} , and P_{fp} . Denoting the expected search length using PWs of length s by \bar{L}_s , the resulting recursive expression is then:

$$\bar{L}_s = (\bar{L}_s + 1) \cdot P_n + (\bar{L}_s + s) \cdot P_{fp} + \bar{F} \cdot P_{tp}, \quad (1)$$

where P_n , P_{tp} , and P_{fp} are the probabilities of choosing a partial walk (out of the w PWs of the node) for which the query for the resource returns a negative (either a TN or a FN, see Table 1), a TP, and an FP result, respectively, with $P_n + P_{tp} + P_{fp} = 1$. \bar{F} is the expectation of the number of final steps taken when traversing the last PW, until an instance of the resource is found (TP). Solving for \bar{L}_s , we obtain:

$$\bar{L}_s = \frac{1}{P_p} \cdot (P_n + s \cdot P_{fp}) + \bar{F}. \quad (2)$$

The probabilities in Eq. 2 are estimated with the following expressions:

$$\begin{aligned} P_{tp} &= \sum_{i=1}^w \sum_{j=0}^{w-i} P(i, j) \cdot \frac{i}{w}, \\ P_{fp} &= \sum_{i=0}^{w-1} \sum_{j=1}^{w-i} P(i, j) \cdot \frac{j}{w}, \\ P_n &= \sum_{i=0}^{w-1} \sum_{j=0}^{w-i-1} P(i, j) \cdot \frac{w - (i + j)}{w} = 1 - P_{tp} - P_{fp}, \end{aligned} \quad (3)$$

where $P(i, j)$ is the probability that, in the w PWs of a node, there are i PWs whose queries return a TP result and j PWs that return an FP result:

$$P(i, j) = B(w, p_{tp}, i) \cdot B(w - i, p_{fp}, j), \quad (4)$$

where

- $B(m, q, n)$ is the coefficient of the binomial distribution

$$B(m, q, n) = \binom{m}{n} \cdot q^n \cdot (1 - q)^{(m-n)}, \quad (5)$$

- p_{tp} is the probability that a given PW at any node returns a TP result,⁴
- and p_{fp} is the probability that a given PW at any node returns an FP result, conditioned on the fact that it does not return a TP.

Therefore, in order to evaluate the estimation of the expected search length given by Eq. 2, we need to obtain the values of p_{tp} , p_{fp} and \bar{F} . Let us provide them:

1. The variable p_{tp} has been defined as the probability that a given PW at any node returns a TP result. This probability can be easily estimated if we condition it on the fact that the PW (of length s) had exactly r instances of the resource at $t=0$. Defining $P_{pw}(r)$ as the probability that a PW has r instances of the resource, and recalling that \bar{R} is the expectation of the number of instances of the resource in the network, we can write:

$$p_{tp} = \sum_{r=1}^{\min\{s, \bar{R}\}} P_{pw}(r) \cdot [(1 - d)^r + d^r \cdot (1 - (1 - a)^{s-r})], \quad (6)$$

where the brackets contain the probability that not all the r instances present at $t=0$ have disappeared (with probability d) at $t=T$ or, if they did disappear, at least one instance appeared (with probability a) in some of the $s - r$ remaining nodes in that interval.

An estimation for $P_{pw}(r)$ can be obtained using the random properties of a random walk in networks built randomly. In particular, we consider that the next hop of a random RW can take it to any of the endpoints in the network (except the endpoints of the current node since we do not allow self-loops). Then we estimate $P_{pw}(r)$ as $B(s, p_{rw}, r)$, where p_{rw} is the probability that the RW visits a node with an instance of the resource in the next hop. In turn, we estimate this probability as:

$$p_{rw} = \frac{\bar{R} \cdot \bar{k}}{S - \bar{k}_{rw}} \cdot \frac{\bar{k}_{rw} - 1}{\bar{k}_{rw}}. \quad (7)$$

The first fraction in Eq. 7 is the ratio of positive endpoints (the ones connected to the \bar{R} nodes that have an instance of the resource) and all endpoints in the network ($S = \sum_k k n_k$) except those of the current node. We use the average degree of the network ($\bar{k} = \sum_k k n_k / N$) as an estimation of the degree of a node that holds the resource (which is assigned or not with uniform probability p_{res} across the network). Similarly, we use the expectation of the degree of a node visited by a random walk as an estimation of the degree of the current node:

$$\bar{k}_{rw} = \sum_k k \cdot \frac{k \cdot n_k}{S} = \frac{1}{S} \cdot \sum_k k^2 \cdot n_k. \quad (8)$$

The second fraction in Eq. 7 corrects the previous ratio taking into account that, when at a node of a given degree, the probability of not going backwards (and therefore having the chance to find the resource) is the probability of selecting any of its endpoints but the one that connects it with the node just visited by the walk. With this, the estimation of $P_{pw}(r)$ is:

$$P_{pw}(r) = \binom{s}{r} \cdot (p_{rw})^r \cdot (1 - p_{rw})^{s-r}. \quad (9)$$

⁴ This probability p_{tp} is not to be confused with P_{tp} (note the different case), defined above as the probability of *choosing* a PW which returns a TP out of the w PWs of the current node.

2. We have defined p_{fp} as the probability that a given PW at any node returns an FP result, conditioned on the fact that it does not return a TP. This conditioning comes from the second binomial coefficient in this equation, which we restrict to the $w - i$ PWs which we know that do not return a TP, since the ones that do are accounted for in the first binomial coefficient. In other words, the second binomial coefficient includes the PWs that return a TN, a FN or an FP result, and p_{fp} is the probability that it returns an FP result conditioned on that. We can then easily write an estimation of p_{fp} as:

$$p_{fp} = \frac{1}{1 - p_{tp}} (1 - P_{pw}(0) - p_{tp}), \quad (10)$$

where we are subtracting $P_{pw}(0)$ (the probability of cases TN and FN in Table 1), and p_{tp} (the probability of case TP).

3. An expression for \bar{F} , the expectation of the number of final steps taken when traversing the last PW until an instance of the resource is found, is still needed to be finally able to estimate the average search length in Eq. 2. For this we rely on $\bar{F}(r)$, the expectation of that variable conditioned on there being r instances of the resource in the PW. Then:

$$\bar{F} = \frac{1}{1 - P_{pw}(0)} \cdot \sum_{r=1}^{\min\{s, \bar{R}\}} \bar{F}(r) \cdot P_{pw}(r). \quad (11)$$

Note that \bar{F} is in fact conditioned on there being at least one instance of the resource in the PW, since it corresponds to a TP (see Eq. 1). This is the reason of the fraction multiplying the summation in Eq. 11.

Now we provide an expression for $\bar{F}(r)$ as the expectation of the position of the first resource in the PW (conditioned on there being r instances of the resource in the PW):

$$\bar{F}(r) = \sum_{i=0}^{s-r} \left[i \cdot \left(\prod_{j=0}^{i-1} \left(1 - \frac{r}{s-j} \right) \right) \cdot \left(\frac{r}{s-i} \right) \right]. \quad (12)$$

Each factor in the product of Eq. 12 is the probability that there is no instance in the j th position of the PW, conditioned on that there is no instance in the previous position. The final factor outside the product is the probability that there is an instance in the i th position conditioned on there are no instances in the previous positions.

2.2. Performance evaluation

In this section, we apply the model presented in the previous section to real networks, and we also validate its predictions with data obtained from simulations. Three types of networks have been chosen for the experiments: regular networks (constant node degree), Erdős-Rényi (ER) networks and scale-free networks (with power law on the node degree). These topologies cover a wide variety of real networks [3], ranging from communication networks [2] to Internet [26] and P2P networks [12].

A network of each type and size $N = 10^4$ has been randomly built with the method proposed by Newman et al. [25] for networks with arbitrary degree distribution, setting their average node degree to $\bar{k} = 10$. Each network is constructed in three steps: (1) a preliminary network is constructed according to its type; (2) its degree distribution is extracted, and (3) the final network is obtained feeding the referenced method with that degree distribution. For each experiment, 10^6 searches have been performed. In every search, the source node has been chosen uniformly at random, and every node in the network has been assigned an instance of the resource looked for with probability $p_{res} = 10^{-2}$ at $t=0$. Therefore, the expected number of resource instances present in the network at $t=0$ for all searches is $\bar{R} = N \cdot p_{res} = 100$.




Fig. 4. Choose-first PW-RW with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in a scale-free network.

Table 2

Choose-first PW-RW with dynamic resources: reduction of the expected search lengths relative to RW searches for several d .

d	Search length reduction (%)
0.0	88.28
0.1	86.37
0.3	81.23
0.5	72.94
0.7	57.26

We show the results for scale-free networks, which are especially interesting since this type of degree distribution is frequently found in real networks. Results for regular and ER networks, which do not differ significantly from scale-free results, are included in the Appendix for the interested reader.

2.2.1. Expected search length vs. PW length

Fig. 4 shows the expected search length in a scale-free network for several values of d (i.e., the resource dynamics). The number of PWs per node is set to $w = 5$, although the performance of the PW-RW mechanism is independent from this parameter,⁵ since only one PW is used to locate the resource. Model predictions (Section 2.1) are plotted with lines and simulation results are shown as points. It can be seen that the model provides an accurate correspondence with the real data, with larger error for higher values of d and s . These deviations are discussed further in Section 2.2.3.

All curves show a minimum point, which marks the optimal PW length (s_{opt}) and the corresponding optimal expected search length (\bar{L}_{opt}). Interestingly, the values of s_{opt} are small and do not depend heavily on d . According to the analytic data, s_{opt} ranges between 13 and 15 for all shown curves.

2.2.2. Reduction of the expected search length

We have compared the performance of the proposed search mechanism for \bar{L}_{opt} with searches based on simple random walks (RW searches), finding large reductions in the average search lengths. Table 2 shows the relative reductions (%) for several

⁵ The extreme case of having just one PW is to be avoided because it yields many unfinished searches, since it is relatively easy to build walks that are loops that do not cover the network. Indeed, if the last node of a PW is a node whose (only) PW has been already used in that total walk, it will take the search to the same place again, resulting in a never-ending loop.

values of d . We can see that the reductions that *choose-first PW-RW* achieves with respect to RW searches are lower for higher d , ranging from around 88% in the case when $d = 0$ to 57% when $d = 0.7$.

2.2.3. Deviations of the predictions of the analytical model

It can be concluded from Fig. 4 that the proposed analytical model succeeds in capturing the general behavior of the search mechanism observed in the simulations. There are, however, significant deviations that deserve further evaluation and discussion. We look first at the region of the graphic corresponding to interesting practical scenarios, i.e., values of s close to s_{opt} . For $s = 20$, relative deviations range between 2.9% ($d = 0.0$) and 7.9% ($d = 0.7$). We now look at the region for large s , to explore the dependency of the deviation on the length of the PWs. For $s = 300$, relative deviations are slightly larger, ranging between 3.6% and 8.2%.

These deviations may be explained by one aspect of the construction of PWs that is not captured by our analytical model: revisits within a PW. That is, when constructing a PW, the next node visited by the random walk can be one already in that PW. This will in general happen for several nodes of the PW, with increasing probability for longer PWs (larger s). Now consider a search traversing a PW with revisited nodes. In the calculation of the probability that a PW contains the resource, our analysis considers each node of the PW independent from the others, which is clearly not true for a revisited node (which appears more than once in the PW but is in fact the same node). This makes our model slightly *optimistic* because it overestimates the probability of a PW containing the resource, and therefore underestimates the probability of a query to that PW resulting in a false positive when resources disappear. This means that there will actually be more PW traversals caused by false positives than expected, increasing the average search length in simulations. This effect is larger for higher d when more resources disappear and therefore more than expected PWs cause false positive queries, which is in accordance with Fig. 4.

3. Check-first PW-RW with dynamic resources

This section describes a variation of the mechanism presented in Section 2. Suppose the search is currently in a node and it needs to pick one of the PWs in that node to decide whether to traverse it or to jump over it. Recall from stage 2 under paragraph “The search mechanism” that the original mechanism first *chooses* one of the PWs at random, and then *checks* its associated information for the desired resource, resembling the behavior of a random walk. The proposed variation, on the other hand, reverses the order of these tasks. It first *checks* the associated resource information of *all* the PWs of the node, and then randomly *chooses* among the PWs with a positive result, if any (otherwise, it chooses among all PWs of the node, as the original version). This *check-first PW-RW* mechanism improves the performance of the original (or *choose-first PW-RW*) since the probability of choosing a PW with the resource increases, with no extra storage space cost.

There is another, less important, difference between the algorithms. In the original version, the nodes whose resources were registered in the information associated to the PW ranged between the current node and the one before the last node. In the *check-first* version, the resource information is registered from the *first* node (the next to the current node) to the *last* node in the PW. This change slightly improves the performance of the new version, since the probability of choosing a PW with the resource increases also in the cases where the resource is held by the last node of the PW. The rest of the operation of the mechanism remains the same.




Fig. 5. *Check-first PW-RW* ($w = 5$) with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in a scale-free network.

3.1. Analysis

Most of the analysis provided in Section 2.1 is still valid for *check-first PW-RW*. We present here the equations that need to be modified to reflect the new behavior. That is the case of Eq. 3 for the probabilities of choosing a PW with a TP, FP and negative result, respectively. Their counterparts follow. Remember that i and j represent the number of PWs of the node that return a TP result and an FP result, respectively:

$$\begin{aligned}
 P_{tp} &= \sum_{i=1}^w \sum_{j=0}^{w-i} P(i, j) \cdot \frac{i}{i+j}, \\
 P_{fp} &= \sum_{i=0}^{w-1} \sum_{j=1}^{w-i} P(i, j) \cdot \frac{j}{i+j}, \\
 P_n &= P(0, 0) = 1 - P_{tp} - P_{fp}.
 \end{aligned} \tag{13}$$

The expression for the expectation of the number of final steps taken when traversing the last PW until the resource is found (Eq. 11) is still valid. It uses $\bar{F}(r)$, the expectation of the position of the first resource in the PW, conditioned on there being r instances of the resource in the PW. Its expression (Eq. 12) needs to be modified, since the range of nodes whose resources are associated with the PW has changed from $[0, s-1]$ to $[1, s]$. The indexes limits and their use in the expression have been updated as necessary in the new expression, which completes the analysis of the *check-first PW-RW* mechanism:

$$\bar{F}(r) = \sum_{i=1}^{s-r+1} \left[i \cdot \left(\prod_{j=1}^{i-1} \left(1 - \frac{r}{s-j+1} \right) \right) \cdot \left(\frac{r}{s-i+1} \right) \right]. \tag{14}$$

3.2. Performance evaluation

For the performance evaluation of the *check-first PW-RW*, we use the same scenarios used in the performance evaluation of the *choose-first PW-RW*. As earlier, the corresponding figures for regular and ER networks can be found in the Appendix.

3.2.1. Expected search length vs. PW length

Fig. 5 shows the expected search length in a scale-free network for several values of d , and for $w = 5$. The shape of the curves is the same as that for the original mechanism (Fig. 4), and the discussion on the observed deviations of the analytical results given in Section 2.2.3 is also applicable to this case.

A substantial decrease in the optimal search length (\bar{L}_{opt}) for *check-first PW-RW* is observed. For example, for $d = 0.3$, \bar{L}_{opt} is




Fig. 6. PW-RW with dynamic resources: expected search length \bar{L}_s vs. PW length s for *choose-first* PW-RW and *check-first* PW-RW ($w = 2, 5, 10$) in a scale-free network with $d = 0.3$.

Table 3

Check-first PW-RW ($w = 5$) with dynamic resources: reduction of the expected search lengths relative to RW searches for several d .

d	Search length reduction (%)
0.0	94.29 %
0.1	93.44 %
0.3	91.04 %
0.5	86.91 %
0.7	78.68 %

around 11, while it was about 21 for the *choose-first* PW-RW mechanism. The optimal PW length also diminishes, from about 14 to 6 in that case. The expected search length decrease is due to the fact that the new mechanism checks all the PWs in the node for the resource and then chooses one only among those with positive result, increasing the probability of choosing a PW that currently holds the resource. Following this reasoning, the more PWs in the node, the higher this probability. It is therefore interesting to explore the dependency of \bar{L} and s_{opt} with w .

Fig. 6 shows the expected search length of the *check-first* PW-RW mechanism, for $w = 2, 5$ and 10 with $d = 0.3$. To make the comparison of performance between both mechanisms easier, a curve corresponding to *choose-first* PW-RW has been added to the graph. As explained above, the performance of the latter is independent from w , but it plays a central role in the *check-first* PW-RW. The range of the axes of this graph has been restricted to focus on the area around s_{opt} . For higher s , the curves for the several w converge. As expected, it is observed that higher w yields lower \bar{L}_{opt} , with a value about 8 for $w = 10$. Another interesting observation is that s_{opt} also diminishes for higher w , falling to 4 in this case.

These values mean a reduction of about 92% in the expected search length of simple random walks, with 10 precomputed PWs of just four nodes. Higher reductions can be achieved, at the expense of increasing the cost of the computation of the PWs, as analyzed in Section 6.

3.2.2. Reduction of the expected search length

Table 3 is provided as a reference, presenting the reductions achieved by *check-first* PW-RW with respect to random walk searches for $w = 5$ and several d . We see that reductions range between 94% and 78%, while those of *choose-first* PW-RW ranged between 88% and 57%. That is an additional reduction of about 50%.

4. Choose-first and check-first PW-RW with dynamic nodes

So far we have analyzed the performance of PW-RW when resource instances dynamically appear in, and disappear from, network nodes. In this section, we explore the case where network nodes themselves leave and join the network. In our model, we assume that at time $t = T$, some nodes have left the network and others have joined it (recall that PWs were computed initially at time $t = 0$). Nodes that have left the network have a double impact on PW-RW: (1) their resource instances are no longer available, so the information on the resources in a PW is degraded, and (2) searches will not be able to visit them when following their PWs. On the other hand, new nodes have no impact on PW-RW: since they are not in the previously computed PWs, they will not be used by searches, even if they hold instances of the resource looked for.⁶ Therefore, a single parameter l , the probability that a node present at $t = 0$ has left at $t = T$, is enough to evaluate the impact on PW-RW performance at $t = T$. At this time, the expected number of resource instances is $Np_{res}(1 - l)$, in a network with expected size $N(1 - l)$. This implies that the expected number of resource instances per node is still p_{res} , which allows a fair comparison with the baseline (static) case. The algorithm for the PW-RW mechanisms needs to be adapted for the dynamic nodes scenario. We call nodes that remain in the network at $t = T$ *active nodes*. When traversing a PW, the walk will visit the first active node at each hop. Similarly, when jumping over a PW, the walk will hop to the last active node in the PW. If no active nodes remain in the PW, the search will proceed visiting any active neighbor or will stop if there are none (only searches that find the resource are included in the performance evaluation results shown later).

4.1. Analysis

The analytical models in Section 2 and 3 are modified for the dynamic nodes scenario as follows. The expected length of a PW is now $s(1 - l)$. Therefore, Eq. 2 becomes:

$$\bar{L}_s = \frac{1}{P_{tp}} \cdot (P_n + s(1 - l) \cdot P_{fp}) + \bar{F}. \quad (15)$$

Eq. 3 for *choose-first* PW-RW are still valid, and so are their counterpart Eq. 13 for *check-first* PW-RW. We need to provide a new version of p_{tp} in Eq. 4, since the reason of information degradation is now the dynamics of nodes:

$$p_{tp} = \sum_{r=1}^{\min\{s, \bar{R}\}} P_{pw}(r) \cdot (1 - l^r), \quad (16)$$

with $P_{pw}(r)$ given by the original Eq. 9, which relied on Eqs. 7 and 8. Eq. 10 for p_{fp} can still be used, with the new definition of p_{tp} above. Finally, a new expression for \bar{F} is needed, since it depends on the number of active nodes u in the last PW, which is now a random variable:

$$\bar{F} = \sum_{u=1}^s \bar{F}(u) \cdot P_{pw}(u), \quad (17)$$

where $\bar{F}(u)$ is the expected number of final steps conditioned on there being u active nodes in the PW, and $P_{pw}(u)$ is the probability that the PW has u active nodes. To obtain the former we rely on $\bar{F}(u, r)$, the expected number of final steps conditioned on there being r instances of the resource in a PW with u active nodes:

$$\bar{F}(u) = \frac{1}{1 - P_{pw}(u, 0)} \cdot \sum_{r=1}^{\min\{u, \bar{R}\}} \bar{F}(u, r) \cdot P_{pw}(u, r), \quad (18)$$

⁶ New nodes with fresh resource instances will become useful from the searches point of view when PWs are recomputed at $t = T$.




Fig. 7. Choose-first PW-RW with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several l in a scale-free network.

where $P_{pw}(u, r)$ is the probability of a PW with u active nodes having r instances of the resource. Now, an estimation for $\bar{F}(u, r)$ can be obtained reasoning as for Eq. 12:

$$\bar{F}(u, r) = \sum_{i=0}^{u-r} \left[i \cdot \left(\prod_{j=0}^{i-1} \left(1 - \frac{r}{u-j} \right) \right) \cdot \left(\frac{r}{u-i} \right) \right]. \quad (19)$$

The expression for $P_{pw}(u)$ can be obtained as in Eq. 9:

$$P_{pw}(u) = \binom{s}{u} \cdot (p_{rw,a})^u \cdot (1 - p_{rw,a})^{s-u}, \quad (20)$$

where $p_{rw,a}$, the probability that a node visited by a RW is active, is:

$$p_{rw,a} = \frac{N(1-l) \cdot \bar{k}}{S - \bar{k}_{rw}}. \quad (21)$$

Likewise, the expression for $P_{pw}(u, r)$ is:

$$P_{pw}(u, r) = \binom{u}{r} \cdot (p_{rw})^r \cdot (1 - p_{rw})^{u-r}, \quad (22)$$

where p_{rw} is given by Eq. 7.

4.2. Performance evaluation

We look now at the performance of the PW-RW mechanisms in networks with dynamic nodes, that operate according to the model described in the previous section. As earlier, the corresponding figures for regular and ER networks can be found in the Appendix.

4.2.1. Expected search length vs. PW length

Fig. 7 shows results for choose-first PW-RW in a scale-free network for several values of l , the probability of a node leaving the network. Predictions of the analytical model in Section 4, plotted as lines, provide accurate estimations for experimental results, shown as points. The observed deviations are discussed in Section 4.2.3.

The effect of node dynamics is in general very similar to that of resource dynamics (see Fig. 4). We note, however, that s_{opt} slightly increases with l in Fig. 7, while it slightly decreased for larger d (the probability of a resource instance disappearing) in Fig. 4. In addition, we see that \bar{L}_{opt} remains lower in the dynamic nodes case as l grows, compared to the dynamic resources case for the same values of d . Both effects have their origin in the fact that PWs are effectively shorter in the dynamic nodes case, since some of their nodes are no longer in the network when searches are performed. This contributes to shorten the searches, obtaining the optimal lengths with originally longer PWs.

Table 4
Choose-first PW-RW with dynamic nodes: reduction of the expected search lengths relative to RW searches for several l .

d	Search length reduction (%)
0.0	88.28 %
0.1	87.08 %
0.3	84.21 %
0.5	80.21 %
0.7	73.58 %




Fig. 8. Check-first PW-RW ($w = 5$) with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several l in a scale-free network.

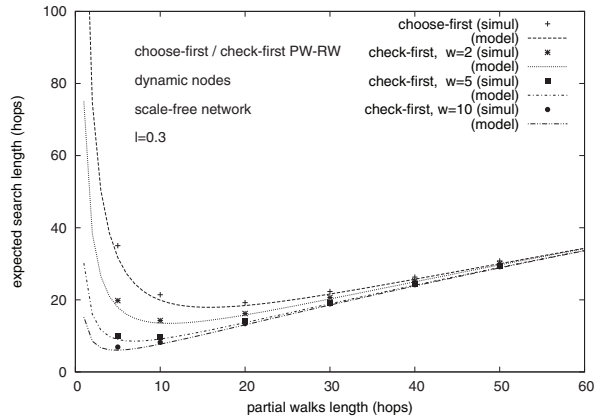


Fig. 9. PW-RW with dynamic nodes: expected search length \bar{L}_s vs. PW length s for choose-first PW-RW and check-first PW-RW ($w = 2, 5, 10$) in a scale-free network with $d = 0.3$.

Table 4 shows the reductions in expected search length achieved by choose-first PW-RW relative to RW searches in scale-free networks with dynamic nodes. As for the dynamic resources case (Table 2), reductions are similar, ranging in this case between 73% and 88%. Reductions for large l are significantly higher than those for large d in the dynamic resources case (about 73% for $l = 0.7$ vs. about 57% for $d = 0.7$). This is explained as above by the effectively shorter PWs in the dynamic nodes case.

Fig. 8 shows the results for the check-first PW-RW mechanism in a scale-free network with dynamic nodes. As in the dynamic resources case (see Fig. 5), both s_{opt} and \bar{L}_{opt} decrease with respect to the choose-first mechanism, since the algorithm checks all PWs available at the node for the resource, increasing the probability of finding it. Fig. 9 shows that both \bar{L}_{opt} and s_{opt} decrease for larger w (the number of PWs per node), as observed also for the dynamic resources case (see Fig. 6).

Table 5

Check-first PW-RW ($w = 5$) with dynamic nodes: reduction of the expected search lengths relative to RW searches for several l .

d	Search length reduction (%)
0.0	94.29 %
0.1	93.76 %
0.3	92.48 %
0.5	90.68 %
0.7	87.73 %

4.2.2. Reduction of the expected search length

Finally, Table 5 shows the reductions in expected search length of check-first PW-RW compared to RW searches. Reduction values range between 87% and 94%. Again, reductions for high l are larger than in the dynamic resources case for large d (around 87% for $l = 0.7$ vs. 77% for $d = 0.7$). Reductions relative to choose-first PW-RW are above 51%, slightly larger than for the dynamic resource case.

4.2.3. Deviations of the predictions of the analytical model

Figs. 7 and 8 show deviations of the proposed analytical model with respect to simulation results. These deviations are very similar to those for the dynamic resources scenario, discussed in Section 2.2.3. For example, deviations of the analytical model for choose-first PW-RW with dynamic nodes range from 3.2% ($l = 0.0$) to 9.4% ($l = 0.7$) when $s = 20$, and from 3.7% to 10.3% when $s = 300$. As for the dynamic resources case, these deviations may be caused by revisits within PWs. Our analytical models do not take this effect into account, resulting in optimistic predictions (see Section 2.2.3).

5. Robustness of PW-RW to parameter perturbations

In this section, we assess the robustness of the proposed mechanisms to variations of some of the parameters, relative both to the network and to the mechanisms themselves. For this, we use the analytical models defined and validated against simulation results in the previous sections. The PW-RW mechanisms have two parameters: the length of the PWs (s) and the number of PWs per node (w). The impact of variations of s on the expected search length has already been shown in Figs. 4 to 9, revealing the existence of an optimal length for PWs (s_{opt}), which achieves minimum expected search length. As for w , it is only relevant for check-first PW-RW, since all the PWs available in the node are checked for the resource. Figs. 6 and 9 show that the minimum expected search length (and also s_{opt}) decreases as w is increased, as expected.

Regarding network parameters, we have already studied the impact of the dynamics of resources (through d , the probability that a resource disappears from a node) and of the dynamics of nodes (through l , the probability that a node leaves the network), observing that the minimum expected search lengths remain reasonably low even for high volatility probabilities (see Figs. 4 to 9).

We now investigate the effect of p_{res} , the probability that a node has an instance of the resource looked for, which in turn determines the average number of instances in the network. We write p_{res} as a function of a multiplicative perturbation factor c for the chosen baseline value (0.01), i.e., $p_{res} = 0.01 \cdot c$. Then we vary c in the range $[0.1, 2]$ and observe the performance of the mechanisms in a scale-free network with dynamic resources ($d = 0.3$). We have found that the value of s_{opt} itself varies when p_{res} is perturbed, ranging from 44 to 10 for choose-first and from 19 to 4 for check-first. We are presented with two options here: calculate the expected search lengths with s fixed to the s_{opt} corresponding to the baseline p_{res} (14 for choose-first and 6 for check-first), or




Fig. 10. PW-RW with dynamic resources: expected search length \bar{L}_s vs. c , the perturbation of p_{res} in a scale-free network with $d = 0.3$.

use the actual s_{opt} for each value of p_{res} (in a hypothetical adaptive PW-RW). Fig. 10 compares both cases for choose-first PW-RW (upper curves) and check-first PW-RW (lower curves). Expected search lengths grow for $c < 1$ (as the number of resources decreases below the baseline), and diminish for $c > 1$ (as the number of resources increases over the baseline), as expected. More interesting is the fact that the proposed mechanisms (with fixed s) and the hypothetical adaptive versions achieve very similar expected search lengths in a wide range of c . This suggests that the PW-RW mechanisms are robust to variations of p_{res} , with little performance degradation due to the suboptimality of s when p_{res} varies around a baseline value. Results are similar for networks with dynamic nodes, and are not shown here.

Finally, we have checked that the expected search lengths (and also s_{opt}) do not depend on the network size (N), as long as p_{res} and the average degree of the network remain unchanged. This is as expected, since fixing p_{res} means that the number of resource instances is proportional to the network size, and fixing the average degree means that the number of choices of a random walk to visit any endpoint in the network (the expected number of neighbors of the current node) remains the same.

6. Cost of the PW-RW mechanisms

In our proposed searching mechanisms, we assumed that at some initial time all network nodes compute a number w of PWs that are used for an interval of time of length T . After that, the searches performed in this interval use these PWs, and thus the cost of their computation must be added to the cost of the searches themselves. In this section we show that such an increase is not relevant in the average cost per search, regardless of the value of T .

First, we observe that since the information associated with those initial PWs is degraded due to the network dynamics, new PWs are computed at $t=T$ to capture the current network state. These new PWs are then used for another interval of length T , and so on. The election of the length of the PW recomputation interval T influences the total cost with two opposite effects. For longer intervals, more searches share the cost of the computation of PWs, lowering the cost per search. However, the information associated with the PWs gets more degraded for longer intervals, causing longer searches. This suggests the existence of some optimal T_{opt} that minimizes the total cost per search.

To investigate this we use C_t , the average total cost per search for interval length T , as defined in [17] for the static network case, as the goodness metric to optimize. This total cost, defined as the number of messages, takes into account the cost of the average




Fig. 11. Choose-first and check-first PW-RW in dynamic networks: relative increment of the expected search length \bar{L}_s vs. (left) the fraction d of resource instances d that disappear, and (right) the fraction l of nodes that leave the network, in a scale-free network.

search (\bar{L}_s messages plus one to notify the source) and the cost of the computation of each of the w PWs of a node (s messages plus one to notify the node), divided by the number of searches b performed per node in the interval:

$$C_t = (\bar{L}_s + 1) + \frac{w}{b}(s + 1). \quad (23)$$

In dynamic networks, however, \bar{L}_s varies with time during the interval as information progressively degrades, so it needs to be averaged over T . In addition, since we intend to study the average total cost as a function of T , the number of searches per node in the interval depends on T as $b = \lambda T$, if we assume that the number of searches per node per time unit is a constant that we denote λ . The previous equation then becomes:

$$C_t(T) = \frac{1}{T} \int_0^T (\bar{L}_s(t) + 1) \cdot dt + \frac{w}{\lambda T}(s + 1). \quad (24)$$

$\bar{L}_s(t)$ is determined by the stochastic process of deaths that is governed by μ_k , the rate of departures of resource instances or nodes when the system is in state k (the number of instances or nodes). A rigorous approach to this problem would require a complete stochastic characterization of the behavior of a system with a large number of states, and is out of the scope of this work. For our purpose, we will focus on intervals of length T in which a relatively small fraction of resources or nodes leave the system, so that the expected search length does not overly increase. For these intervals, we will assume that μ remains constant. This is clearly just an approximation to the behavior of a real system, but it greatly simplifies the analysis allowing us to draw interesting practical conclusions.

To obtain an estimation of $\bar{L}_s(t)$ under these assumptions, we first seek $\bar{L}_s(d)$, the relation between the expected search length and the fraction d of resource instances that disappear. Similarly, we need to find $\bar{L}_s(l)$ for networks with dynamic nodes. The following discussion focuses on the dynamic resources case for conciseness. To find $\bar{L}_s(d)$, which we note is independent from time, we use the analysis of the PW-RW mechanisms presented in the previous sections.

Fig. 11 shows the increments in the expected search length for the choose-first and check-first PW-RW mechanisms in a scale-free network with dynamic resources (left) and dynamic nodes (right). Increments are obtained relative to the expected search length for the static case (e.g., $(\bar{L}_s(d) - \bar{L}_s(0))/\bar{L}_s(0)$). We observe that these dependencies are approximately linear, with a slope $\alpha \approx 2$ for dynamic resources and $\alpha \approx 1$ for dynamic nodes. We can then write $\bar{L}_s(t) \approx \bar{L}_s(0)(1 + \alpha d(t))$. Now, if μ is the departure rate of




Fig. 12. Choose-first and check-first PW-RW with dynamic resources: average total cost per search C_t vs. the length of the PW recomputation interval T , in a scale-free network.




Fig. 13. Choose-first and check-first PW-RW with dynamic nodes: average total cost per search C_t vs. the length of the PW recomputation interval T , in a scale-free network.

resource instances (or nodes), and assuming it constant for the intervals of interest, we obtain an estimation for the fraction of instances that have disappeared at time t as $d(t) \approx \mu t/N$. With all the above we can rewrite Eq. 24 as:

$$C_t(T) = \frac{1}{T} \int_0^T (\bar{L}_s(0) \cdot (1 + \alpha \frac{\mu}{N} t) + 1) \cdot dt + \frac{w}{\lambda T}(s + 1) = (\bar{L}_s(0) + 1) + \frac{\bar{L}_s(0) \cdot \alpha \mu}{2N} T + \frac{w(s + 1)}{\lambda} \frac{1}{T}. \quad (25)$$

A simple calculation on the previous equation allows us to obtain an expression for the interval T_{opt} that minimizes the average total cost per search:

$$T_{opt} = \sqrt{\frac{2Nw(s + 1)}{\bar{L}_s(0) \cdot \alpha \mu \lambda}}. \quad (26)$$

Figs. 12 and 13 show plots of the average total cost per search $C_t(T)$ for the choose-first and check-first PW-RW mechanisms in networks with dynamic resources and with dynamic nodes, respectively, for several values of μ and λ . All curves have a linear behavior for large T (where cost is dominated by longer searches), and a rapid increase for small T (where cost is dominated by PW computation). The minimum cost, marked with a dot, can be found in the region in between. Interestingly, this intermediate region is observed to be quite flat, which allows to reduce the interval T below T_{opt} with no relevant increase in the average cost per search, but with significant reductions in the expected search length (see Fig. 11). This also means that the PW-RW mechanisms are not very

sensitive to the election of the length T of the recomputation interval, given that the system remains in the intermediate region. Finally, the plots of $C_r(T)$ also show that the cost increases for more volatile networks (higher μ) and decreases for more frequent searches (higher λ), as it was expected.

7. Conclusions

We have proposed a mechanism to locate a desired resource in randomly built networks with dynamic resource behavior (resource instances can appear and disappear) or dynamic node behavior (nodes can join and leave the network). The mechanism is based on building a total walk with partial walks that are precomputed as random walks and available at each network node. When precomputing each partial walk, information on the resources held by its nodes is stored and associated to it. This information is used by the searches, so that they can jump over partial walks in which the desired resource is not located. Two versions of the mechanism have been described. In the *choose-first* version, one of the partial walks at the current node is randomly chosen, and then checked for the desired resource. In the *check-first* version, all the partial walks of the node are checked for the resource, and then one is randomly chosen among those in which the resource was found. We have presented an analytical model that predicts the expected search length achieved by the two versions of the mechanism. Simulation experiments have been used to validate the model and to assess the effect of resource and node dynamics. We have found that the *choose-first* version achieves large reductions of the average search length in relation to searches based on simple random walks. These reductions remain significant even in the face of high volatility of resources or nodes. The *check-first* version produces larger reductions as we increase the number of partial walks precomputed at each node (at the corresponding extra cost). Results have been found to be very similar for networks with different degree distributions (k -regular, Erdős-Rényi and scale-free). Finally, we have analyzed the cost of the PW-RW mechanisms, concluding that the choice of the length of the PW precomputation interval does not have a significant impact on the cost in a wide range of interval lengths.

An interesting future work line for this study is to measure the improvement in the search length that can be obtained by using different strategies to choose one of the partial walks available in a node. Another possibility to shorten search lengths is to use more intelligent (and more costly) variations of random walks instead of simple random walks.

Appendix A. Choose-first PW-RW with dynamic resources

This Appendix contains results for regular and ER networks that complement those given for scale-free networks in the main text, as well as some additional results. Results for the three network types are very similar in all cases.

Figs. A.14 and A.15 show experimental and analytical results for the expected search length achieved by *choose-first* PW-RW in regular and ER networks with dynamic resources, respectively, for several values of p , the probability of a resource disappearing from a node.

• Length distributions and dispersion:

The use of partial walks also affects the shape of the probabilistic distribution of search lengths. Figs. A.16 and A.17 show the distributions for simple random walks (RW) searches and for PW-RW searches, for $s = 10$ and for several values of d , obtained from the experiments with regular and ER networks with dynamic resources. Instead of the slowly decaying distribution of RWs, the proposed mechanism exhibits search length

distributions that show a maximum frequency for a small search length and then decay much faster than the random walk distribution. We also note that the search length for the maximum frequency (9 in this case) is independent from the dynamic behaviour of resources (d). The search length distributions of PW-RW have therefore lower standard deviation than random walk searches. For $d = 0.3$, for instance, the standard deviation in regular, ER and scale-free networks are, respectively, 21.06, 21.34 and 21.58. Correspondingly, the mean search lengths are 23.56, 23.72 and 23.81. As a reference for compari-




Fig. A.14. *Choose-first* PW-RW with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in a regular network.




Fig. A.15. *Choose-first* PW-RW with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in an ER network.




Fig. A.16. *Choose-first* PW-RW with dynamic resources: search length distributions for RW searches and for PW-RW searches, in a regular network with $s = 10$ and several d .




Fig. A.17. Choose-first PW-RW with dynamic resources: search length distributions for RW searches and for PW-RW searches, in a ER network with $s = 10$ and several d .

son, random walk searches on the scale-free network registered a standard deviation and mean of search lengths of 116.48 and 113.56.

Appendix B. Check-first PW-RW with dynamic resources

Figs. B.18 and B.19 show the expected search lengths achieved by check-first PW-RW in regular and ER networks with dynamic resources, respectively, for several values of p .




Fig. B.18. Check-first PW-RW ($w = 5$) with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in a regular network.




Fig. B.19. Check-first PW-RW ($w = 5$) with dynamic resources: expected search length \bar{L}_s vs. PW length s for several d in a ER network.

In addition, Fig. B.20 shows the dependency between the expected search lengths achieved by check-first PW-RW and the number of PWs per node (w), in a regular network with dynamic resources. Results for ER networks are very similar.

Appendix C. Choose-first PW-RW with dynamic nodes

Figs. C.21 and C.22 show the expected search lengths achieved by choose-first PW-RW in regular and ER networks with dynamic




Fig. B.20. PW-RW with dynamic resources: expected search length \bar{L}_s vs. PW length s for choose-first PW-RW and check-first PW-RW ($w = 2, 5, 10$) in a regular network with $d = 0.3$.




Fig. C.21. Choose-first PW-RW with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several d in a regular network.




Fig. C.22. Choose-first PW-RW with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several d in a ER network.

nodes, respectively, for several values of l , the probability of a node leaving the network.

Appendix D. Check-first PW-RW with dynamic nodes

Figs. D.23 and D.24 show the expected search lengths achieved by *check-first PW-RW* in regular and ER networks with dynamic nodes, respectively, for several values of l .




Fig. D.23. *Check-first PW-RW* ($w = 5$) with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several d in a regular network.




Fig. D.24. *Check-first PW-RW* ($w = 5$) with dynamic nodes: expected search length \bar{L}_s vs. PW length s for several d in a ER network.




Fig. D.25. PW-RW with dynamic nodes: expected search length \bar{L}_s vs. PW length s for *choose-first PW-RW* and *check-first PW-RW* ($w = 2, 5, 10$) in a regular network with $d = 0.3$.

Finally, Fig. D.25 shows the dependency between the expected search lengths achieved by *check-first PW-RW* and the number of PWs per node (w), in a regular network with dynamic nodes. Results for ER networks are very similar.

References

- [1] L.A. Adamic, R.M. Lukose, A.R. Puniyani, B.A. Huberman, Search in power-law networks, *Phys. Rev. E* 64 (046135) (2001).
- [2] W. Aiello, F. Chung, L. Lu, A random graph model for massive graphs, in: Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing, in: STOC '00, ACM, New York, NY, USA, 2000, pp. 171–180, doi:10.1145/335305.335326.
- [3] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, *Rev. Mod. Phys.* 74 (1) (2002) 47–97, doi:10.1103/RevModPhys.74.47.
- [4] D. Amit, G. Parisi, L. Peliti, Asymptotic behavior of the 'true' self-avoiding walk, *Phys. Rev. B* 27 (1983) 1635–1645.
- [5] Y. Chawathe, S. Ratnasamy, N. Lanham, S. Shenker, Making gnutella-like p2psystems scalable, in: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)Karlsruhe, Germany, 2003, pp. 407–418.
- [6] L. da Fontoura Costa, G. Travieso, Exploring complex networks through random walks, *Phys. Rev. E* 75 (016102) (2007).
- [7] A. Das Sarma, D. Nanongkai, G. Pandurangan, P. Tetali, Efficient distributed random walks with applications, in: Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing, in: PODC '10, ACM, New York, USA, 2010, pp. 201–210. <http://doi.acm.org/10.1145/1835698.1835745>.
- [8] S. Dolev, E. Schiller, J.L. Welch, Random walk for self-stabilizing group communication in ad hoc networks, *IEEE Trans. Mobile Comput.* 5 (7) (2006) 893–905.
- [9] S. Fortunato, A. Flammini, Random walks on directed networks: the case of pagerank, e-print physics (0604203) (2006).
- [10] C. Gkantsidis, M. Mihail, A. Saberi, Random-walks in peer-to-peer networks: algorithms and evaluation, *Perform. Eval.* 63 (2006) 241–263.
- [11] B.D. Hughes, Random walks and random environments, vol. 1, Clarendon Press, Oxford, 1995.
- [12] M.A. Jovanovic, Modeling large-scale peer-to-peer networks and a case study of gnutella, University of Cincinnati, Cincinnati, Ohio, 2001 Master's thesis.
- [13] M.A. Jovanovic, F.S. Annexstein, K.A. Berman, Scalability issues in large peer-to-peer networks - a case study of Gnutella, Technical Report, Technical report, University of Cincinnati, 2001.
- [14] B. Krishnamurthy, J. Wang, Y. Xie, Early measurements of a cluster-based architecture for p2p systems, in: Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement, in: IMW '01, ACM, New York, USA, 2001, pp. 105–109, doi:10.1145/505202.505216.
- [15] C. Law, K.-Y. Siu, Distributed construction of random expander networks, in: Proceedings of the Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2003), 3, 2003, pp. 2133–2143.
- [16] S.H. Lee, P.-J. Kim, H. Jeong, Statistical properties of sampled networks, *Phys. Rev. E* 73 (016102) (2006).
- [17] V. López Millán, V. Cholvi, A. Fernández Anta, L. López, Improve resource location with locally precomputed partial random walks, *Computing* 97 (9) (2015) 871–891, doi:10.1007/s00607-013-0353-x.
- [18] L. Lovász, Random walks on graphs: a survey, in: *Combinatorics, Paul Erdős is eighty*, vol. 2, Keszthely, Hungary, 1993, pp. 1–46.
- [19] Q. Lv, P. Cao, E. Cohen, K. Li, S. Shenker, Search and replication in unstructured peer-to-peer networks, in: ICS '02: Proceedings of the 16th International Conference on Supercomputing, ACM, New York, USA, 2002, pp. 84–95. <http://doi.acm.org/10.1145/514191.514206>.
- [20] Q. Lv, S. Ratnasamy, S. Shenker, Can heterogeneity make gnutella scalable? in: Revised papers from the First International Workshop on Peer-to-Peer Systems, Cambridge, United States, 2002, pp. 94–103.
- [21] I. Mabrouki, X. Lagrange, G. Froc, Random walk based routing protocol for wireless sensor networks, in: Proceedings of the Second International Conference on Performance Evaluation Methodologies and Tools (ValueTools '07), ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2007, pp. 1–10.
- [22] G.S. Manku, M. Naor, U. Wieder, Know thy neighbor's neighbor: The power of lookahead in randomized p2p networks, in: Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing, in: STOC '04, ACM, New York, NY, USA, 2004, pp. 54–63, doi:10.1145/1007352.1007368.
- [23] P. Maymounkov, D. Mazières, Kademia: A peer-to-peer information system based on the xor metric, in: Revised Papers from the First International Workshop on Peer-to-Peer Systems, in: IPTPS '01, Springer-Verlag, London, UK, 2002, pp. 53–65.
- [24] M.E.J. Newman, A.-L. Barabási, D.J. Watts, The Structure and Dynamics of Networks, Princeton University Press, Princeton, New Jersey 08540, USA, 2006.
- [25] M.E.J. Newman, S.H. Strogatz, D.J. Watts, Random graphs with arbitrary degree distributions and their applications, *Phys. Rev. E* 64 (026118) (2001).
- [26] R. Pastor-Satorras, A. Vespignani, Evolution and structure of the internet: a statistical physics approach, Cambridge University Press, New York, NY, USA, 2004.
- [27] S. Ratnasamy, P. Francis, M. Handley, R. Karp, S. Shenker, A scalable content-addressable network, *SIGCOMM Comput. Commun. Rev.* 31 (4) (2001) 161–172, doi:10.1145/964723.383072.

- [28] L. Rodero-Merino, A. Fernández Anta, L. López, V. Cholvi, Performance of random walks in one-hop replication networks, *Comput. Netw.* 54 (5) (2010) 781–796.
- [29] N. Sadagopan, B. Krishnamachari, A. Helmy, Active query forwarding in sensor networks, *Ad hoc Netw.* 3 (2005) 91–113.
- [30] G. Slade, The self-avoiding walk, in: *Surveys in Stochastic Processes*, Proceedings of the 33rd SPA Conference, EMS Series of Congress Reports, 2009. In preparation.
- [31] I. Stoica, R. Morris, D. Karger, M.F. Kaashoek, H. Balakrishnan, Chord: A scalable peer-to-peer lookup service for Internet applications, in: *ACM SIGCOMM Computer Communication Review*, 31, ACM, 2001, pp. 149–160.
- [32] B. Tadić, Adaptive random walks on the class of web graphs, *Eur. Phys. J. B.* 23 (221) (2001).
- [33] S.-J. Yang, Exploring complex networks by walking on them, *Phys. Rev. E.* 71 (016107) (2005).



Víctor M. López Millán is an Assistant Professor at Universidad CEU San Pablo in Spain. He graduated in Telecommunication Engineering from Universidad Politécnica de Madrid in 1990, obtained his MS degree in Computer Science from University of Illinois at Urbana-Champaign in 1993 (Fulbright Grant) , and his PhD from Universidad Rey Juan Carlos (Spain) in 2014. Before joining USP in 2004, he worked for several ICT companies including Telefónica Research and Development, Alcatel and Sema Group, as a software engineer, telecommunication systems engineer, project leader and account manager. His research interests include computer networks and distributed algorithms, with an emphasis on complex networks.



Vicent Cholvi graduated in Physics from the Universitat de València (Spain) and received his doctorate in Computer Science in 1994 from the Polytechnic University of Valencia (Spain). In 1995, he joined the Jaume I University in Castelló (Spain) where he is currently a Professor. His interests are in distributed and communication systems.



Antonio Fernández is a research professor at IMDEA Networks, Spain, since 2010. Previously, he was on the faculty of the Universidad Politécnica de Madrid and at the Universidad Rey Juan Carlos in Madrid. He graduated in Computer Science from the Universidad Politécnica de Madrid in 1991, and got a PhD in Computer Science from the University of Southwestern Louisiana in 1994. His research interests include data communications, computer networks, parallel and distributed processing, algorithms, and discrete and applied mathematics.



Luis López Fernández obtained his PhD in Computer Science at Universidad Rey Juan Carlos in 2003. He has a Telecommunications Engineering degree by the ETSIT-UPM , obtained in 1999, and a Telecommunications Engineering degree by the ENST - Télécom Paris , also in 1999. Before starting his academic career, he worked in several IT companies including Texas Instruments Inc (working as Hardware Engineer) and Visual Tools S.A (working as Development Engineer). He is currently Professor at the DITTE department belonging to Universidad Rey Juan Carlos, where he coordinates/collaborates in several undergraduate and master courses in Computer Science and Telecommunications Engineering. His current research interests are concentrated on the applications of Complex Networks and Game Theory in the field of Computer Science in general, and IP networks in particular. Luis has coauthored more than 50 research publications in top scientific journals and conferences including Physica Review E, Theoretical Computer Science, IEEE Electronics Letters, etc.