Greedy Routing in the Internet: Is it a Solution?

Pedro Rodrigues
CITI and Departamento de Informática
Faculdade de Ciências e Tecnologia, FCT
Universidade Nova de Lisboa
2829-516 Caparica, Portugal
pedro.m.f.rodrigues@gmail.com

José Legatheaux Martins
CITI and Departamento de Informática
Faculdade de Ciências e Tecnologia, FCT
Universidade Nova de Lisboa
2829-516 Caparica, Portugal
jose.legatheaux@di.fct.unl.pt

***Abstract*—The current protocol for interdomain routing (BGP) faces several problems which may compromise its future. Being scalability and convergence the biggest cause of concern, forms of scalable routing are a relevant contribution for this discussion.**

This papers presents *greedy routing*, a form of scalable routing with remarkable properties, *e.g.*, the complexity in each node is $O(\#neighbours)$ as opposed to $O(\#nodes)$. The applicability of greedy routing is supported by studies concerning large-scale networks on which the traditional routing schemes do not scale. Being the Internet AS graph one of those networks [1], the application of a greedy routing scheme for interdomain routing could solve its scalability and convergence problems.

***Index Terms*—Routing, Internet Topology, Scale-free Networks, Greedy Routing**

I. Introduction

BGP was initially defined when the Internet had a few hundred ASes and it was fairly limited to academic usage. Over the years, BGP has been updated to follow the evolution of the Internet. Although BGP continues to support the current functioning of interdomain routing, there are some critical aspects that it does not manage quite well: scalability and convergence due to the explosion of IP prefixes; lack of techniques for inbound load balancing; growth of the number of ASes and links; no security mechanism that prevent an AS from advertising arbitrary prefixes, *i.e.*, BGP does not support prefix authentication, ...

There are two major directions that researchers have been taking in order to overcome the identified issues of BGP. On the one hand, short-term fixes may delay the decline of the present architecture or even solve its main problems for some next years. Preventing withdrawal [2], differentiating advertisements [3] and fine-tuning its timers, as well as *flushing* obsolete paths [4] are examples of techniques proposed to improve convergence time and scalability. On the other hand, to sort out in a more definitive way the convergence, scalability, security, quality of routes problems and adding mobility support to the Internet, a totally new approach and architecture are required. One common factor of the new architectures [5], [6] is the separation of the current IP address scheme into two address spaces: one to represent host location (locator) and other for host identification (identifier). Scalability is greatly improved by hierarchically organising the locator space. This is also a common concern of the proposed alternative routing schemes to BGP [7]–[10].

In parallel with the definition of proposals to the current and future Internet, there has been studies regarding the structure and topological characteristics of the Internet graph. Surprisingly, most of these are common to other large-scale networks, opening the possibility of applying mechanisms from these networks to the Internet. The study of large-scale networks has opened a new exciting field known as *Networking Science*. Researchers acting on this area have proposed a new routing strategy known as *greedy routing* [11]. According to this routing strategy, each node only knows its characteristics and the characteristics of its neighbours. With this information and the characteristics of a given destination, a node is able to *greedily* route a message by selecting the direct neighbour *closest* to the destination. A successful greedy routing scheme is highly scalable, requires less routing state and easily adapts to dynamic networks. The question we pose in this paper is: can greedy routing help to improve the aforementioned issues?

The contributions of this paper are twofold. First, we do a short survey covering the most relevant Networking Science results. Then, we discuss a preliminary approach of applying a greedy routing scheme to interdomain routing.

The rest of this paper is structured as follows. Section II details the most relevant topological characteristics of scale-free networks. Section III describes the greedy routing strategy, its components as well as its suitability for scale-free networks. Section IV presents a preliminary approach of defining a greedy routing scheme for interdomain routing. Section V states the values of evaluation metrics of the greedy routing scheme. Finally, VI discusses the related work and section VII ends this paper with some conclusions and presents future work.

II. Scale-free Networks

The node distribution of several complex networks, such as biological networks, social networks, cellular networks, collaboration networks, citation networks, the Internet backbone [1], follows a power-law distribution. In order to differentiate these networks from other complex networks which degree distributions do not follow a power law, they were designated as scale-free networks since most of their properties are independent of the scale. A power law distribution of the generic degree d is $P(d \geq k) = k^{-\alpha}$. The main property of this distribution is scale invariance: applying a scale factor to the distribution variable leads only to a proportional scaling of the distribution, thus maintaining its properties. The current Internet AS graph also follows a power law with exponent 1,9. The main topological properties that characterise scale-free networks are presented in the following items.

Small World and Network Navigability: In 1969, an interesting experience was performed by Milgram *et al.* [12]. They asked some random individuals (sources) to send a letter to a specific person (destination), from whom they (the sources) only knew his/her name, age, occupation and city of residence. The sources had to pass the letter to people they knew, who were chosen based on the characteristics of the destination in order to maximise the probability of the letter reaching its destination. Surprisingly, 30% of the letters reached their destination and needed only a small number of intermediate people, 5.2 hops on average, even though sources had no global knowledge of the human acquaintance network topology.

Recently, the *small-world* property has been precisely defined as follows: a network holds the small-world property if the shortest paths between any two pair of nodes scales, at most, logarithmically with the network size [13]. Several scale-free networks hold this property, such as social networks and the Internet backbone.

Norros and Reittu [14] have demonstrated that in graphs with N nodes whose degree distribution follows a power law with exponent $\alpha \in [1, 2]$, the distance between any two nodes is of the order $loglogN$. These networks are called *superscalable* or *ultra-small world* graphs.

Network Resilience: Scale-free networks are more resilient to random attacks than random networks since the majority of the network has a low degree [15]. By contrast, in random graphs nodes tend to have a similar and more balanced role in network functioning. Therefore, they are less robust than scale-free networks. However, scale-free networks are less robust under target attacks to high degree nodes.

Assortativity: Assortativity measures selective linking between nodes, *i.e.*, the preference which nodes have to be connected to others of the same type or of other types. While social-networks are highly assortative, since people tend to be related to persons which are similar in some way, other scale-free networks, such as the Internet and Biological networks, are *disassortative* [13]. For instance, the Internet can be simply divided in three groups: high-degree nodes (T1 backbone operators), transit nodes (ISPs) and end nodes (stubs). Although it is very unlikely that stubs are connected to T1s, there are several links between backbone operators and ISPs as well as between ISPs and stubs, which may overcome the number of connections within each group.

Clustering: One characteristic that clearly distinguishes scale-free networks from random graphs is clustering or transitivity [13], *i.e.*, if node X is connected to node Y and node Y is connected to Z, it is highly probable that node X is connected to node Z. In the context of social networks, it means that a friend of your friend is also your friend. We have calculated the average clustering coefficient of the CAIDA AS graph [16], which is $\simeq 0.011$. This is lower than the ones found in other scale-free networks [13] since there is unlikely to have the following relations in the Internet AS graph: (i) two customer of a given provider being connected; (ii) a node having two providers which are connected, if the two providers are not from the core.

Network Construction Model: Derek Price [17] defined the *cumulative advantage* mechanism as the underlying growth principle of scale-free networks construction: the rate at which a node gets new connections is proportional to its degree. This property is sometimes dubbed as ”the rich get richer”.

III. GREEDY ROUTING

The term *greedy routing* was firstly introduced by Jon Kleinberg in [11] to characterise the type of routing used in the experience of Milgram *et al.* [12]: (i) each node has only information regarding its neighbours and the destination; (ii) in each hop, the message is routed to the *nearest* neighbour towards the destination node. The notion of the *nearest* neighbour is given by a distance function among network nodes based on the information associated with each node.

In addition, Jon Kleinberg defined a model in which each node is represented in a coordinate space and it only knows the coordinates of its neighbours. In order to send a message from a source node to a destination node, each node sends the message to the nearest neighbour towards the destination. Geographically-inspired routing is an example of this type of routing strategy.

If successful, a greedy routing algorithm has the following highly interesting properties:

- Small routing state: each node needs only to maintain information regarding its neighbours. There is no routing state maintenance cost in the sense that there is no need to exchange messages in order to perform routing;
- Small routing stretch[1];
- Robustness: in scale-free networks, even if a considerable number of simultaneous failures happen it ensures near full reachability maintaining a small routing stretch value.

In order to build a successful greedy routing algorithm, several interrelated problems must be solved:

- devise a method to map the network topology into a coordinate space, *i.e.*, an embedding of the network;
- construct a distance function acting on nodes coordinates;
- elaboration of a concrete routing algorithm dealing with route optimality criterion and the dead-end problem, *i.e.*, when a message reaches a node where it cannot make any further progress but to get back through an already known path to find an alternative one.

A. Greedy routing in scale-free networks

Boguñá *et al.* [18] have defined a general model, based on the concept of node similarity, as the underlying mechanism to explain the navigability properties and the greedy routing

[1]Stretch is the ratio between the length of paths chosen by a routing algorithm and optimal ones, *e.g.*, shortest paths.

success in scale-free networks. This model does not address the engineering problems of technological networks since it focuses on the topological properties of scale-free networks. The model explains the Milgram's experience and can be applied to other scale-free networks. With this model, nodes characteristics define how similar they are, which is abstracted as a *hidden distance*, which defines a *hidden metric space* guiding the routing on the network and influencing its structure [18]. The (hidden) distance is coupled with the network structure in the following form:

a) the smaller the distance between two nodes, the higher the probability that they are topologically connected;

b) if a given node A is close to node B and node B is close to another node C, then nodes A and C are also close as a consequence of the triangle inequality [2];

c) it is highly probable that the triangular relationship ABC exists in the network topology, which explains the strong clustering of scale-free networks.

The metric space plays a central role in the success of greedy network *navigation* and has a major impact in the embedding and distance function definition.

The navigability of scale-free networks can be illustrated using an example from passenger air travel, using the greedy routing strategy guided by a metric space with a distance function based on a combination of geography and airport size. At each airport it is chosen the next-hop airport which is geographically *closest* to the destination. The navigation process has two symmetric phases. The first phase is a coarse-grained search, in which a *zoom-out* mechanism is applied: from a small local airport to a larger hub at a larger distance [3]. Large hub airports are connected to the majority of other large hub airports. The turning point between the two phases is when the navigation process reaches the closest large hub to the destination. From here begins the second phase, a fine-grained search towards the destination airport.

The navigation process, *i.e.*, the zoom out/zoom in mechanisms, works efficiently if the airport network topology and the underlying metric space exhibits the following two properties:

- the network has enough hub airports to provide an increasing degree sequence during the zoom-out phase;
- the next greedy hop from a remote low-degree node is a node with a higher degree so that greedy paths normally move first to the highly connected network core.

These conditions are fundamental to ensure that local loops do not occur. Specifically to the air travelling example, an airport network without enough clustering would result in a path with several hops among small nearby airports, reaching the destination after many hops. In the worst case, when travelling through those small airports, it is possible to reach one that does not have any other connections closer to the destination, facing a dead end. This suggests that scale-free networks are suitable for greedy routing since they comprise a large number of hubs, *i.e.*, high-degree nodes, as well as strong clustering.

[2] $d(A,C) \leq d(A,B) + d(B,C)$, being d the distance function.

[3] Note that there are other flights (hops) between the small airport and the large hub using medium airports.

IV. Greedy Routing for Interdomain Routing : a Preliminary Approach

The application of greedy routing in the Internet comprises: *a*) the construction of a mapping of the network topology into a coordinate space; *b*) the definition of a distance function to be used in the coordinate space; *c*) the definition of a concrete greedy routing algorithm. There are two metrics that are commonly used to evaluate the suitability of a greedy routing algorithm to the Internet:

- stretch using shortest paths as reference;
- success ratio, the percentage of nodes which are reachable through the greedy routing algorithm.

While in some networks, *e.g.*, wireless networks, overlay, *etc.* it is possible to compute geographic coordinates, or synthetic ones for their nodes, based on latency, these methods are not well suited for the Internet AS Graph for various reasons:

- there are several ASes which do not have a well-defined geographic location, *e.g.*, tier-1 ASes;
- typically, coordinates devised from latency are dynamically computed, thus the coordinate system does not ensure convergence;
- latencies are not symmetric, *i.e.*, many routes are asymmetric ;
- it does not always ensure the triangle inequality [19].

We have defined a method to assign synthetic coordinates based on routing requirements. We follow the approach of LISP [5] which divides the IP address space into two address ones: locator and identifier space. Although hosts continue to use IP addresses (End-point Identifiers-EID) to communicate with each other, EIDs concern only with the domain (AS) where the host resides. In order to reach a host a Routing Locator (RLOC) is needed, *i.e.*, IP addresses hierarchically organised and bounded to a domain (AS). Mapping EIDs to RLOCs is outside the scope of this paper.

We take NIRA [10] as a model to organise the locator space. We consider a set of ASes which follow a definition of the presence of a core in scale-free networks, as the Internet AS graph. In NIRA tier-1 (core) ASes that have globally unique IP prefixes from which they allocate non-overlapping sub-prefixes to their customers. This induces a provider-customer hierarchy from each tier-1 AS composed by the set of its customers, direct and indirect, *i.e.*, which have a sub-prefix derived from the prefix of the tier-1 AS.

In our model each AS has a coordinate for each provider-hierarchy it pertains, which represents a way of reaching the core. As each prefix is bounded to one AS, a mapping component from a prefix to the corresponding set of coordinates is needed in order to perform routing at the inter-AS level. We assume that such component is present in an architecture where our routing scheme could be applied. In addition, if two ASes having a provider-customer or peering relationship have more than one direct link between them, we only consider one link in our routing scheme.

A. Provider-Customer Hierarchies

We use the CAIDA AS graph [16], a snapshot of the real AS graph with 33508, as our working foundation. This graph comprises four types of links: *a*) customer-provider; *b*) provider-customer; *c*) peer-peer; *d*) sibling-sibling link[4].

We divide the locator space in non-disjoint provider-customer hierarchies, *i.e.*, an AS can be in more than one provider-customer hierarchy. Each hierarchy is rooted in a different AS from the core. Following the kernel definition presented in [14], the core can be modelled as follows:

- $K_n\ (V_n,\, E_n)$ - a clique composed by peer-peer links
- $\forall v \in V_n : v \in V_o,\ degree(v) > \sqrt{\#(V_o)},\ \nexists e(v, v_i, -1) \in E_o, v_i \in V_o$

This definition of the core comprises all the ASes that have a degree $> \sqrt{\#(V_o)}$ [14] and that are transit-only, *i.e.*, which are not customer of any other AS. We extend the above definition to include ASes with degree $> \sqrt{\#(V_o)}$, that are not transit-only but are only customers of ASes from the initial definition of the core. As a result, the core is composed by 14 ASes.

In table IV-A it is shown the number of ASes which pertain to a given number n of hierarchies, $n \in \{1..\#(V_k)\}$. Two patterns can be identified in table IV-A: 17 % of total nodes pertain to few hierarchies (1 or 2) and 79 % of total nodes are in all hierarchies.

TABLE I
NUMBER OF NODES IN n HIERARCHIES, $n \in \{1..\#(V_k)\}$

# Nodes	Hierarchies
2700	1
2507	2
697	3
899	4 - 9
26496	14

B. An Euclidean Metric Space

We have defined a Metric Space (ξ, ϱ) in the Euclidean Plane. The set ξ is described in the next section, as a sub-set of $\mathbb{R}^2$. In the following subsection the metric ϱ will be defined.

1) Coordinate Distribution Model: The distribution of coordinates comprises two phases: firstly, the coordinates of core nodes are manually assigned; secondly, from the coordinate of the core node, coordinates are set along the correspondent provider-customer hierarchy. This process can be defined as follows:

- $\forall v \in V_k$, its coordinates are manually assigned;
- $\forall v_k \in V_k, \forall v \in V_{v_k}$ with coordinates = (v_x, v_y), $\forall v_s \in successors_{provider-customer}(v, G_{v_k})$, v assigns $(f(v_x, v_s), f(v_y, v_s))$ to v_s

We concentrate our design in the $\mathbb{R}^{2+}$ part of the Euclidean Plane, though it can be applied to the all $\mathbb{R}^2$ as it will be discussed later. Core nodes are disposed in a semi-circumference with radius = $\#(V_k)^4$ [20]. Each core node is associated with an arc of the semi-circumference and is placed in the middle of it, as can be seen in the figure 1. The lines which pass through the ends of each core node arc delimit the region associated with that node, where the coordinates of its provider-customer hierarchy nodes will be assigned.

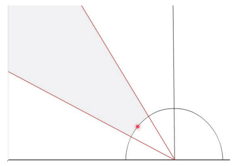

Fig. 1. Region of a core node

For simplicity, we will only define the coordinate assignment method for a core node, and then describe the difference for further levels of the provider-customer hierarchy. Each core node has the following information: *a*) distance to centre: core radius; *b*) its angle; *c*) *growth factor* of distance to centre; *d*) boundaries of its region in the x-Axis: min, max and its width $\chi = (max - min) =$ kernel_radius$/\#(V_k)$.

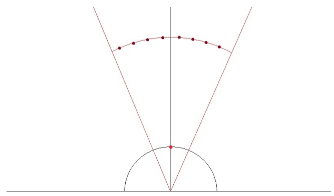

Fig. 2. Coordinate assignment to customers of a core node

The kernel node customers are placed in an arc of the semi-circumference centred in the origin with $radius =$ distance_to_centre $\times$ growth_factor, as in figure 2. Consider n_c as the number of customers of the kernel node, the region width χ is divided in $(3 \times n_c + n_c + 2)$ spaces with length δ. Starting from $min + \delta$, each customer is placed at the middle of its region and within $3 \times \delta$ to the consecutive customer. As all customers are placed in a semi-circumference centred in the origin, they all have the same norm. That spaces are needed to ensure that in each hop, when moving towards a destination node, it is chosen the provider which will lead to the destination node, *i.e.*, the closest node to the destination is the provider of the provider of the ... of the destination node. Moreover, each customer has a region of $3 \times \delta$ width and correspondent boundaries $min_c = min + \delta + 3 \times (i-1) \times \delta$ and $max_c = min + \delta + 3 \times (i-1) \times \delta + \delta$, being $i \in \{1, ..., n_c\}$. Finally, a rotation of the kernel node angle is applied to determine the final coordinates of each customer.

In further levels of a provider-customer hierarchy the same method is applied, though with different values from each parent node. The *growth factor* is the same for all levels as well as the kernel node angle in a given provider-customer hierarchy. The distance to centre and the boundaries of each node region are the values which are specific to each node.

[4]A sibling-sibling link connects two ASes managed by the same company.

2) Metric: Given two points in $\xi \subset \mathbb{R}^2$, $a(x_a, y_a)$ and $b(x_b, y_b)$, the distance between a and b is given by the following expression:

$$\varrho(a, b) = (\parallel a \parallel + \parallel b \parallel) \times \epsilon(a, b)$$

where $\epsilon(a, b)$ corresponds to the euclidean distance in $\mathbb{R}^2$ and $||a|| = \epsilon(a, (0, 0))$.

C. Greedy Routing in an Euclidean Metric Space

The classical greedy forwarding strategy does not consider the distance between the current node and its neighbours, only the distance between the neighbour and the destination node. We have made a slight modification to the classical greedy forwarding strategy. In our greedy routing algorithm not only the closest neighbour to the destination is selected but the one which is also closer to the current node. Formally, in each hop it is selected the node which matches the following condition:

$$min(\ \varrho(current, neighbour) + \varrho(neighbour, destination)\ )$$

We divide the functioning of the greedy routing algorithm for interdomain routing in two modes: routing in different hierarchies and routing in the same hierarchy. As regards to routing in different hierarchies, a normal route is of the form: a chain of customer-provider links towards the core, a peering link in the core followed by a chain of provider-customer links towards the destination. The distribution of coordinates along with the metric ϱ lead to the following route: shortest path from the source node to the core, one hop in the core, followed by the shortest path from the core to the destination. The choice of coordinates determines from which provider-customer hierarchies the message goes through, towards the core and towards the destination, as in NIRA [10]. The complete proof that the metric ϱ leads to the mentioned path between ASes from different hierarchies is presented in [20].

In what concerns routing in the same hierarchy, the shortest path would be a chain of customer-provider links, an *inversion* of the path, *i.e.*, from customer-provider links to provider-customer links, in a common provider of source and destination nodes, or a common provider of source and destination providers, and so on, followed by a chain of provider-customer links towards the destination. However, with metric ϱ, the choice of the *inversion* node can be faulty, *i.e.*, it can lead to a dead-end as can be seen in figure 3. Node $AS5$ is the neighbour of $AS2$ that is closer to the DST node, though it is impossible to reach DST node from $AS2$. As source and destination nodes are too close, we enforce the following path: shortest path to the core node of that hierarchy followed by the shortest path towards the destination. This path is identified in figure 3 by blue links: $SRC{-}AS4, AS4{-}AS2, AS2{-}AS1, AS1{-}AS3, AS3{-}AS6$ and finally $AS6 - DST$.

V. Evaluation

Using the CAIDA AS graph [16] we have verified that our approach achieves full success ratio [20], *i.e.*, it is possible to reach every pair of nodes in the network without reaching a dead-end. However, the average stretch obtained for the base greedy routing scheme was 1.401 [20]. This value is due to the fact that the above base greedy routing algorithm only explores strict hierarchically paths, ignoring intra and inter-hierarchies peering links.

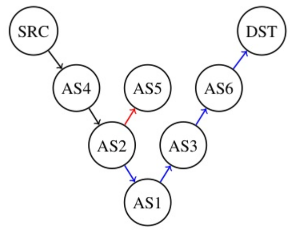

Fig. 3. Invalid Chosen of the *Invert* AS

Optimisation of the Base Scheme

Each AS has its own region from which it assigns coordinates and sub-regions to its clients, which in turn repeat the same process to their clients. Having a coordinate and a region, determining if the coordinate pertains to the region is straightforward. Therefore, the two ASes involved in a peering link, either inter or intra-hierarchies, can exchange their own regions in order to verify the possibility of usage of that peering link. Remember that a peering link can only be used by the direct and indirect customers of the two ASes involved in the peering link.

When choosing the next hop, if the neighbour chosen is connected via a peering link, the following conditions have to be verified:

- the source coordinates belong to a region of the current node, *i.e.*, it is a customer of the current node;
- the destination coordinate refers to a region of the other node in the peering link.

Each AS can control to which customer(s) it allows the peering link to be used, by sending sub-regions of its own region. Although this increases the amount of data that each node has to maintain, it is only exchanged between the two ASes involved in the peering link. Besides rare modifications, the data concerning the control of peering links does not increase the network traffic. This optimisation, albeit not being a pure greedy routing strategy and slightly increasing routing state, would reduce the overall routing stretch since several packets would follow a *shortcut* using a peering link, instead of having to pass through the core.

VI. Related Work

Earliest versions of greedy routing applications relied upon real geographic position information [21], [22], *e.g.*, as determined by a GPS device and wireless ad-hoc routing scenarios seemed to be the ideal context to study if the approach would be viable. However, there are several problems concerning with wireless communication which complicate the application of greedy routing in those networks: *a*) decreased signal strength; *b*) unknown obstacles; *c*) weather conditions; *d*) hidden terminal problem; *e*) multipath propagation. Nonetheless, even if the embedding problem has been solved, another problem arises: how a node knows the coordinates of a destination node? Several proposals have been presented in the literature, though they are limited by the aforementioned problems.

To our best knowledge, there is no solid proposal of a greedy routing scheme for interdomain routing. There are some embedding proposals for arbitrary graphs [23], [24] that rely upon a spanning tree, though this leads to a non-utilisation of several links. Moreover, some geographically inspired proposals concerning the Internet [25], [26] rely upon volatile distance measures, *i.e.*, latency, which leads to the continuing computation of coordinates.

VII. Conclusions and Future Work

This proposal is a first essay of using a greedy routing approach for interdomain routing in the Internet. Next we discuss how some of the features of BGP could be performed using our greedy routing scheme as well as how it improves some of the identified critical issues of BGP. We conclude revising several open problems.

BGP uses policy filters in order to guarantee valley-free paths and correct usage of peering links. The paths induced by our greedy routing scheme do not violate those navigability restrictions and do not lead to dead-ends.

Each AS can define a preference mechanism similar to the Local Preference attribute in BGP: an AS can check if the destination coordinate is one of its direct or indirect customers, or one of the customers of one of its peers, by verifying if a given coordinate pertains to a given region.

In our greedy routing scheme ASes do not continuously exchange messages concerning the routing protocol. Only configuration messages are rarely exchanged between pairs of ASes. In addition, the amount of data that each AS has to maintain is in the order of the number of provider-customer hierarchies times the number of neighbours.

Since our greedy routing scheme allows an AS to be reached via alternative paths from the core, when available, it is possible to control inbound load-balancing using the mapping system. LISP [5] has a mechanism which allows to perform inbound load-balancing as well as route differentiation in the network edge, instead of including it in the routing scheme. Furthermore, monitoring of link faults should also be done at the edge. Several links can be used in parallel in order to improve fault tolerance [7].

In order to prevent re-computation of coordinates when adding new customers, free spaces can be left in regions at different levels. The number of free spaces needed can be measure based on the evolution of the AS relationships in the last years [27].

Finally, in order to complement the greedy routing scheme we need to define an enhanced distance function which would lead to shortest paths for routing within the same hierarchy. Moreover, the main follow up work is the definition of a new architecture for interdomain routing which comprises traffic engineering, mobility and security mechanisms, while using greedy routing to perform interdomain routing.

References

[1] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *SIGCOMM*, 1999, pp. 251–262.

[2] V. Schrieck, P. Francois, C. Pelsser, and O. Bonaventure, "Preventing the unnecessary propagation of bgp withdraws," in *NETWORKING '09*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 495–508.

[3] A. Lambert, M.-O. Buob, and S. Uhlig, "Improving internet-wide routing protocols convergence with mrpc timers," in *CoNEXT '09*. New York, NY, USA: ACM, 2009, pp. 325–336.

[4] D. Pei, M. Azuma, D. Massey, and L. Zhang, "Bgp-rcn: improving bgp convergence through root cause notification," *Comput. Netw. ISDN Syst.*, vol. 48, no. 2, pp. 175–194, 205.

[5] I. N. R. Group, "Locator/id separation protocol (lisp," draft-farinacci-list-09.txt, October 2008.

[6] A. Feldmann, L. Cittadini, W. Mühlbauer, R. Bush, and O. Maennel, "Hair: hierarchical architecture for internet routing," in *ReArch '09*. New York, NY, USA: ACM, 2009, pp. 43–48.

[7] D. Zhu, M. Gritter, and D. R. Cheriton, "Feedback based routing," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 1, pp. 71–76, 2003.

[8] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica, "Hlp: a next generation inter-domain routing protocol," in *SIGCOMM '05*. New York, NY, USA: ACM, 2005, pp. 13–24.

[9] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *SIGCOMM '09*. New York, NY, USA: ACM, 2009, pp. 111–122.

[10] X. Yang, D. Clark, and A. W. Berger, "Nira: a new inter-domain routing architecture," *IEEE/ACM Trans. Netw.*, vol. 15, no. 4, pp. 775–788, 2007.

[11] J. Kleinberg, "Navigation in a small world," *Nature*, vol. 406, p. 845, 2000.

[12] J. Travers, S. Milgram, J. Travers, and S. Milgram, "An experimental study of the small world problem," *Sociometry*, vol. 32, pp. 425–443, 1969.

[13] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, pp. 167–256, 2003.

[14] I. Norros and H. Reittu, "Network models with a 'soft hierarchy': a random graph construction with loglog scalability," *Network, IEEE*, vol. 22, no. 2, pp. 40–46, 2008.

[15] R. Albert, H. Jeong, and A.-L. Barabasi, "Error and attack tolerance of complex networks," *Nature*, vol. 406, no. 6794, pp. 378–382, 2000.

[16] CAIDA, "As relationships," http://www.caida.org/data/active/as-relationships/index.xml.

[17] D. D. S. Price, "A general theory of bibliometric and other cumulative advantage processes," *Journal of the American Society for Information Science*, pp. 292–306, 1976.

[18] M. Boguna, D. Krioukov, and K. C. Claffy, "Navigability of complex networks," *Nat Phys*, vol. 5, no. 1, pp. 74–80, 2009.

[19] H. Zheng, E. K. Lua, M. Pias, and T. G. Griffin, "Internet routing policies and round-trip-times," in *In PAM*, 2005.

[20] P. Rodrigues, "Scalable interdomain routing," *MSc Dissertation, Universidade Nova de Lisboa*, 2010.

[21] P. Bose, P. Morin, I. Stojmenović, and J. Urrutia, "Routing with guaranteed delivery in ad hoc wireless networks," in *DIALM '99*. New York, NY, USA: ACM, 1999, pp. 48–55.

[22] B. Karp and H. T. Kung, "Gpsr: greedy perimeter stateless routing for wireless networks," in *MobiCom '00*. New York, NY, USA: ACM, 2000, pp. 243–254.

[23] A. Cvetkovski and M.Crovella, "Hyperbolic embedding and routing for dynamic graphs," *INFOCOM*, 2009.

[24] R. Kleinberg, "Geographic routing using hyperbolic space," in *IEEE INFOCOM 2007*. IEEE, May 2007, pp. 1902–1909.

[25] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A decentralized network coordinate system," in *In SIGCOMM*, 2004, pp. 15–26.

[26] T. S. E. Ng and H. Zhang, "Predicting internet network distance with coordinates-based approaches," in *In INFOCOM*, 2001, pp. 170–179.

[27] S. Shakkottai, M. Fomenkov, D. Krioukov, R. Koga, and kc claffy, "Evolution of the internet as-level ecosystem," 2006. [Online]. Available: http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0608058