B1 - Scientific Context and Objectives

The project ALADDIN deals with the fundamental aspects of large interaction networks enabling massive distributed storage, efficient decentralized information retrieval, quick inter-user exchanges, and/or rapid information dissemination. The project is mostly oriented towards the design and analysis of algorithms for these (logical) networks, by taking into account properties inherent to the underlying infrastructures upon which they are built. The infrastructures and/or overlays considered in this project are selected from different contexts, including communication networks (from Internet to sensor networks), and societal networks (from the Web to P2P networks).

During the last decade, lots of progresses have been made concerning the understanding of the structural properties of the above mentioned logical or physical networks. In particular, one can mention at least two major achievements of the structural analysis of these networks : (1) the observation that several of them share common statistical properties, including *small world* properties and *scale free* properties, to mention just a few, and (2) the design of structural models capturing these statistical properties. These achievements enabled a better understanding of common behaviors shared by these networks, mostly in terms of dissemination (e.g., rumor or virus spreading), vaccination, and fault tolerance. On the other hand, little has been done for exploiting these properties in order to design specific algorithms solving relevant problems. (A notable exception is information retrieval using highest-degree-first lookup schemes). Moreover, the aforementioned properties are mostly *quantitative* (e.g., clustering coefficient, power of the degree distribution, etc.) and suffer from possible biases of the experimental measurements.

The main objective of this project is to design algorithms and to analyze algorithmic phenomena in the broad context of interaction networks, and hence to make a breakthrough in this field. The algorithms will be based on *qualitative* properties, considered *from the nature and/or functionality of the networks*, such as bounded growth rate, low doubling dimension, minor excluding, or hyperbolic metrics.

Typical problems that will be tackled by the participants of this project are motivated either for their fundamental interests in the field of distributed information systems, or for their practical interests in the considered field of applications (e.g., for large content delivery networks). These problems include (but are not limited to) the search for ”good” overlay networks (e.g., multiplicative or additive spanners, Steiner trees, pre-defined topologies, etc.), the design of informative labeling (e.g., distance, adjacency, etc.), the design of protocols for mobile entities (e.g., exploration, searching, etc.), as well as classical networking and navigability problems such as routing and closest server finding, and fundamental algorithmic problems such as TSP and k-center.

The originality of our approach is based on the following methodology and concerns :

1. Several of our hypotheses regarding the networks will be built upon the analysis of the *nature* of the networks (i.e., its structural context, from computer science to sociology) and the analysis of the *functionality* and *usage* of the networks (i.e., what it is made for). It has indeed been recently observed that from this nature and functionality could *emerge* general properties that can be used in turn for algorithm design and analysis. The interest of this approach is two folded. First, it does not suffer from any bias of the networks experimental measurements ; Second, the emerging properties are of qualitative flavor (e.g., the network has a bounded doubling dimension), and are thus very malleable. Part of our investigations will be dedicated to the analysis of new emerging properties from the nature and functionality of the considered networks.

2. Our assumptions regarding the networks are very novel. For instance, the notion of doubling dimension has been only very recently introduced with the objective of tacking hard problems such as TSP (the most recent achievements using this notion are restricted to few problems only, including routing). Notions like *minor excluding*, or *hyperbolic metrics* have been seldom considered for the design and analysis of network algorithms, whereas it is strongly believed that they are critical in this context. The main advantages of the aforementioned notions are their generality and their power. They apply to a large class of networks, and they enable designing efficient solutions to hard problems. (For instance, *any* network with bounded doubling dimension is navigable, i.e., greedy routing performs in poly-logarithmic number of steps).
3. We will tackle some network problems in contexts in which the network is only implicitly and/or partially known. The concept of *implicit knowledge* aims at capturing the fact that the amount of knowledge each node of a distributed network has about any aspect of the topology (from the full knowledge of the whole structure or global information such as number of nodes, to local information such as neighborhood topology) is inherently limited. Implicit knowledge is modeled by an *oracle*, whose complexity is measured in term of the amount of information (i.e., the number of bits) given to the entities or nodes composing the networks. The concept of *partial knowledge* aims at capturing the fact that nodes do not need to know all their connections. This is modeled by the *probe-complexity* of the problems, i.e., the minimum amount of probing that nodes must performs to solve a problem.

The topics addressed by this project requires competences from several scientific fields of computer science, including distributed computing, algorithms design and analysis, and various algorithmic aspects of graph theory. The participants to the project are recognized actors in these fields, at the international level. They will gather their expertise to fulfill the objectives of the project. They will take advantage of previous collaborations in projects granted by the French Ministry of Research, in particular the project PairAPair of the ACI ”Masses de données”, which ended last summer.

This new project aims at going much further and broader in the investigation of network algorithm design and analysis than what the project participants had the opportunity to do in the past. The ambition of the project is to investigate the very fundamental aspects of problems lying at the heart of a large number of applications, from P2P file sharing system design to sociological phenomenon analysis, and from Internet algorithms to logical interaction networks. It is expected that the investigations carried on by the participants of the project will enable deeper understanding of a large spectrum of problems arising in the field of content delivery networks. In particular :

– The design of efficient algorithms for various types of generic problems will provide as many insight for the future design of protocols dedicated to specific related problems.
– Universal lower bounds or impossibility results will provide as many insight on the limit of what can be achieved by interaction networks.

To make a long story short, this project is aiming at participating to the establishment of the fundamental bases of the interaction networks, in connection with what is currently investigated in the field of complex networks, and in complement with what is currently done by several communities such as physicists (statistical physics) and network engineering (metrology).

B2 - Description of Tasks

B2.1 - General description

As said in Section B1, ALADDIN deals with the fundamental aspects of large interaction networks, including P2P file sharing systems, content-delivery networks, sensor networks, distributed data bases, the Web, etc. The project is mostly oriented towards the design and analysis of algorithms for these networks, including algorithms for information retrieval, publish and search, routing, overlay network design, data gathering, distributed location, informative labeling, etc. One of the main characteristics of the project is that it aims at taking into account properties inherent to the underlying infrastructures upon which the (logical) content delivery networks are built. The infrastructures and/or overlays considered in this project are selected from different contexts, including communication networks (from Internet to sensor networks) and societal networks (from the Web to P2P networks), and the algorithms will be based on qualitative properties, considered from the nature and/or functionality of the networks, such as bounded growth rate, low doubling dimension, minor excluding, or hyperbolic metrics.

To achieve its objective, the project will investigate three main topics :

Topic 1. Impact of the nature and/or functionality of a network on its structural features ;

Topic 2. Algorithm design for interaction networks with specific properties (e.g., satisfying the specific structural features identified in Topic 1) ;

Topic 3. Distributed algorithms design for limited knowledge and/or limited probing capacity.

Topic 1 is dealing with *emerging* properties, precisely on the analysis of the structural characteristics of networks that are (directly or indirectly) caused by their nature or functionalities. A large part of physics consists in studying how constraints shape the nature (a typical example of such constraints is energy minimization). We believe that a similar kind of study must be conducted in computer science and networking. Indeed, the nature of a network (e.g., geographically constrained for wired networks, uniformly spread over a field for sensors, etc.) or its functionality (e.g., navigability for societal networks, content-addressable for P2P networks, etc.) has a strong impact on its structure. Topic 1 aims at identifying high level features that can be used for the design of efficient algorithms. As mentioned before, we will mostly focus on qualitative properties, in complement to all what has already been done in network metrology.

Observe that although Topic 1 may seem to be of a rather theoretical flavor, its practical impact in term of algorithm design (and therefore, in term of protocol design) can be significant. Indeed, in the framework of network algorithms the worst case analysis is often too pessimistic because worst cases are often very specific and do not occur in practice. On the other hand, the (uniform) average case analysis is often too optimistic because networks often present characteristics for which the average is not representative (e.g., the power law degree distribution of interaction networks is not captured by the $\mathcal{G}_{n,p}$ random graph model). A way to pass over this dilemma is to perform average case analysis using non uniform network distributions. However, a difficulty of this approach is the design of the appropriate distribution. We believe that a probably more fruitful approach is to identify classes of networks (possibly characterized by very abstract properties) for which one has strong evidence that the network under investigation fits in. Then, algorithms designed for these specific classes will definitely be efficient in practice, even if these classes appear to be of rather theoretical flavor (such as bounded doubling dimension networks, bounded growth networks, minor excluding networks, hyperbolic networks, etc.).

Topic 2 is considering the classes of networks identified in Topic 1, and is aiming at designing algorithms for these classes. Typical problems that will be tackled by the participants of the project include (but are not limited to) the search for "good" overlay networks (e.g., multiplicative or additive spanners, Steiner trees, pre-defined topologies, etc.), the design of informative labeling (e.g., distance, adjacency, etc.), the design of protocols for mobile entities (e.g., exploration, searching, etc.), as well as classical networking and navigability problems such as routing and closest server finding, and fundamental algorithmic problems such as TSP and k-center.

Finally, Topic 3 will revisit the problems under investigation in Topic 2 in contexts in which the network is only implicitly and/or partially known. As mentioned before, the concept of implicit knowledge aims at capturing the fact that the amount of knowledge the nodes of a distributed network have about any aspect of the topology is inherently limited, and the concept of partial knowledge aims at capturing the fact that nodes do not need to know all their connections. Our main objective will be to design algorithms requiring a limited amount of knowledge on the networks structure. One tool for this framework is the recent notion of advice, whose complexity is measured in term of amount of information (i.e., the number of bits) given to the entities or nodes composing the networks. Another tool is the notion probe-complexity, measuring the minimum amount of queries that nodes must performs to solve a problem. Last but not least, we will also pay a specific attention to the design of informative labeling schemes, that are compact distributed data structures enabling answering queries (e.g., distance) by considering solely the labels of the nodes involved in the query.

B2.2 - State of the Art and Roles of the Partners

A detailed state of the art of each task that ALADDIN will carry out will be given in Sections B2.3 where each task is described. Globally, as we had already the opportunity to say, the assumptions investigated in Topic 1 regarding the networks are very novel. For instance, the notion of doubling dimension has been only very recently introduced with the objective of tacking hard problems such as TSP (Klein, FOCS 2005). Notions like *minor excluding*, or *hyperbolic metrics* have been seldom considered for the design and analysis of network algorithms, whereas it is strongly believed that they are critical in this context (Abraham, Gavoille, 2006). The main advantage of the aforementioned notions is their generality and their power. They apply to a large class of networks, and they enable designing efficient solutions to hard problems.

The problems investigated in Topics 2 and 3 are crucial in the context of interaction networks. To mention just a few, the design of light spanners in the context of arbitrary networks is a hot topic (e.g., Bollobás, Coppersmith, Elkin, 2005), but much remains to be done for extending this theory to networks satisfying specific features (doubling dimension, etc.). The design of efficient informative labeling schemes finds direct applications, e.g., in the context of XML (Abiteboul et al., 2001). Finally, the notion of an oracle providing limited information about the network to its nodes will enable to generalize lower bounds and impossibility results such as the ones in (Lynch, 1989 ; Fich, Rupper, 2000).

The participants of the ALADDIN project are among the main actors in the aforementioned scientific fields, with a recognized international expertise. They recently chaired or participated to the program committees of the major conferences dealing with fundamental aspects of network computing (PODC, DISC, ICDCS, etc.). Non anecdotally, their contribution to network navigability and to small worlds are cited several times in the talk of Jon Kleinberg (Cornell University) given at the International Congress of Mathematicians, Madrid, August 22-30, 2006, where he was honored by the prestigious Nevanlinna price. Jon Kleinberg was the initiator of the formal analysis of Milgram's experiment that demonstrated the existence of short chains of acquaintances between individuals, and his influencing publications in the conference STOC and the journal Nature in year 2000 are seminal.

In fact, most of the participants to the project ALADDIN developed their expertise in the framework of three ending projects granted by the ACIs ”masses de données”, ”sécurité”, and ”ATIP CNRS” : PairAPair, FRAGILE, and LOCOGLOBO. These three projects enabled their participants coming from different fields (graph theory, networking, algorithmic, etc.) to share their knowledge, hence enabling them to become active contributors to some topics at the heart of interaction networks, specifically P2P file sharing systems, sensor networks, and ad hoc networks. ALADDIN aims at pushing much further the study of interaction networks by capitalizing on the expertise gained by the participants during their carry out of the three aforementioned projects.

B2.3 - Tasks and milestones

Here, we summarize the tasks corresponding of the three aforementioned topics. A complete description of these tasks can be found in the Appendix, Section B2.7.

Note that ALADDIN is part of a very ambitious research framework whose long-term objectives are far beyond the duration of a 4-year project. What ALADDIN aims at doing is instantiating some problems, and, for each instance, focussing on the development of solutions for it. ALADDIN is thus organized into a set of (not necessarily independent) tasks that are at least partially solvable within four years.

Topic 1. Impact of the nature and/or functionality of a network on its structural features

Several of our hypotheses regarding the networks will be built upon the analysis of the *nature* of the networks (i.e., its structural context, from computer science to sociology) and the analysis of the *functionality* of the networks (i.e., what it is made for). It has indeed be recently observed that from this nature and functionality could *emerge* general properties that can be used in turn for algorithm design and analysis. The interest of this approach is two folded. First, it does not suffer from any bias of the networks experimental measurements ; Second, the emerging properties are of qualitative flavor (e.g., the network has a bounded doubling dimension), and are thus very malleable. Part of our investigations will be dedicated to the search for new emerging properties from the nature and functionality of the considered networks.

Task 1.1. Doubling Dimension and Graph Metrics (Leader : Emmanuelle Lebhar). Our objective is to explore how geometric parameters interacts with each other and with the algorithms running on the network.

Task 1.2. Minor Exclusion (Leader : Ioan Todinca). The main objective of this task is to determine which type of networks among those appearing as model for interaction networks can be characterized by the exclusion of some particular minor. (For instance, a search tree is a K_3-free network).

Task 1.3. Geographical embedding (Leader : Nicolas Bonichon). A first objective in this task is to characterize the properties of the topologies induced by the power efficient schemes in sensor networks. A second objective is to study the Internet latencies and how they fit with doubling metrics characteristics.

Task 1.4. Structural properties of P2P networks vs. their functionalities (Leader : Fabien de Montgolfier). This task is two-folded. A first objective resides in studying the properties of the structure of existing P2P networks. In particular, affinity based models is a recent domain, and is a promising approach for giving more theoretical insight in the structure of the resulting networks. A second objective will focus on the structure of data based networks, including those using skiplist-like logical structure for connecting the data.

Task 1.5 Algorithms Structuring Networks (Leader : Nicolas Schabanel). Our objective is to explore how the analysis of the performances of the algorithms running on interaction networks enable understanding the structures of these networks. A first topic of study is the navigability property which has a lot of applications in information retrieval.

Topic 2. Algorithm design for interaction networks with specific properties

Our assumptions considered in Topic 1 regarding the networks are very novel. For instance, the notion of doubling dimension has been only very recently introduced with the objective of tacking hard problems such as TSP (the most recent achievements using this notion are restricted to few problems only, including routing). Notions like *minor excluding*, or *hyperbolic metrics* have been seldom considered for the design and analysis of network algorithms, whereas it is strongly believed that they are critical in this context. The main advantages of the aforementioned notions are their generality and their power. They apply to a large class of networks, and they enable designing efficient solutions to hard problems. The ambition of the project is to investigate the very fundamental aspects of problems lying at the heart of a large number of applications, from P2P file sharing system design to sociological phenomenon analysis, and from Internet algorithms to logical interaction networks. It is expected that the investigations carried on by the participants of the project will enable deeper understanding of a large spectrum of problems arising in the field of interaction networks. In particular : (1) the design of efficient algorithms for various types of generic problems will provide as many insight for the future design of protocols dedicated to specific related problems, and (2) universal lower bounds or impossibility results will provide as many insights on the limit of what can be achieved by interaction networks.

Task 2.1. Design of overlay networks and spanners (Leader : Cyril Gavoille). A first objective of this task is to design efficient static solution for light spanner construction in graphs having specific properties like bounded doubling dimension or excluding a fixed minor (keeping in mind that TSP is a particularly chalenging problem for these two properties (cf. Talwar, STOC 2004 ; Klein, STOC 2005). The second objective of this task is to determine whether the "good" static spanners can be used to design light overlay networks in a dynamic environment in which nodes join and leave at will (as in, e.g., P2P systems), or in which nodes are moving (e.g., ad hoc networks).

Task 2.2. Nearest server and k-center problems (Leader : Ralf Klasing). The objective of this task is to revisit the nearest server problem and the k-center problem in network models such as the ones derived from the analysis carried on in Topic 1 (doubling dimension, minor excluding, etc.).

Task 2.3. Impact of correlation in probabilistic models (Leader : Nicolas Hanusse). Correlations are known to be difficult to tackle in probalistic models. For instance, it is still an open problem to generate a random network with a given clustering coefficient (or with a given number of triangle). Our goal is to study different situations where dealing with correlations can improve quality of service. A primary objective consists in improving services in content delivery networks by taking into account correlations between requests.

Task 2.4. Exploration and searching (Leader : Pierre Fraigniaud). The first objective of this task is to define a setting for the exploration and search problem, so that they fit with the framework of information retrieval. Once this setting will have been defined, both problems will be investigated in a distributed and dynamic context. We will mostly concentrate our work on the minimum amount of information to be given to the entities or to the nodes so that both problems can be performed rapidly[1].

Topic 3. Distributed algorithms with limited knowledge and/or limited probing capacity

For many network problems (such as exploration, information retrieval, spanners, wakeup, broadcast, etc.), the quality of the algorithmic solutions often depends on the amount of ”knowledge” given to nodes of the network about their environment, or on the ability of the nodes to retrieve information about this environment. The impact of knowledge concerning the environment is significant in many areas of distributed computing, and hundreds of impossibility results and lower bounds for distributed computing are depending on whether or not the nodes are provided with partial knowledge of the topology of the network. Topic 3 deals with the design of distributed algorithms with limited knowledge and/or limited probing capacity.

We will mostly focus on three types of limitations and uncertainty : (1) bounded size information labeling, (2) bounded size oracle, and (3) limited probing capacity.

Information labeling is a method for network representation (Peleg, 2000). The goal is to cheaply store at nodes useful information about the network and make it readily and conveniently accessible. For this purpose, nodes are given labels (aimed to be short) so that information about them (e.g., inter-node distance, adjacency, similarity, etc.) can be efficiently computed based only on their labels.

Advising is also a method for network representation, quite similar to information labeling (Fraigniaud, Ilcinkas, Pelc, PODC 2006). Aware of global information about the environment, oracles provide information (aimed to be compact) to the nodes, or alternatively to mobile entities moving in the network, so that they can perform sophisticated tasks (e.g., broadcasting, wakeup, exploration, etc.) efficiently.

The two aforementioned models are ways to measure the minimal amount of knowledge that should be given to the nodes so that information about the network can be locally retrieved (information labeling) or complex tasks can be distributively computed efficiently (advising). The last model that we investigate in the framework of Topic 3 measures the ability that the nodes possess to learn about their environment.

Limited probing constraints the nodes to perform a limited amount of probes (typically the number of pings for IP) about their surrounding. Network protocols achieving their tasks using a limited number of pings are susceptible to work well in a rapidly changing environment.

Task 3.1. Design of informative labeling schemes (Leader : Cyril Gavoille). The goal of this task is to construct informative labeling schemes for queries such as distance, routing,

[1]Note that we may investigate these two problems for specific classes of networks such as the ones identified in the tasks of Topic 1, but at the time this proposal is edited, we believe that it is beyond the scope of ALADDIN.

ancestry, etc., in network topologies with specific properties (low doubling dimension and excluding a fixed minor for instance). An extension to a dynamic setting will be considered.

Task 3.2. Design of bounded size oracle (Leader : Pierre Fraigniaud). Task 3.2 will focus on the design of small oracles for solving problems efficiently, such as sparse spanner construction, small (connected) dominating sets for sensors, and overlay network design for, e.g., P2P systems. In the framework of this task, we will also tackle problems in the context of mobile computing (cf. Task 2.4), specifically cooperative network exploration and searching. This task will be achieved in the context of implicitly described networks (i.e., networks with small doubling dimension, or excluding a minor).

Task 3.3. Protocol with limited probing capacity (Leader : Laurent Viennot). The project aims at designing algorithms assuming plausible structural features, as the ones identified in Topic 1. A basic problem in this context consist in finding nearest neighbors in decentralized P2P networks. We will be mostly interested in establishing tradeoffs between the efficiency of the algorithm and the number of probes. In particular, we will be focussing on the minimum number of probes that are required to find close neighbors. The quality of the solution in this context is measured by the stretch factor of the computed solution, i.e., by analyzing the worst case distance of the neighbors selected by the algorithm compared to the nearest neighbor. More generally, this task will focus on finding algorithms using a limited amount of probes, and on deriving lower bounds on the probing cost necessary to solve basic problems. Another objective of this task is to understand the connections between the dynamic of the network and the probing cost of protocols designed for this network.

Task 3.4. Sub-topologies preserving structuring properties (Leader : Olivier Beaumont). This task aims at characterizing sub-topologies constructible in a distributed manner so that to preserve bi-connectivity (and hence fault tolerance). In the specific case of the construction of sub-topologies ensuring message transmissions along short routes, this task intersects with Task 2.1. Nevertheless, the two tasks remains different enough because, in this task, nodes can also use their local links as opposed to the spanner theory which restricts the set of used links to belong to the spanner. We also plan to consider the problem in a broader context. For instance, data or nodes in a P2P system are often connected via a structure depending on some proximity measure between the data, or on some affinities between the nodes. This task will therefore also consider the search for optimized spanners such that each node can perform efficient tasks based solely on the spanner *and* its local links.

B2.4 - Project management

This ambitious project is planed to run over four years. It is important to note that most of the partners involved in ALADDIN had the opportunity to collaborate in the past while they were involved in various more narrow national and European projects. Therefore there will be no need of a warmup, and the project will start immediately after its notification.

The success of the project is based on an intensive collaboration between the two sites, involving frequent meetings and inter-site visits. ALADDIN is therefore planing four 2-day meetings every year during which all partners will present their most recent results related to the project. One of these meetings will be dedicated to the compilation of the results achieved during the year, to make sure that all partners synchronize. All dates for deliverables will fit with the dates of these annual special meetings.

In addition, our budget requests a grant for a 1-week visit per year per permanent member (that can be used by the member himself or by his students). The chair of the project will pay a specific attention that frequent exchanges between the partners are done.

The management of the web page of the project will play an important role. Indeed, the fundamental nature of the project requires that the results will be made available to the community as quickly as possible. Our web page will be a repository for preliminary versions of papers by the partners, enabling rapid diffusion of our preliminary results inside the project (via a private web site), and a rapid advertising of the finalised results outside the project (via a public web site).

B2.5 - Results and evaluation criteria

Again, the objectives of ALADDIN are :

1. Understanding the impact of the nature and/or functionality of a network on its structural features ;
2. Designing algorithms for interaction networks satisfying the specific structural features identified in item 1 ;
3. Designing distributed algorithms in the context of limited knowledge and/or limited probing capacity.

As far as the first objective is concerned, the quality of our contribution can be measured by our ability or inability to identify large classes of networks (including all current types of content delivery networks) fitting in at least one of the categories ”bounded doubling dimension networks”, ”minor excluding networks”, etc. Alternatively, we would be fully satisfied by results stating that all but a small fraction of the nodes fit in such categories. Actually, even probabilistic results, valid for a fraction of nodes would probably be satisfactory.

In fact, the success of the first objective cannot be evaluated independently from our second objective. Indeed, a class $\mathcal{C}$ of networks is ”interesting” if one can design specific algorithms for this class, that can be proved to be more efficient than algorithms designed for arbitrary networks. Probably, for enabling easy design and analysis of algorithms, this class $\mathcal{C}$ must be malleable and simple to define, but the main criterion remains the ability to produce informative results about networks in this class, and especially specific efficient algorithms for them.

The quality of the contributions of ALADDIN corresponding to the third objective are mostly function of the ability of the distributed algorithms to be turned into network protocols. Indeed, these constrained algorithms are aimed at being implemented on large networks, and thus they must satisfy several constraints including, as far as this project is concerned : locality (nodes must exchange a limited amount of information between their neighbors), and limited knowledge (nodes must act under the constraint that they are only partially aware of the network topology).

ALADDIN aims at deriving lower bounds and impossibility results. In general such results are very informative because they define limits to models. However, they are really informative if one has evidence that ”real instances” (i.e., instances that correspond to real applications) are in the class of problem that are proved intractable. For instance, it is known that shortest path routing requires routing tables of size $\Omega(n \log n)$ for n-node networks (Gavoille, Perennes, 1994). It is however unclear whether this result holds when restricted to the class of networks with bounded growth, or networks with a power law distribution of the degrees. ALADDIN will thus pay a specific attention to the fact that all the impossibility results derived during the execution of the project must correspond to networks that are used in practice for distributed massive storage.

Last but not least, since the project is oriented toward fundamental research, the quality of the project will also be evaluated by the ability of the partners to present and advertise their results in the most appropriate and prestigious conferences of the field.

B2.6 - Perspectives

ALADDIN is aiming at participating to the establishment of the fundamental bases of the interaction networks theory, in connection with what is currently investigated in the field of complex networks, and in complement with what is currently done by several communities such as physicists (statistical physics), and network engineering (metrology). In particular, several techniques (and possibly results) that will be derived by the partners during their investigations in ALADDIN will probably be extendable to broader contexts, including other types of interaction networks appearing in the framework of molecular interactions, linguistic, etc. One important perspective of this project is to carry on the researches described in this document in collaboration with partners from physics, biology, and social sciences. We are confident in this perspective by noticing that several partners in ALADDIN had already the opportunity to successfully interact with searchers from other disciplines through their participation to common events. Cf., e.g., the *Ecole thématique sur les grands réseaux d'interaction* (Paris, April 2005), the *Summer School on Complex Systems, Theoretical Informatics, and Systems Biology* (Valapraiso, January 2005), as well as the project "Carpelle Virtuel" involving biologists and some members of ALADDIN for the description of the transcriptional network of the *Arabidopsis thaliana*.

B2.7 - Appendix (to be read at the referee's discretion)

This section precisely describes the context and state of the art of each of the tasks summarized in Section B2.3.

Context of Task 1.1. Doubling Dimension and Graph Metrics

The geometry of a network, especially the growth of its ball sizes, has a high impact on the design of efficient algorithm. Precisely, it has recently been observed that, under the assumption that a network has a low *doubling dimension*, efficient distance labeling and routing schemes can be designed. The doubling dimension of a network is the smallest α such that any ball of radius $2r$ can be covered by at most 2^α balls of radius r, for any $r > 0$. This reflects how close to a constant dimension ℓ_p metric the shortest paths metric of the network is. The key of efficiency of algorithms designed on networks of low doubling dimension is the possibility to decompose them hierarchically into nested clusters of exponentially decreasing size, while guaranteeing a constant number of clusters in each level of the hierarchy.

In parallel, it has been observed empirically that numerous large real networks present a high local clustering, i.e., the nodes have a high proportion of their neighbors that are neighbors of each other. This property is clearly related to the local growth of the network, and thus to its dimensionality. One goal of task 1.1 is to relate this observation to the existing combinatorial measures of dimensionality (doubling dimension and ball growth). A first step for this investigation consists in trying to generate random graph models presenting a fixed number of triangles. The number of triangles is indeed one formal way to measure the local clustering (Watts and Strogatz 98), however it is unsatisfactory with respect to several aspects (one can build graphs with a high number of triangles and still a majority of nodes having no connected neighbors). A second goal of task 1.1 is to propose an alternative formal measure of the local clustering that will capture the aspects missed by the number of triangles measure.

The second main parameter observed on large real networks which is strongly related to their geometry is their power-law degree distribution. It has been shown that such a distribution could be produced by a dynamic construction of the network based on the preferential attachment principle : new nodes get preferentially attached to the highest degree nodes (Albert and Barabasi 99). Other contributions (Carlson, Doyle, 1999 ; Fabrikant, Koutsoupias, Papadimitriou, 2002 ; Berger, Borgs, Chayes, D'Souza, Kleinberg, 2004) extend this approach by showing that other trade-off optimization processes yield also power laws. It is then likely that power law-like structures are strongly related to the underlying engineering optimization processes. This is a typical example of algorithms structuring a network.

Relating geometrical properties observed empirically to specific metric classes could provide efficient tools to design algorithms on large graphs with incomplete information of their global structure, in particular thanks to the approximation of a network by a well known metric through randomized embeddings.

Context of Task 1.2. Minor Exclusion

Network topologies can be classified from various qualitative properties. Among them, one can cite : hierarchical networks (tree-like topologies), backbone structures (outer-planar or series-parallel networks), two dimensional maps or geographic based networks (planar or bounded genus topologies). The nature of such properties are very different, and this is why solutions for problems on trees and on meshes fundamentally differ so much : Divide-and-Conquer techniques are well suited for trees because a small number of nodes suffice to halve the topologies, whereas node labeling based solutions are better suited for meshes by exploiting the fact that a coordinate systems for nodes can reflect the plane embedding of the topology.

Nevertheless, these two extreme topologies (trees and meshes) share a common fundamental graph property : both exclude a fixed minor. A graph G exclude a minor H if H cannot be obtained from G by edge contractions, or by edge/node deletions. The theory of graphs excluding a fixed minor has led to very important and deep results. For instance, the Graph Minor Theorem (cf. Robertson and Seymour series) states that every infinite family of graphs that is close under minor can be characterized in terms of graphs excluding a *finite* list of finite *minors*. For instance, trees exclude K_3 (a clique with three nodes), outer-planar networks exclude K_4 and $K_{2,3}$ (a complete bipartite graph with parts of size two and three), series-parallel graphs exclude K_4, planar graphs exclude K_5 and $K_{3,3}$, and so on. One can also cite the non-trivial property that the graphs excluding a planar minor have a tree-like structure (bounded treewidth).

Context of Task 1.3. Geographical embedding

Networks are often physically embedded in the plane or the sphere. Ad hoc networks and sensor networks for example exhibit a strong correlation between the physical proximity of the nodes and their probability to be connected. For instance, using a simple open field radio propagation model, the network forms a unit disk graph where two nodes are connected when their distance is less than one. In academic simulations, uniform distribution of nodes in a rectangular field are often used. However, many practical scenarios are far from uniformity (e.g., spreading of a population over a territory).

In wireless networks (typically sensor networks as far as this project is concerned), specific connection structures can be constructed when considering energy efficient schemes (gathering, etc.). This is particularly the case for sensor networks where energy is the scarcest resource of the system. The Gabriel graph is typical of such structures spanning the original networks : two nodes u and v of the Gabriel graph are connected if no other node lies in the disk with diameter uv. This spanner is well suited when assuming a simple open field radio propagation model where the power necessary to reach a node at distance d is proportional to d^α, $\alpha > 2$. More practical and energy-efficient constructions increases the power transmission of each node until some local criterion is met.

On the other hand, the Internet is embedded on the Earth and is hence physically embedded in a sphere. Recent works (e.g., Slivkins, 2005) tried to model Internet latencies according to the inter-node distances in a virtual space (e.g., an Euclidean or an hyperbolic space with low dimension). The basic idea is to assign virtual coordinates to nodes and to estimate the latency between two nodes based on their coordinates. Surprisingly the spheric nature of the physical embedding is ignored in such work. Interestingly though, numerous traces of latency measurements can be used to study the propagation delay metric in the Internet (pointing out problems in the frequent violation of the triangular inequality when considering latencies as metric distances). If nodes are uniformly distributed in a geographical space, then bounded growth metrics are good candidates for modeling distances between nodes. When the distribution is not uniform, doubling metrics appear more appropriate.

Context of Task 1.4. Structural properties of P2P networks vs. their functionalities

P2P networks[2] are highly distributed content distribution networks. A large number of nodes share resources (text, audio, video files), and each node is in charge of selecting other nodes to connect with in order to create an overlay network supporting the search, publish, and sometimes the transfer procedures. A straightforward method for achieving this consists in connecting to random nodes. For example, Gnutella nodes find neighbors by random walk. In BitTorrent, a tracker maintains a list of participants and advertises randomly chosen contacts to nodes. Interestingly, a node in BitTorrent uses only a subset of its connections based on the bandwidth rate of data exchanges with its neighbors. On the other hand, some protocols such as Kademlia (used in eDonkey and some extensions of BitTorrent) propose to organize the overlay network according to specific structures supporting efficient routing strategies. Kademlia uses prefix routing and has a pseudo-hypercube topology. Other solutions where proposed using, e.g., a d-dimensional tore (CAN), a ring with doubling shortcuts (Chord), a de Bruijn graph (D2B, Koorde, etc.), etc. The basic approach for constructing such structures consists in assigning random IDs to nodes, or using distributed hashing. Each node then selects its neighbors according to their IDs. For instance, to support prefix routing, each node connect to nodes whose IDs share increasingly longer common prefix with its ID. The resources are assigned to nodes using distributed hashing, creating a Distributed Hash Tables (DHT).

Depending on the P2P application deployed in the network, nodes get naturally connected to nodes they exchange data with. This exchange graph can then be used to enhance the overlay graph. In file sharing applications, this graph has properties similar to social networks (e.g., power-law degree distribution). More generally, the choice of the neighbors can be captured by affinity models where each node has a preference list. This latter approach is closely related to marriage and roommate theory. An open field of research resides in studying the structure induced by the preference lists obtained via a given affinity function. Classical affinity functions include bandwidth (neighbors with better bandwidth capacity are preferred), latency (neighbors with lower response delay are preferred), data complementarity (neighbors with more un-possessed data are preferred), etc. More complex affinity functions can be obtained by combining several criteria.

Moreover, recent methods for indexing highly distributed data propose to build an overlay network on the data itself. For example, SkipGraph considers totally ordered data elements. All elements are then searched through a skiplist structure. More generally, multi-field data items can be seen as points in a d-dimensional space. Nodes are then connected according to the proximity of their data items in this space. Some other types of data may exhibit different structures. For example ontology classified elements or surface mail addresses have a natural tree-like structure. Identifying and exploiting the structure induced by some classes of data is a promising track for the construction of highly distributed content network.

Context of Task 1.5 Algorithms Structuring Networks

A large part of physics consists in studying how constraints shape the nature. A typical example of such constraints is energy minimization. A similar situation exists in interaction networks where communication algorithms play the role of constraints. The choice of the content sharing protocol has indeed a huge impact on the structure of P2P networks. For instance,

[2]This project will not address the problems related to intellectual property rights, which are far beyond the scope of ALADDIN. We note however that the P2P file sharing paradigm can be used in complete respect of these intellectual property rights, and that this paradigm still fully deserves to be investigated as it is currently the most promising alternative to centralized systems, and it has a large number of potential applications in other fields, including routing in radio networks.

in the early days, the highly inefficient protocol Gnutella (based on systematic flooding) prevents the communities from getting larger than a thousand nodes, above that threshold the network crashed. Afterwards, other protocols were designed that structured the P2P networks (and traffic) differently. More generally the fact that a certain kind of algorithms runs on a given structure, puts constraints on that structure. A recent striking example is the navigability in social networks : the ”social feature” of such a network is that in spite of a huge number of individuals, each one can easily reach any other with a very partial information. Interestingly enough (Lebhar, Schabanel, 2005), this feature does not seem to be related with the famous power law property. Kleinberg (2000) (generalized by Slivkins, 2005) showed that only specific link distributions would allow navigability. It has also been recently shown (Fraigniaud, Lotker, Lebhar, 2006) that some network cannot support navigability, independently of the protocol chosen. It is likely that other important specific properties can be deduced on navigable networks that would allow to understand better how interaction networks (such as P2P networks) are structured and then to improve their performances.

Context of Task 2.1. Design of overlay networks and spanners

Determine a “good skeleton” of a topology and solve the problem on it is a very general approach for hard problems in Computer Science. Specifically, a k-*spanner* of a graph G is a spanning subgraph H approximating orginal distances : $dist_H(x, y) \leq k \cdot dist_G(x, y)$ for any pair x, y of nodes of G. The general goal is to construct efficiently “light” k-spanners, that is a k-spanner with the smallest number of edges, or more generally having a total edge-weight as small as possible (for edge-weighted topologies the total edge-weight is the sum over all edges of the spanner of the edge weights). There is obviously a trade-off between the lightness of the spanner and the approximate ratio k.

There are huge numbers of applications using light spanner techniques directly or indirectly. In distributed computing, a lot of tasks (including broadacsting, routing, distance estimation, etc.) have a cost directly depending on the number of edges on which the algorithm works (cf. Peleg’s SIAM monograph, 2000). There are also applications in Computational Geometry (cf. Chapter 7 of Handbook of Computational Geometry 2nd Ed.), and interestingly solving Traveling Salesman Problem (TSP) uses some suitable k-spanners. In particular, based on a $(1 + \epsilon)$-spanner of total weight $O(\mathrm{MST}/\epsilon)$ where MST denotes the total weight of a minimum spanning tree, Papadimitriou et al. (FOCS, 1995) showed that TSP on planar graphs can be $(1 + \epsilon)$-approximated in polynomial time for every fixed $\epsilon > 0$.

Context of Task 2.2. Nearest server and k-center problems

In large content delivery networks, data are often replicated to make sure they are permanently quickly available. Finding the closest copy of an item is thus an important field of investigation. A basic problem in that context is to find the nearest server among a set of given servers. This problem typically occurs when a client is trying to fetch some data from one web or ftp mirror. Since Internet is an important field of application for that problem, Task 2.2 will be considered in conjunction with Task 1.3, hence using models suited for Internet latencies. Existing solutions typically use virtual coordinates enabling the construction of distance labeling schemes. Other labeling schemes could however be deduced from real structural properties (cf. Topic 1) instead of approximating distances in a metric space.

The k-center problem is closely related to the nearest server problem. In the k-center problem, k servers have to be placed in a network so that every node in the network is within a given distance range from at least one server. This problem has several practical applications, in particular the placement of data collectors in sensor network for bounding the length of the routes along which data are collected.

Context of Task 2.3. Impact of correlation in probabilistic models

Since the discovery of power laws in the Internet router graph (Faloutsos et al., 1999), it is now clear that uniformly random graphs are not suitable models for real networks. This was naturally not a surprise but the important (and surprising) fact was that the power law structure drastically modifies the observed phenomena (such as virus spreading, routing algorithm, robustness to attack, etc.). Correlations (as opposed to uncorrelated uniform random structures) are thus an essential feature of real networks that significantly changes their behavior. Power laws networks are indeed an example of such correlations : nodes tend to connect to highly connected nodes rather than to arbitrary random nodes, which for instance implies that a virus can spread very fast in such a network even if the average number of links per node is only 2. Another example of correlations is the clustering coefficient : two neighbors of a node have an high probability to be connected with eachother. Correlations are also observed in traffic traces : from time to time, one can observe traffic bursts, i.e., a large number of requests are sent during a very sharp period of time after and before which traffic returns to normal. Correlations also naturally arise between the targets of the requests : for instance in P2P networks, people liking the same song are likely to have other tastes in common (and hence are likely to try to download them from the same sources). It has been recently shown that taking into account correlations between requests (Huan, Chen, 2004 and 2005 ; Dey, Schabanel, 2006) can considerably improve network quality of service.

Context of Task 2.4. Exploration and searching

This task will investigate fundamental aspects of traversing a network by a team of mobile agents (Alpern and Gal, 2003). According to the standard terminology, *exploration* is the task in which a set of mobile entities has to visit all nodes of a networks, while *searching* is the task in which a set of mobile entities has to capture an hostile and possibly harmful intruder in the network. In other worlds, exploration can be seen as a special case of searching in which the intruder is immobile. Although the two problems look very similar, they were studied by different communities, mostly because, as far as theoretical computer science is concerned, the former is related to complexity (NL vs. L) and logic (pebble automaton), whereas the latter is related to graph theory (treewidth). The first objective of Task 2.4 is to revisit both theories to make them applicable to the context of content delivery networks, for information retrieval.

Exploration is a difficult task in absence of any a priori information on the network (Rollik, 1980 ; Cook and Rackoff, 1980). In large scale networks such as the Web, it is currently impossible to give a map (even unprecise) of the network to the entities performing the exploration, either because this map would be too big, or because the network is simply not known. Therefore, the exploration problem will mostly be tackled in the frameworks defined in Tasks 3.2 and 3.3, that is we will look for the minimum amount of information to be given to the entities or to the nodes so that exploration can be performed efficiently.

Searching has been mostly considered in a centralized setting (Megiddo et al., 1986), which is clearly inappropriate in the context of content delivery networks. We will consider the search problem in the distributed setting. Preliminary results have been achieved in (Blin et al., 2006), but the search strategies defined there require a huge amount of local storage, and... they perform in exponential time !

Context of Task 3.1. Design of informative labeling schemes

When querying XML documents on the Web, it is desired to enable complex queries, like "books authored by X that cost at most Y euros". An XML document can be viewed as a tree whose nodes are the document items, and whose edges correspond to the relationship between these items. A query of the above form thus requires finding nodes with particular tags (book, author, price) sharing some relationship. In this example, interesting documents are book nodes that are ancestors of the required author and price nodes. The heart of a typical implementation of a search engine is a large hash table containing all words occurring in the database, where, for each word, the identifiers of the documents in which it appears is stored. For XML documents it is possible to refine the structure by maintaining tags at the nodes (like "book", "author", etc.), and by associating to each tag the labels of all nodes with this tag. Now, if one gives to the nodes some *meaningful labels* reflecting the hierarchical structure of the documents, queries of the above form can be answered by just using an index, without access to the documents. To enable good performances it is essential that the index resides in main memory, and so it is crucial to minimize the maximum label length associated to each node of the XML tree. This is one motivation for the "ancestry labeling scheme" problem in trees (Abiteboul et al. 2001). The objective of this problem is to label the nodes of an n-node tree so that ancestry queries between two nodes can be answered from the two corresponding node labels, without any other source of information. The best current label length for this problem is $\log n + O(\sqrt{\log n})$ bits (Abiteboul et al., 2005).

In fact, Peleg (2000) introduced the much more general notion of *informative labeling scheme* whose goal is to assign short labels to the nodes of a given graph so that specific queries between nodes can be answered from their associated labels, without access to any other source of information. This notion captures not only classical adjacency labeling schemes (Naor et al., STOC 1988), but also, e.g., distance labeling schemes (Thorup and Zwick, JACM 2005) and its variant *triangulations with beacons* (Kleinberg et al., 2004), as well as label based routing scheme (Talwar, STOC 2004 ; Slivkins PODC 2005).

Context of Task 3.2. Design of bounded size oracle

In this task, the knowledge of a node concerning the network is modeled by an *oracle*, which is a function $\mathcal{O}$ whose arguments are networks, and the value $\mathcal{O}(G, v)$, for a network G and $v \in V(G)$ is a binary string assigned to v. Intuitively, the oracle looks at the entire network and assigns to every node some information, coded as a string of bits. The *size* of the oracle on a given network G is the sum $\sum_{v \in V} |O(G, v)|$ of the lengths of all the strings it assigns to nodes. Hence this size is a measure of the amount of information about the network, available to its nodes. For instance, if the current number of nodes is known by every node in the network, the corresponding size of the oracle is $O(\log n)$ bits for each node, hence $O(n \log n)$ bits in total. Since the exact knowledge of the size of the network is an information that is difficult to get in practice for rapidly evolving networks (e.g., P2P) or extremely large networks (e.g., the Web), Task 3.2 is aiming at designing algorithms using oracles of much smaller size, for instance algorithms relying on rough approximations of the size of the network.

Precisely, solving a network problem $\mathcal{P}$ using oracle $\mathcal{O}$ consists in designing an algorithm that is *unaware* of the network G at hand but solves the problem $\mathcal{P}$ for it, as long as every node v of the network G is provided with the string of bits $\mathcal{O}(G, v)$. The formulation of the problem $\mathcal{P}$ may include a demand on the efficiency of the solution, thus we may be interested in solving problems within a prescribed time, or using at most a prescribed number of messages.

Given the problem $\mathcal{P}$, we ask what is the minimum size of an oracle for solving it. This minimum oracle size can be considered as a measure of the difficulty of the problem $\mathcal{P}$. The novelty and significance of the use of an oracle to model knowledge about the network is that it enables asking *quantitative* questions about the required knowledge, regardless of what *kind* of knowledge is supplied. This should be contrasted with the traditional approach that assumes availability of particular items of information, such as the neighborhood of a node. The fruitfulness of this approach has been recently demonstrated in various settings, including broadcasting and network exploration, for separating the difficulty of problems. For instance, linear broadcasting requires oracles of size $\Theta(n)$ bits whereas wakeup requires oracles of size $\Theta(n \log n)$ bits.

Context of Task 3.3. Protocol with limited probing capacity

In dynamic networks, the limited probing capacity of nodes is often one major constraint on protocol. Probing is however crucial for discovering ”good” neighbors in content distribution networks. For example, it is important to choose neighbors offering good connections with respect to either latency or bandwidth. For enabling high mobility or high versatility of both nodes and data, low probing algorithms have to be designed.

Context of Task 3.4. Sub-topologies preserving structuring properties

Dense sensor or ad hoc networks have too many links for nodes to be aware of all of them, and thus a working subset has to be selected (these networks may have a quadratic number of links whereas a linear number of links can reasonably be selected). The topology used by each node is then composed of the advertised links plus the links locally known to the node. The *link state routing problem* consists in selecting a subset of links that will be globally advertised, so that each node can compute shortest paths using solely its neighborhood and the advertised links. In fact, this problem leads to the sub-structure of *multipoint relays* (i.e., a set of neighbors that dominate the two-hop neighborhood), such as the ones used in the OLSR protocol, and introduced to optimize flooding. Indeed, the fact that two hop neighbors can compute optimal routes implies that the advertised links define a multipoint relay set. More generally, defining sparse advertised sub-topologies is an original approach we propose to extend.