Final Report on Project Number 124171): Optimal Mapping of Communication Services onto Cloud Computing Architectures Based on Discrete Mathematical Methods

(Mivel a publikációs követelmények kielégítése kizárólag angol nyelven történt, a végső szakmai beszámolót is angol nyelven készítettük el.)

The project is aimed at achieving high-impact results in the field of communication services over cloud computing architectures.The specific areas of research in the plan are network science, data structures, complexity, and graph and matroid theory, where we focus on virtual networks in data-centres and the implications of cloud-based services in terms of security, privacy, resilience, robustness, monitoring, and verification.
In the beginning of the project, we initiated research on *theoretical computer science using graph and matroid theory*. We have built game-theoretical tools for *measuring the reliability of networks*. We have defined and analyzed attacker-defender games and provided new graph reliability metrics and in some cases it can shed a new light on some well-known ones. According to plan, we also investigated the optimal pebbling number of some induced subgraphs of the square grid graph. We determined the graph parameters when the subgraph is induced by at most seven positive diagonals. In the project, we suggested several possibilities to represent the Fourier operator in the data structure. Although they are of exponential size, the fact that one can save a constant factor compared to the case when the entire matrix is stored still could be useful in simulations. In order to evaluate the expected availability of a service, a network administrator should consider all possible failure scenarios under the specific service availability model stipulated in the corresponding service-level agreement. Given the increase in natural disasters and malicious attacks with geographically extensive impact, considering only independent single link failures is often insufficient.

With respect to *reliable communication services*, in the project, we built a stochastic model of geographically correlated link failures caused by disasters, in order to estimate the hazards a network may be prone to, and to understand the complex correlation between possible link failures. With such a model, one can quickly extract information, such as the probability of an arbitrary set of links to fail simultaneously, the probability of two nodes to be disconnected, the probability of a path to survive a failure, etc. Furthermore, we introduced a pre-computation process, which en-ables us to succinctly represent the joint probability distribution of link failures. In particular, we generated, in polynomial time, a quasilinear-sized data structure, with which the joint failure probability of any set of links can be computed efficiently. In the project we presented a novel framework for disaster resilience, called FRADIR, which incorporates reliable network design, disaster failure modeling and protection routing in order to improve the availability of mission-critical applications.

In the next stage of the project, in the research work we have focused on the area of *cloud resilience and virtual network embedding*, through the mathematical modeling of pebbling movements on well-defined network structures. By studying Sidorenko's Conjecture, we have reached some interesting conclusions on the construction of special network structures. Sidorenko's Conjecture asserts that every bipartite graph H has the Sidorenko property, i.e. a quasi random graph minimizes the density of H among all graphs with the same edge density. We studied a stronger property, which requires that a quasirandom multipartite graph minimizes the density of H among all graphs with the same edge densities between its parts; this property is called the step Sidorenko property. We showed that many bipartite graphs fail to have the step Sidorenko property and used our results to show the existence of a bipartite edge-transitive graph that is not weakly norming.

For *cloud network monitoring,* we presented some deployment of monitoring-flows based network verification and failure localization approach for SDN networks. It not only minimizes the number of static forwarding rules but also significantly reduces the control plane load, i.e., reduces the total number of messages needed for network verification and failure localization. Our flexible hybrid implementation consists of MikroTik RB2011iLSIN and Open vSwitch (Mininet) switches, enabling the user to test network verification and failure localization in a more complex manner even if the number of physical devices is limited. We have studied the problem of finding a pair of shortest (node- or edge-) disjoint paths that can be represented by only two forwarding table entries per destination. Building on prior work on minimum length redundant trees, we showed that the complexity of the underlying mathematical problem is NP-complete and we presented fast heuristic algorithms. By extensive simulations we found that it is possible to very closely attain the absolute optimal path length with our algorithms (the gap is just 1–5%), eventually opening the door for wide-scale multipath routing deployments. We also showed that even if a primary tree is already given it remains NP-complete to find a minimum length secondary tree concerning this primary tree.
In order to evaluate the expected *availability of a service*, a network administrator should consider all possible failure scenarios under the specific service availability model stipulated in the corresponding service-level agreement. In the research, we built a stochastic model of geographically correlated link failures caused by disasters, in order to estimate the hazards a network may be prone to, and to understand the complex correlation between possible link failures. With such a model, one can quickly extract information, such as the probability of an arbitrary set of links to fail simultaneously, the probability of two nodes to be disconnected, the probability of a path to survive a failure, etc. Furthermore, we introduced a pre-computation process, which enables us to succinctly represent the joint probability distribution of link failures. In particular, we generated, in polynomial time, a quasilinear-sized data structure, with which the joint failure probability of any set of links can be computed efficiently. We proposed a data structure, called EGH filter, that supports the Bloom filter operations and besides it can guarantee false positive free operations for a finite universe and are restricted number of elements stored in the filter. Our data structure in the work was based on recently developed combinatorial group testing techniques.

Besides working on topics related to *virtual networks and resilience*, we have put more stress on *cloud network management, data forwarding mechanisms, and service migration*. More specifically, we published many results on greedy navigation and other navigation patterns in complex networks, entropy-based memory requirement calculations of hop-by-hop routing in computer networks, migrating legacy Ethernet Switches to SDN, full-stack software-defined networking, batch-scheduling data flow graphs with service-level objectives, survivable routing algorithms for diversity coding and disaster-resilient routing schemes. In the area of entropy-based memory requirement calculations of hop-by-hop routing in computer networks, we introduced a method that represents forwarding tables as consecutive character chains, which then is suitable to be analyzed as formed information package structure. As opposed to earlier research works, it is not dealing with worst-case scenarios but concentrates on the exact quantification of memory necessary to store selected network topologies and forwarding rules. We have created an optimization task for addressing scheme to set the namespace to optimize the information-theoretic print of the forwarding tables. We have developed well-defined constraints to predefined graph types and suggested a simple form to the namespaces of graphs. The results show that well-formed subnetworks, like ones occurring frequently in the Internet, can be stored using considerably less storage using memory optimization algorithms on the new namespaces. *The above work is the first attempt to solve the problems of the management of increasingly large computer networks with analyzing information content of routing namespaces.*

The constructed *cloud service migration mechanism* is to transfer Ethernet switch services for software-defined purposes. The method is capable of supporting the introduction of SDN architecture with the help of legacy devices. We have published a demo, where we demonstrate the efficiency of the service migration process showing that the modified layers do not come together with significant service quality degradation or packet forwarding performance degradation. At this stage, with respect to *generalized diversity coding*, we have dealt with the formulation of the problem of revealing important structural properties of the minimum cost survivable routings.

On network embedding, we investigated geometric routing of networks because of its locality and simplicity. This can operate in geometrically embedded networks in a distributed manner, distances are calculated based on coordinates of network nodes for choosing the next hop in the routing. Based only on node coordinates in any metric space, the Greedy Navigational Core (GNC) can be identified as the minimum set of links between these nodes which provides 100% greedy navigability. We performed results on structural greedy navigability as the level of presence of Greedy Navigational Cores in structural networks of the cloud networks.

Finally, during *network formation research* two-dimensional hyperbolic space turned out to be an efficient geometry for generative models of complex networks. The networks generated with this hyperbolic metric space share their basic structural properties (like small diameter or scale-free degree distribution) with several real networks. During the research, we presented a new model for generating trees in the two-dimensional hyperbolic plane. The generative model is not based on known hyperbolic network models: the trees are not inferred from the existing links of any network; instead, the hyperbolic tree is generated from scratch purely based on the hyperbolic coordinates of nodes. We showed that these hyperbolic trees have scale-free degree distributions and are present to a large extent both in synthetic hyperbolic complex networks and real ones that link Internet autonomous system topology embedded in the hyperbolic plane.

The original aim of the project funding was solely to support the publishing activity of the research group: attending conferences. covering the costs of journal papers, like overlength fees, or other types of printing and editing charges. We were aiming at achieving high-impact results in the field of computer science, mostly on the recent challenges in the communication services over cloud computing architectures. For that purpose, we have outlined some well-defined research areas in the project proposal, which we have planned to contribute to during the years supported by the Hungarian Scientific Research Fund. In the second half of the project, in the COVID era, due to lack of traveling possibilities and limited personal cooperations the funding was shifted towards covering expenses of journal papers, open access fees and involving younger researchers.