Publications of George Skretas

This document lists all peer-reviewed publications of George Skretas, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on August 13, 2022. An up-to-date version is available online at hpi.de/friedrich/docs/publist/skretas.pdf.

Journal articles

[1] Michail, O., Skretas, G., Spirakis, G. P., On the transformation capability of feasible mechanisms for programmable matter. In: *Computer and System Sciences*, pp. 18–39, 2019.

In this work, we study theoretical models of programmable matter systems. The systems under consideration consist of spherical modules, kept together by magnetic forces and able to perform two minimal mechanical operations (or movements): rotate around a neighbor and slide over a line. In terms of modeling, there are n nodes arranged in a 2-dimensional grid and forming some initial shape. The goal is for the initial shape A to transform to some target shape B by a sequence of movements. Most of the paper focuses on transformability questions, meaning whether it is in principle feasible to transform a given shape to another. We first consider the case in which only rotation is available to the nodes. Our main result is that deciding whether two given shapes A and B can be transformed to each other is in P. We then insist on rotation only and impose the restriction that the nodes must maintain global connectivity throughout the transformation. We prove that the corresponding transformability question is in PSPACE and study the problem of determining the minimum seeds that can make feasible otherwise infeasible transformations. Next we allow both rotations and slidings and prove universality: any two connected shapes A,B of the same number of nodes, can be transformed to each other without breaking connectivity. The worst-case number of movements of the generic strategy is $\Theta(n^2)$. We improve this to $\mathcal{O}(n)$ parallel time, by a pipelining strategy, and prove optimality of both by matching lower bounds. We next turn our attention to distributed transformations. The nodes are now distributed processes able to perform communicate-compute-move rounds. We provide distributed algorithms for a general type of transformation.

Conference papers

[2] Michail, O., Skretas, G., Spirakis, G. P., Distributed Computation and Reconfiguration in Actively Dynamic Networks. In: *Principles of Distributed Computing (PODC)*, pp. 448–457, 2020.

In this paper, motivated by recent advances in the algorithmic theory of dynamic networks, we study systems of distributed entities that can actively modify their communication network. This gives rise to distributed algorithms that apart from communication can also exploit network reconfiguration in order to carry out a given task. At the same time, the distributed task itself may now require a global reconfiguration from a given initial network G_s to a target network G_f from a family of networks having some good properties, like small diameter. With reasonably powerful computational entities, there is a straightforward algorithm that transforms any G_s into a spanning clique in $\mathcal{O}(\log n)$ time, where time is measured in synchronous rounds and n is the number of entities. From the clique, the algorithm can then compute any global function on inputs and reconfigure to any desired target network in one additional round. We argue that such a strategy, while time-optimal, is impractical for real applications. In real dynamic networks there is typically a cost associated with creating and maintaining connections. To formally capture such costs, we define three reasonable edge-complexity measures: the total edge activations, the maximum activated edges per round, and the maximum activated degree of a node. The clique formation strategy highlighted above, maximizes all of them. We aim at improved algorithms that will achieve (poly)$log(n)$ time while minimizing the edge-complexity for the general task of transforming any G_s into a G_f of diameter (poly)$\log(n)$. There is a natural trade-off between time and edge complexity. Our main lower bound shows that $\Omega(n)$ total edge activations and $\Omega(n/\log n)$ activations per round must be paid by any algorithm (even centralized) that achieves an optimum of $\Theta(\log n)$ rounds. On the positive side, we give three distributed algorithms for our general task. The first runs in $\mathcal{O}(logn)$ time, with at most $2n$ active edges per round, an optimal total of $\mathcal{O}(n \log n)$ edge activations, a maximum degree $n-1$, and a target network of diameter 2. The second achieves bounded degree by paying an additional logarithmic factor in time and in total edge activations, that is, $\mathcal{O}(\log^2 n)$ and $\mathcal{O}(n \log^2 n)$, respectively. It gives a target network of diameter $\mathcal{O}(\log n)$ and uses $\mathcal{O}(n)$ active edges per round. Our third algorithm shows that if we slightly increase the maximum degree to polylog(n) then we can achieve a running time of $\wr(\log^2 n)$. This novel model of distributed computation and reconfiguration in actively dynamic networks and the proposed measures of the edge complexity of distributed algorithms may open new avenues for research in the algorithmic theory of dynamic networks. At the same time, they may serve as an abstraction of more constrained active-reconfiguration systems, such as reconfigurable robotics which come with geometric constraints, and draw interesting connections with alternative network reconfiguration models, like overlay network construction and network constructors. We discuss several open problems and promising future research directions.

[3] Michail, O., Skretas, G., Spirakis, G. P., On the Transformation Capability of Feasible Mechanisms for Programmable Matter. In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 136:1–136:15, 2017.

In this work, we study theoretical models of programmable matter systems. The systems under consideration consist of spherical modules, kept together by magnetic forces and able to perform two minimal mechanical operations (or movements): rotate around a neighbor and slide over a line. In terms of modeling, there are n nodes arranged in a 2-dimensional grid and forming some initial shape. The goal is for the initial shape A to transform to some target shape B by a sequence of movements. Most of the paper focuses on transformability questions, meaning whether it is in principle feasible to transform a given shape to another. We first consider the case in which only rotation is available to the nodes. Our main result is that deciding whether two given shapes A and B can be transformed to each other is in P. We then insist on rotation only and impose the restriction that the nodes must maintain global connectivity throughout the transformation. We prove that the corresponding transformability question is in PSPACE and study the problem of determining the minimum seeds that can make feasible otherwise infeasible transformations. Next we allow both rotations

and slidings and prove universality: any two connected shapes A,B of the same number of nodes, can be transformed to each other without breaking connectivity. The worst-case number of movements of the generic strategy is $\Theta(n^2)$. We improve this to $\mathcal{O}(n)$ parallel time, by a pipelining strategy, and prove optimality of both by matching lower bounds. We next turn our attention to distributed transformations. The nodes are now distributed processes able to perform communicate-compute-move rounds. We provide distributed algorithms for a general type of transformation.