On the quantum simulation of complex networks

The Problem: The authors address the problem of implementing e^{-iAt} on a quantum computer (that is, approximating it up to a given precision parameter ϵ), where A is the adjacency matrix of the given graph. The problem is known to be solved when A is row-sparse and efficiently row-computable. However, it is not known to be the case for an arbitrary adjacency matrix. The authors attempt to address the problem for a more general class of graphs.

My opinion on the problem addressed: The problem addressed in the paper seems to be of fundamental interest to the community since e^{-iAt} implements quantum walks which is a framework to approach several graph problems via quantum computers. The main motivation for using quantum computers for such graph problems is for graphs encoding large complex networks, which often are NOT row-sparse and row-computable. Thus there is a clear need for extending algorithms implementing quantum walks for a larger class of graphs and thus I feel that the community definitely benefits from the investigation of the problem.

Results: The authors solve the case when the adjacency matrix A represents a graph with a small number of hubs, and are sparse otherwise. They call this class of graph *hub-sparse networks*. In essence, you take a sparse graph and add to them a small fraction of nodes that are connected to almost the entire network.

My opinion on the results: I feel that the main novelty of the paper is to find a class of graphs where the sparsity no longer holds but can still can be simulated efficiently. Given that efficient simulation of random walks for sparse graphs is already known, solving a class of graphs where we have a small fraction of densely connected nodes on top of it seems like modest progress. However, what I like a lot is that the hub-sparse graph, for which they manage to extend the results, also seems to be a very realistic subclass of networks from the application point of view.

Techniques: Ideally, the authors would like to use the Quantum Singular Value Transformation Theory, which says that if we can efficiently block-encode H/α, and we can implement e^{-itH} up to a constant error in $O(t\alpha)$ time. However, this technique is limited since one can never block encode H/α with

$\alpha < ||H||$. The spectral norm of the adjacency matrix of a hub-sparse network scales as $||A|| \sim \sqrt{N}$. So, the cost associated with block encoding is exponential in $\log(N)$, so this direction is hopeless.

To solve this, they split the adjacency matrix A as $A = G - A_- + A_h + A_r$. The advantage this offers is as follows: the matrices A_-, A_h, and A_r are sparse and hence have small spectral norms and thus have efficient block encoding; whereas, G, even though it is dense (with a high spectral norm and hence cannot be efficiently block encoded), has a very simple structure and hence a straightforward spectral decomposition. Using the spectral decomposition, they then go on to give a method to prepare eigenvectors efficiently which allows them to get the desired simulation efficiently. For preparing eigenvectors of G, they use Low and Wiebe's interaction picture simulation method to solve the time-dependent Hamiltonian simulation based on the truncated Dyson series. For block encoding the time-dependent Hamiltonian simulation, they do it for the terms separately and combine via the known technique of Linear Combination of Unitaries.

My opinion on the technique used: For getting the technique work, the core contribution is the observation that the adjacency matrix can be split in such a way that we have some sparse components whereas the dense component, being so simple, can still be simulated efficiently using carefully combining known techniques. So, in that sense, the paper's strength in techniques is not in inventing new techniques but in the careful application of known techniques to a cleverly defined interesting extension of the problem.

One limitation of the technique that is also present in the previous works which makes the real-life implementation of the technique questionable is the oracles assumed in the paper.

My opinion on the writing of the paper: Overall, in my opinion, the paper is well written with a good description of proof ideas of an otherwise quite technical proof. This helps an interested reader to navigate better. The problem, its motivation, and how the paper fits in the literature are also well explained.

Overall evaluation of the paper: Overall, I think this is a well-written paper that tackles a fundamental problem and makes modest progress toward it by cleverly finding a solvable subproblem and solving it with a careful combination of known techniques. I recommend acceptance.