

Researchers develop the fastest possible flow algorithm

June 28 2024, by Florian Meyer



The two thinkers behind the almost maximally fast flow algorithm: Rasmus Kyng and Maximilian Probst Gutenberg. Credit: ETH Zurich / Nicola Pitaro

In a breakthrough that brings to mind Lucky Luke—the man who shoots faster than his shadow—Rasmus Kyng and his team have developed a



superfast algorithm that looks set to transform an entire field of research.

The groundbreaking work by Kyng's team involves what is known as a <u>network flow</u> algorithm, which tackles the question of how to achieve the maximum flow in a network while simultaneously minimizing transport costs.

Imagine you are using the European transportation network and looking for the fastest and cheapest route to move as many goods as possible from Copenhagen to Milan. Kyng's algorithm can be applied in such cases to calculate the optimal, lowest-cost traffic flow for any kind of network—be it rail, road, water or the internet.

His algorithm performs these computations so fast that it can deliver the solution at the very moment a computer reads the data that describes the network.

Computations as fast as a network is big

Before Kyng, no one had ever managed to do that—even though researchers have been working on this problem for some 90 years. Previously, it took significantly longer to compute the optimal flow than to process the network data.

And as the network became larger and more complex, the required computing time increased much faster, comparatively speaking, than the actual size of the computing problem. This is why we also see flow problems in networks that are too large for a computer to even calculate.

Kyng's approach eliminates this problem: using his algorithm, computing time and network size increase at the same rate—a bit like going on a hike and constantly keeping up the same pace however steep the path



gets.

A glance at the raw figures shows just how far we have come: until the turn of the millennium, no algorithm managed to compute faster than $m^{1.5}$, where m stands for the number of connections in a network that the computer has to calculate, and just reading the network data once takes m time.

In 2004, the computing speed required to solve the problem was successfully reduced to m^{1.33}. Using Kyng's algorithm, the "additional" computing time required to reach the solution after reading the network data is now negligible.

Like a Porsche racing a horse-drawn carriage

The ETH Zurich researchers have thus developed what is, in theory, the fastest possible network flow algorithm. Two years ago, Kyng and his team presented mathematical proof of their concept in a groundbreaking paper. Scientists refer to these novel, almost optimally fast algorithms as "almost-linear-time algorithms," and the community of theoretical computer scientists responded to Kyng's breakthrough with a mixture of amazement and enthusiasm.

Kyng's doctoral supervisor, Daniel A. Spielman, Professor of Applied Mathematics and Computer Science at Yale and himself a pioneer and doyen in this field, compared the "absurdly fast" algorithm to a Porsche overtaking horse-drawn carriages.

As well as winning the 2022 Best Paper Award at the IEEE Annual Symposium on Foundations of Computer Science (FOCS), their paper was also <u>highlighted</u> in the computing journal *Communications of the ACM*, and the editors of popular science magazine Quanta named Kyng's algorithm one of the <u>ten biggest discoveries in computer science in 2022</u>.



The ETH Zurich researchers have since refined their approach and developed further almost-linear-time algorithms. For example, the first algorithm was still focused on fixed, static networks whose connections are directed, meaning they function like one-way streets in urban road networks.

The algorithms published this year are now also able to compute optimal flows for networks that incrementally change over time. Lightning-fast computation is an important step in tackling highly complex and datarich networks that change dynamically and very quickly, such as molecules or the brain in biology, or human friendships.

Lightning-fast algorithms for changing networks

On Thursday, Simon Meierhans—a member of Kyng's team—<u>presented</u> a new almost-linear-time algorithm at the Annual ACM Symposium on Theory of Computing (<u>STOC</u> 2024) in Vancouver. This algorithm solves the minimum-cost maximum-flow problem for networks that incrementally change as new connections are added.

Furthermore, in a second paper <u>accepted</u> by the IEEE Symposium on Foundations of Computer Science (FOCS) in October, the ETH researchers have developed another algorithm that also handles connections being removed.

Specifically, these algorithms identify the shortest routes in networks where connections are added or deleted. In real-world traffic networks, examples of such changes in Switzerland include the complete closure and then partial reopening of the Gotthard Base Tunnel in the months since summer 2023, or the recent landslide that destroyed part of the A13 motorway, which is the main alternative route to the Gotthard Road Tunnel.



Confronted with such changes, how does a computer, an online map service or a route planner calculate the lowest-cost and fastest connection between Milan and Copenhagen? Kyng's new algorithms also compute the optimal route for these incrementally or decrementally changing networks in almost-linear time—so quickly that the computing time for each new connection, whether added through rerouting or the creation of new routes, is again negligible.

But what exactly is it that makes Kyng's approach to computations so much faster than any other network flow algorithm? In principle, all computational methods are faced with the challenge of having to analyze the network in multiple iterations in order to find the optimal flow and the minimum-cost route. In doing so, they run through each of the different variants of which connections are open, closed or congested because they have reached their capacity limit.

Compute the whole? Or its parts?

Prior to Kyng, computer scientists tended to choose between two key strategies for solving this problem. One of these was modeled on the railway network and involved computing a whole section of the network with a modified flow of traffic in each iteration. The second strategy—inspired by power flows in the electricity grid—computed the entire network in each iteration but used statistical mean values for the modified flow of each section of the network in order to make the computation faster.

Kyng's team has now tied together the respective advantages of these two strategies in order to create a radical new combined approach. "Our approach is based on many small, efficient and low-cost computational steps, which—taken together—are much faster than a few large ones," says Maximilian Probst Gutenberg, a lecturer and member of Kyng's group, who played a key role in developing the almost-linear-time



algorithms.

A brief look at the history of this discipline adds an additional dimension to the significance of Kyng's breakthrough: flow problems in networks were among the first to be solved systematically with the help of algorithms in the 1950s, and flow algorithms played an important role in establishing theoretical computer science as a field of research in its own right.

The well-known algorithm developed by mathematicians Lester R. Ford Jr. and Delbert R. Fulkerson also stems from this period. Their <u>algorithm</u> efficiently solves the maximum-flow problem, which seeks to determine how to transport as many goods through a network as possible without exceeding the capacity of the individual routes.

Fast and wide-ranging

These advances showed researchers that the maximum-flow problem, the minimum-cost problem (transshipment or transportation problem) and many other important network-flow problems can all be viewed as special cases of the general minimum-cost flow problem.

Prior to Kyng's research, most algorithms were only able to solve one of these problems efficiently, though they could not do even this particularly quickly, nor could they be extended to the broader minimumcost flow problem.

The same applies to the pioneering flow algorithms of the 1970s, for which the theoretical computer scientists John Edward Hopcroft, Richard Manning Karp and Robert Endre Tarjan each received a Turing Award, regarded as the "Nobel Prize" of computer science. Karp received his in 1985; Hopcroft and Tarjan won theirs in 1986.



Shift in perspective from railways to electricity

It wasn't until 2004 that mathematicians and computer scientists Daniel Spielman and Shang-Hua Teng—and later Samuel Daitch—succeeded in writing algorithms that also provided a fast and efficient solution to the minimum-cost flow problem. It was this group that shifted the focus to power flows in the electricity grid.

Their switch in perspective from railways to electricity led to a key mathematical distinction: if a train is rerouted on the railway network because a line is out of service, the next best route according to the timetable may already be occupied by a different train.

In the electricity grid, it is possible for the electrons that make up a power flow to be partially diverted to a network connection through which other current is already flowing. Thus, unlike trains, the electrical current can, in mathematical terms, be "partially" moved to a new connection.

This partial rerouting enabled Spielman and his colleagues to compute such route changes much faster and, at the same time, to recalculate the entire network after each change. "We rejected Spielman's approach of creating the most powerful algorithms we could for the entire network," says Kyng.

"Instead, we applied his idea of partial route computation to the earlier approaches of Hopcroft and Karp." This computation of partial routes in each iteration played a major role in speeding up the overall flow computation.

A turning point in theoretical principles



Much of the ETH Zurich researchers' progress comes down to the decision to extend their work beyond the development of new algorithms. The team also uses and designs new mathematical tools that speed up their algorithms even more.

In particular, they have developed a new data structure for organizing network data. This makes it possible to identify any change to a network connection extremely quickly. And this, in turn, helps make the algorithmic solution so amazingly fast.

With so many applications lined up for the almost-linear-time algorithms and for tools such as the new data structure, the overall innovation spiral could soon be turning much faster than before.

Yet laying the foundations for solving very large problems that couldn't previously be computed efficiently is only one benefit of these significantly faster flow algorithms—because they also change the way in which computers calculate complex tasks in the first place.

"Over the past decade, there has been a revolution in the theoretical foundations for obtaining provably fast algorithms for foundational problems in theoretical computer science," writes an international group of researchers from University of California, Berkeley, which includes among its members Rasmus Kyng and Deeksha Adil, a researcher at the Institute for Theoretical Studies at ETH Zurich.

More information: Li Chen et al, Almost-Linear Time Algorithms for Incremental Graphs: Cycle Detection, SCCs, s-t Shortest Path, and Minimum-Cost Flow, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (2024). DOI: 10.1145/3618260.3649745

Almost-Linear Time Algorithms for Decremental Graphs: Min-Cost Flow and More via Duality. 65th IEEE Symposium on Foundations of



Computer Science (FOCS) 2024. <u>focs.computer.org/2024/accepte ...</u> <u>apers-for-focs-2024/</u>

Provided by ETH Zurich

Citation: Researchers develop the fastest possible flow algorithm (2024, June 28) retrieved 17 July 2024 from <u>https://techxplore.com/news/2024-06-fastest-algorithm.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.