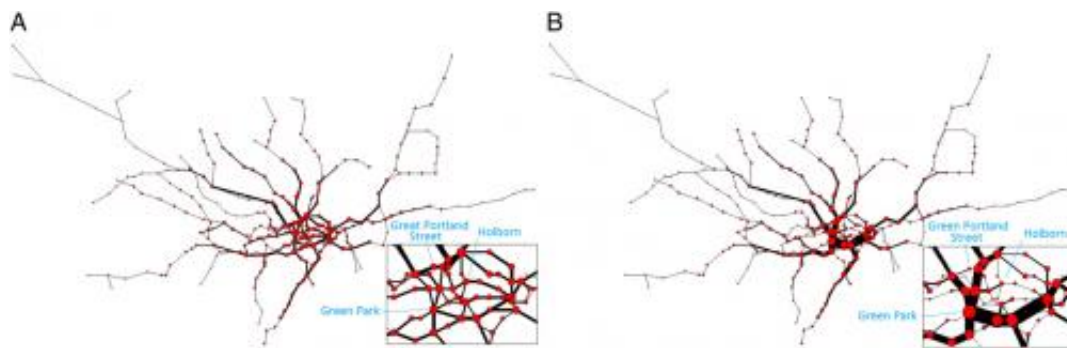# Planes, trains and molecules: Deriving a generic routing algorithm from the physics of interacting polymers

August 20 2013, by Stuart Mason Dambrot

Optimized traffic on the London subway network. A total of 218 real passenger source–destination pairs are optimized, corresponding to 5% of the data recorded by the Oyster card system between 8:30 AM and 8:31 AM on one Wednesday in November 2009 [cf. citation (35) in paper]. The network consists of 275 stations. (B). Red nodes correspond to stations with nonzero traffic. The size of each node and the thickness of each edge are proportional to traffic through them. (Insets) Zoomed-in views of the central region. Nodes are drawn according to their geographic position. Copyright © PNAS, doi:10.1073/pnas.1301111110

Finding a single optimal route is easy, but optimizing the *combination* of multiple routes is a challenge found in a wide range of applications including Internet instant messaging, peer-to-peer networks, subway traffic, airport flight management, water distribution systems, sensor deployment, military convoy logistics, and trip planning. Historically,

due to the computational complexity of deriving a global path optimization (that is, one that simultaneously considers all path possibilities), existing routing algorithms typically optimize each paths in isolation. As a consequence, the resulting solutions are less than optimal. Recently, however, scientists at Aston University, United Kingdom and The Hong Kong University of Science and Technology used the physics of interacting polymers (large molecule composed of many repeated subunits, known as monomers) and disordered systems to analyze macroscopic properties of generic path optimization problems. By so doing, they derived a simple yet global, routing algorithm capable of simultaneously considering all individual path alternatives. The researchers then demonstrated the algorithm utility by applying it to Internet-like random graphs, travel on the London Underground, and the global airport network. Moreover, their analysis revealed phase transitions, scaling laws, non-monotonic growth (that is, not always stable or increasing), and other new routing phenomena related to physics.

Research Fellow Chi Ho Yeung discussed the research he and his colleagues, Profs. David Saad and K. Y. Michael Wong, conducted – and the challenges they face – with *Phys.org*. "While, employing tools in physics to solve the system analytically was indeed our most difficult task," Yeung tells Phys.org, "the analogy between polymers and paths is actually easy to understand. A polymer is a long molecule chain likes a string with two ends," he illustrates, "Suppose I represent my travel path by a polymer: the two ends will be fixed representations of my starting point and destination, and the polymer body will be flexible depending on my path choice. If every traveler represents their path this way, we'd have a system of polymers on a transportation network – meaning that to suppress congestion, we'd introduce a repulsive force between polymers to discourage users using the same route. On the other hand, to encourage passengers share their path we'd introduce an attraction."

Turning now to their analytic work, Yeung points out that polymer paths are non-local variables, which are more difficult to analyze compared to local variables and interactions in conventional physics models. In addition, he notes that all polymers share the same network infrastructure and any two of them may have overlapping paths. "In our transportation network, when polymers overlap they either interact through an attraction or repulsion. It is thus equivalent to say that any two polymers may interact," Yeung continues, "and the extent of that interaction depends on the extent of overlap, which is again a non-local consideration regarding all polymers. With all these complications, we had to select the best paths out of all possible individual choices as well as their mutual overlapping."

Compared to ordinary polymer systems (which do not allow overlap), they researchers had a much larger pool of possible states, and thereby a much more difficult question to solve. "After deriving our theoretical results," Yeung adds, "we obtained the algorithm directly – and testing it on several datasets, found very good results. Once the system was analytically solved, it was straightforward to find its macroscopic properties, such as average path length and energy, by ordinary physics techniques in our area."

One key insight the scientists had, says Yeung, was that while some may think that the shortest path is always the best choice, this is not the case – and in fact, usual choice of going through the shortest path is a bad one when everyone takes the same route. "This isn't difficult to understand, as some observers may have already noticed. For example," Yeung illustrates, "during peak hours, some popular routes which lie on the shortest path may be overloaded, causing delays and making this path slower than a slightly longer one." Yeung points out that their simulations with the London metro data show that – compared to when all passengers travel through the shortest path – if they introduce a repulsive force between passenger paths, and if passengers follow the

suggested path, 20% of the assumed cost can be saved at a price of only 6% increase in average path length.
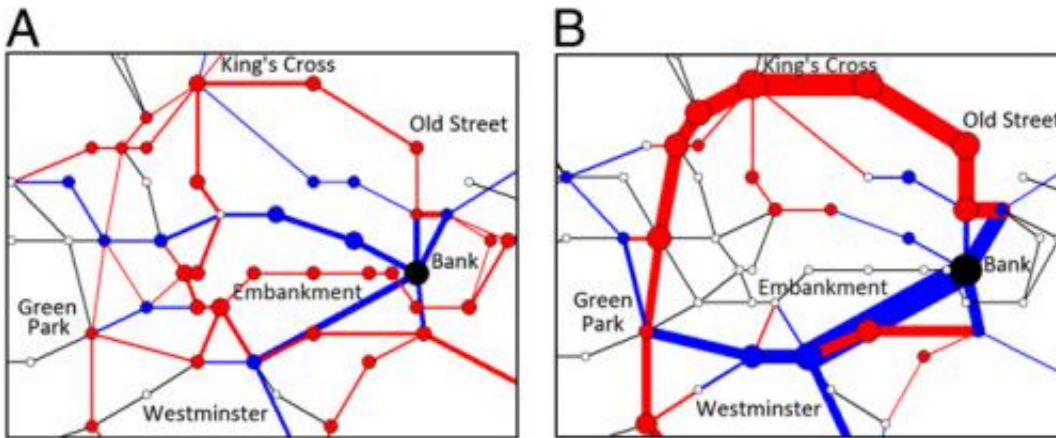
Changes in optimized traffic in the central London subway network after the removal of the station "Bank" (black node). The corresponding costs are $\gamma = 2$ (A) and $\gamma = 0.5$ (B). Nodes and edges that show an increase (decrease) in traffic appear in red (blue), where their size and thickness correspond to the magnitude of increase (decrease). Nodes and edges with no traffic changes appear in white and black, respectively. Passenger source–destination pairs are identical to those of Fig. 3 [cf. paper], except for the removal of pairs starting or ending destinations in Bank. Copyright © PNAS, doi:10.1073/pnas.1301111110

"On the other hand," Yeung continues, "although congestion does not occur in off-peak hours, the shortest path is *still* a bad choice, because the whole network has to remain active even for reduced usage on the less popular routes. Suppose we encourage off-peak commuters to travel on some common routes and share a large part of their journeys; then, the less popular bus routes or train lines can be discontinued which saves a lot of energy." Yeung adds that their results show that they can obtain such optimized shared path configurations through the analogy to an attractive polymer system.

Other than these results, the scientists also found that when they gradually change the polymer interaction from slightly repulsive to slightly attractive, there is a sharp increase in the number of idle nodes. "This is similar to a discontinuous phase transition observed in other physical systems," Yeung says. "Surprisingly, while the average path length does not change much, it does have an important implication on transportation or communication networks – that is, one can greatly increase the number of idle nodes without significantly lengthening the average path length, by introducing a slight attractive force between passenger paths. This may save a lot of resources in sparse traffic scenarios."

A key aspect of the researchers' results was demonstrating the algorithm's efficacy by applying it to random graphs resembling Internet *overlay networks* – that is, computer networks built on top of another network, in which nodes can be thought of as being connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network. "Networks representing websites interconnected by hyperlinks, or friends linked by instant messengers, are usually not bounded by physical location," Yeung notes, "and are well described by some specific random structures. We show in a simple random network how we can find the best choice of communication paths according to the attractive or repulsive strength we introduced."

In the case of repulsion, Yeung explains, individual communication paths avoid each other and at last almost everyone has its own path not overlapping with the others. In the case of attraction, the communications go through a small common region of the network, sharing their paths and leaving a lot of other nodes and links idle. "If we consider the idle nodes as routers," he points out, "a lot of energy can be saved by switching them off."

The researchers also applied their findings to travel on the London Underground network based on Oyster card data. "If we compare to the case where everyone takes their shortest path," says Yeung, "our simulations show substantial improvement on the London Metro network. Again, at a price of only 6% increase in average path length, 20% and 4% of the assumed cost are saved on the London metro network when one aims to balance or consolidate traffic, respectively. Of course, in practice," he acknowledges, "some experienced users would adopt a smarter route than the shortest path and the benefit from our algorithm would be less. However, I believe that in many transportation or communication networks there is still a large room of improvement in terms of energy saving if individual paths are well coordinated." Yeung adds that they did a very similar experiment, and obtain a similar result, in the global airport network.

Regarding other innovations that the scientists might develop and apply to the current experimental design, Yeung says that since physicists usually start with a generic model of physical systems, they've also assumed a model of interacting polymers which accommodates different type of interaction. "It turns out that we obtain a single algorithm which achieves various goals by tuning a single parameter controlling the attractive and repulsive strength between polymers," he explains. "Indeed, our approach can take into account interactions other than attraction and repulsion, and which may have other interesting applications. We welcome networking experts to suggest other specific routing problems which our algorithm may be able to tackle.

On the practical side, Yeung continues, one idea may be to develop a real-time application, based on their algorithm, to globally coordinate paths for individuals who start their journey at roughly the same time. "It's not the same as the usual route-finder applications that simply identify the shortest path for individuals without their interactions with others," he explains. "Rather, the envisioned application would

coordinate routes for many individuals who travel at the same time in order to achieve goals like balancing highway or tunnel usage, or to encourage train or plane sharing in off-peak hours or seasons."

Yeung also describes the planned next steps in their research. "Our path solution is static." He notes. "In other words, it provides an optimized path configuration given a set of destination pairs, and so suits many applications – but *not* those where the amount of traffic between individual destination pairs is rapidly changing. The next step is, perhaps, to develop routing algorithms based on our framework which address a dynamical routing task."

Yeung notes that other areas of research that might benefit from their study. "Our generic routing algorithm is applicable to any application that involves the path selection and coordination of individual paths," he tells Phys.org. "I hope our work can contribute to routing problems in transportation and communication networks, as well as sustainability research where the fixed infrastructure of existing transportation or communication systems is better utilized, thereby reducing the needs for further construction. In a more general respect," Yeung concludes, "I hope that our work demonstrates the power of rigorous physical tools when applied to interdisciplinary areas outside the realm of physics."

More information: From the physics of interacting polymers to optimizing routes on the London Underground, *PNAS* Published online before print July 29, 2013, doi:10.1073/pnas.1301111110