Euclidean Matchings in Ultra-Dense Spatial Communication Networks

Alexander Kartun-Giles[1]

[1]Complex Systems Group
School of Mathematical Sciences
Queen Mary University of London

MoN16 2017

Outline

1. Introduction to Ultra-Dense Networks
   - Introduction
2. Euclidean matching theory
   - Theorems
   - Using Data Capacity instead of Euclidean Length
3. Mean Field Models
   - The Random Assignment Problem
   - A Unique Feature of the Spatial Case: The AEU Property
4. Conclusions

Outline

1. Introduction to Ultra-Dense Networks
   - Introduction
2. Euclidean matching theory
   - Theorems
   - Using Data Capacity instead of Euclidean Length
3. Mean Field Models
   - The Random Assignment Problem
   - A Unique Feature of the Spatial Case: The AEU Property
4. Conclusions

What is an Ultra-Dense Network?

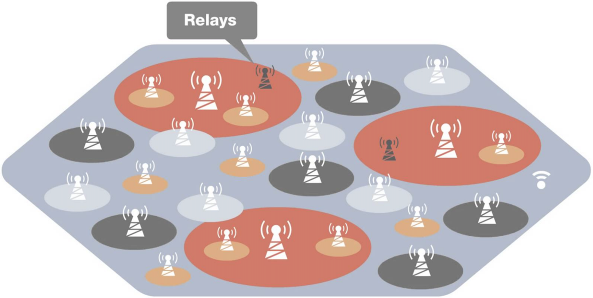

- **Cellular networks densify over time:** As traffic continues to grow rapidly in the coming years, today's largely macro cell networks will evolve, becoming more tightly packed and eventually becoming ultra dense.

(Picture source: BitvilleLearning).

What is an Ultra-Dense Network?

- **Cells approach 10m of each other:** An ultra-dense network (UDN) is one with sites on every lamp post or with indoor sites placed within 10m of each other.
- **Key 5G concept** By 2025 or 2030, Nokia expects UDNs to be covering most urban indoor and outdoor areas with small cells providing cell edge data rates of 100 Mbps to everyone.

Interference

Key engineering challenges:

- **Dense networks suffer from interference:** This constricts their capacity by limiting the concurrency of transmissions (see e.g. Gupta and Kumar 2000).
- **Beamforming is a key solution:** Recently, directional antennas have emerged as a promising technology due partly to their ability to enhance this concurrency, quantified as *spatial reuse ratio* i.e. the ratio between co-channel transmitter (T-T) and transmit-receiver (T-R) distances (see e.g. Li et al. 2011).

A classic stochastic geometry task: Optimise use of the spatial resource.

Outline

MST, Minimal Matchings and TSP

Three classic topics in geometric combinatorial optimisation.

Minimal Spanning Tree (MST)

With $x = \{x_1, x_2, \dots, x_n\}$ a finite set of points in $\mathbb{R}^d$ for $d \geq 2$, a minimal spanning tree T of x is a connected graph with vertex set x such that the sum of the edges lengths of T is minimal, i.e.

$$\sum_{e \in T} |e| = \min_G \sum_{e \in G} |e|$$

where $|x_i - x_j|$ is the Euclidean length of the edge e, and the minimum is over all connected graphs G.

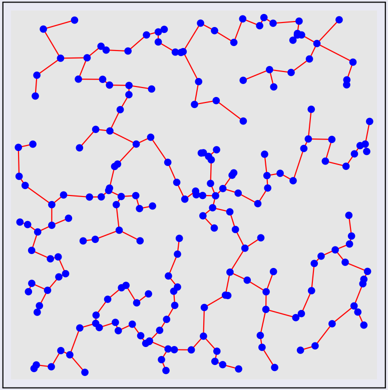

Minimal Matching

With $x = \{x_1, x_2, \ldots, x_n\}$ a finite set of points in $\mathbb{R}^d$, and a matrix (d_{ij}) of distances between them, find the *perfect matching* between the points (a set of unoriented links such that each point belongs to one and only one link), of shortest length. More formally, if d_{ij} is the Euclidean distance between nodes i and j, one looks for a set of occupation numbers $n_{ij} \in \{0, 1\}$ such that $\sum_j n_{ij} = 1$ with $\sum_{i<j} n_{ij} d_{ij}$ minimised.

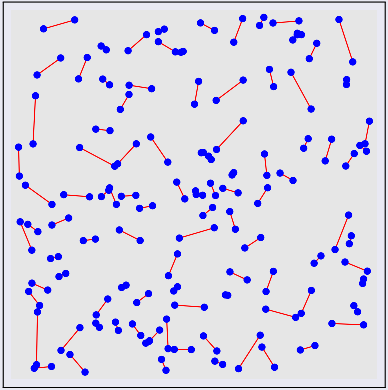

MST, Minimal Matchings and TSP

Travelling salesman tour

With $x = \{x_1, x_2, \ldots, x_n\}$ a finite set of points in $\mathbb{R}^d$ for $d \geq 2$, and a permutation σ of the points, such that the first point under the permutation is denoted $\sigma(1)$, the second $\sigma(2)$, and so on, a *shortest tour* though the points of x is a cycle $x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(n)}$ defined by σ, of minimum Euclidean length.

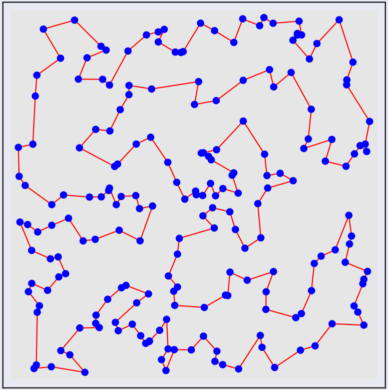

Scaling laws

Theorem

Beardwood, Halton and Hammersley (1959)

If X_i, $1 \leq i \leq \infty$ *are independent and identically distributed random variables with bounded support in* $\mathbb{R}^d$, *the the length* L_n *under the usual Euclidean metric of the shortest path through the points* $\{X_1, X_2, \ldots, X_n\}$ *satisfies*

$$\frac{L_n}{n^{(d-1)/d}} \to \beta_{TSP,d} \int_{\mathbb{R}^d} f^{(d-1)/d}(x)\mathrm{d}x \qquad (1)$$

almost surely, where $f(x)$ *is the density of the absolutely continuous part of the distribution of the* X_i, *and* $\beta_{TSP,d}$ *is a positive constant that depends on* d *but not on the distribution of the* X_i.

Scaling laws

Theorem

Steele (1989). If X_i, $1 \leq i \leq \infty$ are independent and identically distributed with compactly supported density f, and $0 < \alpha < d$,

$$n^{-(d-\alpha)/d} \sum_{e \in T} |e|^\alpha \to \gamma_{MST,\alpha,d} \int_{\mathbb{R}^d} f^{(d-\alpha)/d}(x) \mathrm{d}x \quad (2)$$

almost surely, where the positive constant $\gamma_{MST,\alpha,d}$ depends only on α and d.

Theorem

Ajtai, Komlos and Tusnady (1984). If X_i and Y_i are independent and uniformly distributed in $\mathbb{R}^2$ for $1 \leq i \leq n$, there are constants K_1 and K_2 such that

$$K_1 \sqrt{n \log n} \leq M_n = \min_\sigma \sum_{i=1}^{n} |X_i - Y_{\sigma(i)}| \leq K_2 \sqrt{n \log n} \quad (3)$$

with probability one as $n \to \infty$.

Outline

Data Capacity Application

Example

- For $N \in \mathbb{N}$, consider the *binomial point process* $\mathcal{X}_{2N} \subset [0,1]^d$ of $2N$ points. Form a *perfect matching* of the points, denoted $\mathcal{M}$.
- Call the Euclidean lengths of these edges $d_1, d_2, \ldots, d_N$. For $\eta > 0$, we then assign each edge its own data capacity $C_i = \log_2\left(1 + d_i^{-\eta}\right)$, based on a power law propagation model where signal power $P_i = Cd_i^{-\eta}$ taking η for the path loss exponent, and C a constant.
- For each matching, we therefore have a length $L_{\mathcal{M}} = \sum_{i=1}^{N} d_i$ and a data capacity

$$C_{\mathcal{M}} =: \sum_{i=1}^{N} \log_2(1 + d_i^{-\eta}). \tag{4}$$

Data Capacity Example

Example

Order of Capacity: *Given the capacity of a wireless link over a distance* d_i *is given by* $C(d_i) = \log_2(1 + d_i^{-\eta})$*, and that there exists a perfect matching of the points of* $\mathcal{X}_{2N}$ *where the Euclidean lengths of the edges in the matching are each of order* $N^{-1/d}$*, the one-hop throughput capacity is* $\mathcal{O}(\log N)$*.*

Proof: We have that $L_{\mathcal{M}}$ is of order $N^{1-1/d}$, since there are N edges each with length of order $N^{-1/d}$. Wireless links have a corresponding limiting capacity of order $C_i = \mathcal{O}\left(\log_2\left(1 + N^{\eta/d}\right)\right)$. Since N links can transmit simultaneously, the one-hop throughput capacity of the network is

$$\mathcal{O}\left(\frac{1}{N}\sum_{i=1}^{N}\log_2\left(1 + N^{\eta/d}\right)\right) = \mathcal{O}\left(N\log N\right). \tag{5}$$

Data Capacity Example

Example

The interpoint distances of the point process are $\mathcal{O}(N^{-1/d})$. The capacity $C_{\mathcal{M}} = \sum_{i=1}^{N} \log_2(1 + d_i^{-\eta})$ and therefore

$$\mathbb{E}C_{\mathcal{M}} = \mathcal{O}\left(N \log\left(1 + N^{-\eta/d}\right)\right) = \mathcal{O}\left(N \log N\right)$$

so rescale and study $C'_{\mathcal{M}} = (N \log N)^{-1} C_{\mathcal{M}}$. In other words, with the rescaling $d'_i = N^{1/d} d_i$ we look at

$$C'_{\mathcal{M}_{\text{opt}}} = \lim_{N\to\infty} \max_{\pi} \frac{1}{N} \sum_{i=1}^{N} \log_2\left(1 + d_i'^{-\eta}\right) \tag{6}$$

assuming a permutation of maximum capacity is arranged. We conjecture that this approaches a limit $C'_{\mathcal{M}_{\text{opt}}} \to \gamma_{MM,d,\eta}$ which depends only on d and η. Also, due to "self-averaging", $C'_{\mathcal{M}_{\text{opt}}} \to \mathbb{E}C'_{\mathcal{M}_{\text{opt}}}$ as $n \to \infty$.

Data Capacity Example

Example

Lower Bound: *Given the capacity of a wireless link over a Euclidean distance d_i is given by $C(d_i) = \log_2(1 + d_i^{-\eta})$, and that there exists a perfect matching on $2N$ points in the unit hypercube of lengths $d_1, d_2, \ldots, d_N$, whatever matching is chosen will imply a per-node capacity which satisfies*

$$\frac{1}{N}\sum_{i=1}^{N} \log_2(1 + d_i^{-\eta}) \geq \log_2\left(1 + \left(\frac{\sum_{i=1}^{N} d_i}{N}\right)^{-\eta}\right) \tag{7}$$

Proof: Taking $\Lambda \geq 0$, the maximum value of the product $\prod_{i=1}^{N} d_i$ under the condition $\sum_{i=1}^{N} d_i = \Lambda$ is obtained when $d_i = d_j$ for all i, j. Thus $\sum_{i=1}^{N} \log_2(1 + d_i^{-\eta}) = \log_2\left(\prod_{i=1}^{N}\left(1 + d_i^{-\eta}\right)\right)$ when all edges are of equal length, given $\eta > 0$, and the capacity will be *minimised*.

Outline

Random Assignment Problem

Now: fit a **mean field model** to the geometrically constrained Euclidean matching problem (**monopartite** case)

Theorem

Mézard and Parisi (1988). If X_i are independent and uniformly distributed points lying within a bounded region of $\mathbb{R}^2$, but we approximate the inter-point distances $d_1, d_2, \ldots, d_n$ as independent exponential random variables $\delta_1, \delta_2, \ldots, \delta_n$ with unit mean, then

$$\mathbb{E}\left[\min_{\sigma} \sum_{i=1}^{n} \delta_i\right] \to \frac{\zeta(2)}{2\pi} \tag{8}$$

as $n \to \infty$, where ζ is the Riemann zeta function. We use the exponential distribution because it is easier to work with, but the results apply to δ with any distribution which is strictly positive at zero, since they only need to agree at very short lengths $d_i \downarrow 0$.

Random Assignment Problem

Now: fit a **mean field model** to the geometrically constrained Euclidean matching problem (**bipartite** case).

Theorem

Aldous (2000). If X_i and Y_i are independent and uniformly distributed in a bounded region of $\mathbb{R}^2$, but we approximate the inter-point distances $d_1, d_2, \ldots, d_n$ as independent exponential random variables $\delta_1, \delta_2, \ldots, \delta_n$ with unit mean, then

$$\mathbb{E}\left[\min_{\sigma} \sum_{i=1}^{n} \delta_{i,\sigma(i)}\right] \to \frac{\pi^2}{6} = \zeta(2) \tag{9}$$

as $n \to \infty$. This is **the zeta(2) limit in the random assignment problem**. *We use the exponential distribution because it is easier to use, but the results apply to δ with any distribution which is strictly positive at zero, since they only need to agree at very short lengths $d_i \downarrow 0$.*

Outline

Data Capacity Example

According to Parisi (2008)

- Boltzmann statistical mechanics can be considered an example of a successful reductionistic program in the sense that it gives an microscopic derivation of the presence of emergent (collective) behaviour of a system which has many variables. This phenomenon is known as "phase transition".
- If the different phases are separated by a first order transition, just at the phase transition point a very interesting phenomenon is present: phase coexistence. This usually happens if we tune one parameter: the gas liquid coexistence is present on a line in the pressure-volume plane, while the liquid-gas-solid triple point is just a point in this plane. This behaviour is summarised by the Gibbs rule which states that, in absence of symmetries, we have to tune n parameters in order to have the coexistence of $n+1$ phases.

Definition of a Complex System

- The Gibbs rule is appropriate for many systems, however in the case of complex systems we have that the opposite situation is valid: the number of phases is very large (infinite) for a generic choice of parameters. This last property may be taken as a definition of a complex system.
- This is useful when you're able to control the microscopic details of the ground state (states of maximum data capacity), since one at least has a choice of configurations which satisfy a capacity bound.
- Key question: Are all these "phases" (in our case, Euclidean matchings) very similar? In the random assignment problem, they all are. This is called the AEU property (Asymptotic Essential Uniqueness). Aldous 2000: AEU asserts that every almost-optimal matching coincides with the optimal matching except on a small proportion of edges.

Definition of a Complex System

Studying AEU is interesting for two reasons, again Aldous (2000):

- First, it can be defined for many optimisation over random data problems, providing a theoretical classification of such problems (AEU either holds or fails in each problem) somewhat in the spirit of computational complexity theory.
- Second, in the statistical physics of disordered systems it has been suggested that the minima of the Hamiltonian should typically have an "ultrametric" structure, suggesting that in the associated optimisation problem the AEU property should fail. Since the random assignment problem (studied here as a toy model of ultra-dense networks) has qualitatively different behaviour than that predicted for more realistic models, such as the Euclidean matching problem with correlated edge weights.
- Open question: does the monopartite or bipartite Euclidean matching problem fail the AEU property? Very difficult proof.

Multihop transport

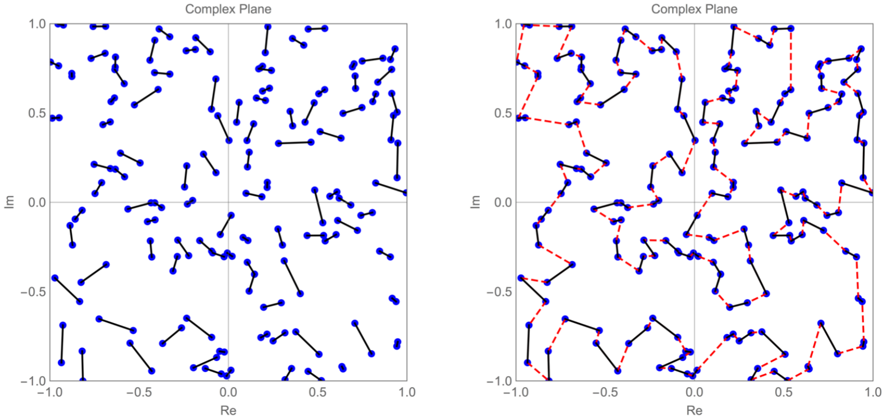

Augmenting paths (dashed line, red, right panel), switch between "highly orthogonal" matchings, allowing information to move between more than a single device pair. These updates form a chain of matchings. See e.g. The *switch chain* of Diaconis, Graham and Holmes (2001), and Dyer, Jerrum and Müller (2015).

Conclusions

Conclusions:

1. Euclidean matching theory provides a fascinating framework for both ultra-dense and D2D communication.
2. The multiple-valleys associated with complex systems here correspond to multiple orthogonal maximum capacity matchings. Determining the communication theory of this AEU property is an open problem.

Outlook: Can someone understand the role of the AEU property in communication networks? Does it have any specific uses?

Thank you.