

# Hitting Buneman Circles

Michael Paul Fourmanh  
Michael.Fourman@ed.ac.uk

School of Informatics, The University of Edinburgh

**Abstract.** We discuss Peter Buneman’s suggestion that a fibre connection to the internet — a hub — should be available within every circle enclosing a population of at least 2,000 people (a b-circle). This poses the problem of finding a small set,  $H$ , of hubs, such that *every* b-circle contains a hub. We show that a greedy algorithm does not lead to an optimal set of hubs. Instead it models market forces, which are naturally greedy. An unfettered market will exploit the most profitable communities and, just like a greedy algorithm, leave gaps that it is uneconomic to fill. We describe a geometric heuristic for the discovery of efficient hub placements satisfying a purely combinatorial analogue of Buneman’s criterion, and apply it to illustrate the inherent inefficiency of gap-funding in a market-led broadband policy.


**Key words:** hitting set, approximation algorithm, facilities location, broadband, market forces, gap funding.

## Introduction

In Scotland, as elsewhere, many communities are too far from an optical connection to the internet to allow them to benefit from the low-latency, high-bandwidth, symmetric internet connections that will soon be routine in metropolitan areas.

This work arose from the Royal Society of Edinburgh’s Digital Scotland inquiry [8]. It was easy for the inquiry to decide that Scotland should have enough “fibred points of presence” (FPOPs, or *hubs*) to provide every community in the country with adequate backhaul.

Quantifying where these should be and where the fibre should be laid, without embarking on detailed infrastructure planning, was more difficult. Our diagram shows the areas served by a collection of hubs satisfying a simple, but ingenious criterion, suggested by Peter Buneman. Such provision



would allow every substantial community in Scotland to access next generation broadband.

The population of Scotland, and other countries, is not randomly distributed. Most people live in relatively dense clusters, many of which are widely separated. Because communities are hard to identify objectively, it is difficult to specify an equitable policy for the provision of a fibre backhaul connection — a *hub* — to every community.

Buneman’s criterion,

*“Draw any circle (I mean any) on a map of Scotland. If that circle contains more than 2000 people, then the circle must also contain a fibre point-of-presence.”*<sup>1</sup>

was adopted as the key recommendation of the Digital Scotland report, intended to ensure that every community will have access to Scotland’s digital infrastructure. The rationale for this recommendation is that the aggregate bandwidth demand from a population of this size will, at peak times, exceed the capacity of a copper or wireless backhaul connection<sup>2</sup>.

The Digital Scotland report assumes that a population of 2,000 corresponds to approximately 800 premises, and that a copper or wireless backhaul connection is limited to  $\sim 512$  Mb/s. With these assumptions, the available bandwidth per subscriber is at most  $\sim 16$  Mb/s, at a contention ratio<sup>3</sup> of 25 : 1.

To a first approximation, the bandwidth available for a communication channel is proportional to the frequency of the carrier signal. Since the electromagnetic frequencies used for optical signals are around  $10^5$  times those used for electrical or wireless communications, fibre can carry correspondingly higher bandwidths. Thus, a single fibre could provide the entire population of the UK with more bandwidth *per caput* than a wireless or copper channel can deliver to a community of 2,000.

Once a community has access to backhaul there are many technologies available for the creation of a local access network — striking examples are given by the Tegola network developed by Peter and colleagues [1], and other projects it has inspired.

The Buneman criterion provides a novel approach to facilities location, in that it analyses the infrastructure requirements for distribution of a utility good, rather than the more usual focus on placement for profit.

The methods introduced here should also have application in areas other than backhaul provision. For example, to find solutions to the wireless base-station placement problem that, unlike much earlier work (see e.g. [5]), impose limits on the number of clients served by each base station. Such constraints are increasingly relevant as contention for spectrum increases.

---

<sup>1</sup> Personal communication, 2010.

<sup>2</sup> A *backhaul connection* is the essential link from a local access network to the internet.

<sup>3</sup> Backhaul is typically provisioned on the assumption that most subscribers will be idle most of the time. Contention is the ratio of the total number of subscribers to the number that can be served concurrently, at the advertised bandwidth.

## Discussion

Buneman circles (*b-circles*) are those circles on the map of Scotland that include a settled population of at least 2,000 individuals. Every b-circle should include a hub. The beauty of Buneman’s criterion is that *any* collection of 2,000 or more individuals (we call such a collection *substantial*) may be viewed as a community, and thus every substantial community is guaranteed at least one fibre hub within any circumscribing circle. Each individual will belong to many such communities.

We now set out a mathematical context for discussion of this criterion. To begin, let  $\mathcal{S}$  (Scotland) be a bounded subset of the plane,  $\mathbb{R}^2$ .

**Definition 1.** A population distribution is a finite measure,  $\pi$  on  $\mathcal{S}$ .

We identify two special cases:

A population is the counting measure on a finite set of points  $P \subseteq \mathcal{S}$ .

A census is a discrete measure, given by a finite set of census points,  $A \subseteq \mathcal{S}$ , together with a count  $\pi(a) > 0$  of the population ascribed to each  $a \in A$ .

It is often helpful to think of a population as a sample from the probability distribution associated with a more abstract population distribution, and a census as the result of aggregating a population to a (relatively) small number of census points. Throughout this paper, we assume a given population distribution,  $\pi$ , and speak of  $\pi(C)$  as the *population* of  $C \subseteq \mathcal{S}$ .

The data used for the Digital Scotland analysis of backhaul requirements is a census, derived from postcode data by multiplying the number of residential addresses for each postcode by a factor of 2.5 inhabitants per household.

**Definition 2.** A Buneman circle is a disc,  $D$ , with population  $\pi(D) \geq k$  (hereinafter, a b-circle).<sup>4</sup>

We write  $\mathcal{B}_x$  for the minimal-radius b-circle centred at  $x$ , and  $\beta(x)$  for the radius of this circle — the b-radius of  $x$ .

Buneman computed the b-circle  $\mathcal{B}_p$  for every one of the 196,273 postcodes,  $p$ , in Scotland. The calculated b-radii range from 84 metres (AB25 1FE, a high-rise in Aberdeen) to 55 kilometres (ZE2 9JU, the Fair Isle). Buneman’s computation places the data in a k-d tree, and then uses standard k-nearest neighbour queries and binary search to determine  $\beta(x)$  for any point  $x$ . For UK postcode data, k-d trees provide adequate performance on standard hardware. See [4] for a recent review of other methods, suitable for massive data sets.

Recall that any set which includes least one element in every set in a collection  $X$  is called a *hitting set* for  $X$ . The Digital Scotland recommendation stipulates that the hubs should form a hitting set for the collection of *all* b-circles. Dually, for each potential hub location,  $h$  we can consider its *client set*  $\mathcal{C}_h = \{x \mid h \in \mathcal{B}_x\}$ . Clearly,

$H$  is a hitting set for  $\{\mathcal{B}_x \mid x \in \mathbb{R}^2\}$  iff  $\{\mathcal{C}_h \mid h \in H\}$  covers  $\mathbb{R}^2$ .

---

<sup>4</sup> For the purposes of this note,  $k$  can be considered to be a fixed integer, substantially smaller than the total population,  $\pi(\mathcal{S})$ .

## Properties and anomalies

**Lemma 1.**  $\beta$  is uniformly continuous.  $\mathcal{C}_h$  is star-convex.

*Proof.* Given two points,  $x, y$ , we claim that  $|\beta(x) - \beta(y)| \leq |x - y|$ . It suffices to show that  $\beta(y) \geq \beta(x) - |x - y|$ .

The result then follows by symmetry.

Observe that  $\forall r, z. \pi(D_{r,z}) < k$  iff  $r < \beta(z)$  (where  $D_{r,z}$  is the disc of radius  $r$  centred at  $z$ ).


For  $y \in \mathcal{B}_x$ , suppose  $s < \beta(x) - |x - y|$ , then,

$$s + |x - y| < \beta(x), \text{ and}$$

$$D_{s,y} \subseteq D_{s+|x-y|,x}$$

$$\text{so } \pi(D_{s,y}) \leq \pi(D_{s+|x-y|,x}) < k.$$

$$\text{Thus, } s < \beta(y).$$



The same construction also shows that  $\mathcal{C}_h$  is star-convex. If  $x \in \mathcal{C}_h$  (equivalently, if  $h \in \mathcal{B}_x$ ) and  $y$  lies on the radius (of length  $\beta(x)$ ) from  $x$  through  $h$ , then  $|y - h| \leq \beta(y)$ , so  $y \in \mathcal{C}_h$ .  $\square$

Buneman’s intention [personal communication] was universal provision: that a “circle drawn on the map of Scotland” really should include *any* circle — “*even one centred in Newfoundland*” — on an infinitely extended map. We call this the *strong Buneman criterion*.

Interactions between the geometry of circles and a discrete population can be surprising. One consequence of the strong criterion is that there must be hubs near the boundaries of  $\mathcal{S}$ . For any  $\varepsilon > 0$  and any half-plane  $H$  with  $\pi(H) \geq k$  we must place a hub within  $\varepsilon$  of  $H$  — since b-circles with arbitrarily distant centres must be hit. In specially coincidental cases,  $\mathcal{B}_x$  may have an arbitrarily large population — for example, if  $P$  contains a large number of individuals all on a circle, such as the shores of a circular loch. In general, a population will be in general position — a sample from a distribution almost certainly will — but a census is certainly not. If we take a community,  $C \subseteq \mathcal{S}$ , even one of the form  $\mathcal{B}_x$ , then the circumcircle of the population of  $C$  may contain many more individuals than does  $C$ .

Clearly, the Buneman criterion is satisfied if *every* disc  $\mathcal{B}_x$  contains a hub. However, this set of discs is infinite, and so is the set of possible hub locations.

We have not attempted to produce hub sets satisfying the strong criterion. Instead, we weaken the condition, by interpreting a “circle drawn on the map of Scotland” to mean a disc  $D_{r,x}$  whose centre,  $x \in \text{supp}(\pi)$ , belongs to the support of  $\pi$  — we call this the *egocentric* interpretation. For a discrete population, or census, this simply means that  $x$  ranges over the sample points, so we now have only finitely many ( $O(n)$ ) circles to hit.

This replaces universal provision with egocentric guarantees. Wherever you live in Scotland, your nearest hub should be among your 800 nearest neighbours. It is also natural to strengthen the criterion by requiring that hubs should be built only in habitable regions — so we look for a hitting set that is also contained in  $\text{supp}(\pi)$ .

A more appealing, but less tractable, community-focussed guarantee would stipulate that for every substantial community  $C \subseteq \mathcal{S}$  the circumcircle of  $C$  should contain a hub. A substantial community distributed around the shores of a circular loch would

then be guaranteed a loch-side hub — or a floating one — whereas the egocentric criterion might only place hubs well inland.

Nevertheless, in the interest of tractability, and because we are not aware of a need for a community-centric guarantee arising in practice, we use the egocentric version, with hubs located in populated sites. Let  $P = \text{supp}(\pi)$ . For  $p \in P$  we define

$$\mathcal{N}_p = \mathcal{B}_p \cap P \quad \mathcal{K}_p = \mathcal{C}_p \cap P$$

$H \subseteq P$  is a hitting set for  $\{\mathcal{N}_p \mid p \in P\}$  iff  $\{\mathcal{K}_h \mid h \in H\}$  covers  $P$ .

We want to find small hitting sets,  $H$ , such that  $\{\mathcal{K}_h \mid h \in H\}$  covers  $P$ . For a discrete population distribution, this is a purely combinatorial problem.

## Selecting hubs

Standard approaches to the hitting-set problem include a naïve greedy algorithm and linear programming relaxation (LPR) (see [10], Ch. 1). We tried both on Buneman’s postcode data. They did not perform well. A geometrically-motivated heuristic produces much better results.

We used various types of synthetic data, as well as the postcode data, to understand why these approaches fail and to evaluate alternative methods. These data sets included:

- a uniform lattice of  $n^2$  points in the unit square,
- samples of size  $n^2$  drawn from the uniform distribution on the unit square,
- samples of size  $n^2$  drawn from the bivariate normal distribution  $\mathcal{N}_2(\mathbf{0}, \mathbf{1})$ .

The randomised examples are chosen to model the kinds of local variation (and in some areas local uniformity) that we see in our population data. Uniform density across a region with a sharp boundary is often seen near a coastline, and a unimodal peak characterises many isolated communities.

Consider covering a uniform  $n \times n$  square lattice with discs containing  $k$  lattice points.. Take  $k = n$  and let  $n \rightarrow \infty$ , then, away from the boundary, we approach the well-known geometric problem of covering the plane with uniform discs. The familiar regular hexagonal cover always gives a baseline solution (which may sometimes be improved upon by exploiting quantisation gaps in the lattice).

The greedy algorithm first chooses a maximal disjoint set of circles, and then fills in the gaps, normally producing a sub-baseline solution. LPR uses randomised rounding to derive a cover from a non-integral solution to an integer linear program (ILP) expressing the constraints. Away from the boundary, our problem is regular. Solution of the LP relaxation gives equal weight to indistinguishable candidates, and so provides no useful information — randomised rounding amounts to random selection.

*A geometric heuristic:* Our heuristic is initially designed to produce good solutions for the uniform lattice, where we are guided by a clear geometric picture. We find that, with a minor modification described below, it also performs well on randomised and real-life data.


Consider the continuous limiting case of a uniform population distribution extending infinitely. The b-radius is uniform. Without loss, assume that  $\beta(x) = 1$ . We write  $\mathcal{B}_x$  for the b-circle  $D_{1,x}$  of  $x$ , and note that each client set is also a unit disc  $\mathcal{C}_h = D_{1,h}$ . For this idealised setting, we want to produce a regular hexagonal covering of the plane

by discs  $\mathcal{C}_h$ , using an algorithm that only has access to the combinatorial relationship  $x \in \mathcal{C}_h$  or, equivalently,  $h \in \mathcal{B}_x$ .

Suppose we have already placed two hubs,  $a, b$ , which are (by magic)  $\sqrt{3}$  apart. What discs might we add to extend our cover?

It is clear that we need to cover the uncovered points near the intersections of our two circles. We identify these as points that are hard to cover.


We write  $U$  for the set of uncovered points. In general, a point  $p$  can be covered by any disc  $\mathcal{C}_h$  centred at  $h \in \mathcal{B}_p$ , but we can characterise the particular points we want to identify if we only consider discs with uncovered centres,  $h \in \mathcal{B}_p \cap U$ . For  $p$  near a covering circle, many points in  $\mathcal{B}_p$  are already covered by that circle. For uncovered points,  $p \in U$ , near to *both* circles few points in  $\mathcal{B}_p$  remain uncovered.



For uncovered points,  $p$ , near a cusp where the two circles meet  $\mathcal{B}_p \cap U$  is smaller than it is for uncovered points further from the cusp — the points nearer the cusp are harder to cover. So, we say that, ‘covering  $p$  is hard’, if  $p$  is not yet covered and there are few uncovered  $h$  such that  $h \in \mathcal{B}_p$  (or, equivalently,  $p \in \mathcal{C}_h$ ) (the fewer, the harder). Then we restrict our attention, temporarily, to those  $h$  that cover the hardest  $p$ , and from these select, greedily, an  $h$  such that  $\mathcal{C}_h$  covers as many uncovered  $q$  as possible.

The diagram shows our two original discs, with centres  $a, b$ , together with a few of the potential  $\mathcal{C}_h$  we have to consider — with their centres,  $h$ , sized to indicate the relative size of  $\mathcal{C}_h \cap U$ . The candidate covering circles,  $\mathcal{C}_h$ , shown all cover all uncovered points in a neighbourhood of the upper cusp. Choosing greedily from among these candidates selects the central one, because it includes the largest uncovered area.

If we start from our magical initial configuration and repeat this procedure indefinitely we will incrementally generate the regular covering of the plane. If we start instead from a pair of discs that overlap, but have less than the magic separation, what then? Our new circle, centred at  $h$ , includes the hard to reach points and passes through one intersection of the original two circles. It is placed symmetrically with respect to the originals. If the angle between the two original radii is  $\theta$ , then the angle between the new radius and either original is



$$\theta' = \theta + (2\pi/3 - \theta)/2.$$

The same formula applies if we start with two discs marginally too far apart. So, as we iterate our construction, successive values of  $\theta$  approach the ideal value,  $2\pi/3$ . This explains how our algorithm adapts to variations in population density, and why it produces reasonable solutions for our synthetic data.

In the next section we transfer this idea to our discrete setting and report on our real-world application.

## An algorithm

We use the heuristic introduced above to select new hubs, while keeping track of the set  $U$  of uncovered sites.

*Iteration:* One step of the algorithm consists of choosing a site,  $p$ , to cover; choosing a hub,  $h \in \mathcal{N}_p$ , so that  $\mathcal{K}_h$  will cover  $p$  (we will say, ‘ $h$  covers  $p$ ’); then updating  $U$ . We take such steps until  $U$  is empty.

First, for each uncovered  $p$ , the set of uncovered candidates that would cover it is given by  $N(p) = U \cap \mathcal{N}_p$ . The difficulty of covering  $p \in U$  is given by the reciprocal of the number of uncovered sites potentially available to cover it  $d(p) = 1/|N(p)|$ . For  $q \notin U$  set  $d(q) = 0$ . We will cover some  $p$  of maximal difficulty, chosen from the set  $M = \{p | d(p) = \max_{q \in U} d(q)\}$ .

Second, we must choose  $h$  from  $\bigcup_{p \in M} \mathcal{N}_p$ . For the purely geometric version of our problem described in the previous section, a greedy selection, based solely on maximising the number of newly covered sites, will choose the ‘correct’ hub for each  $p$ . For the discrete version of the problem, we find that quantisation introduces enough noise to perturb this effect: a purely greedy choice will often leave a hard-to-cover site that can only be inefficiently covered. To overcome this effect, we again favour choices that cover difficult sites, using  $d(q)^2$  as a measure of the value of covering  $q$ .<sup>5</sup> So, we choose a hub,  $h \in \bigcup_{p \in M} \mathcal{N}_p$ , to maximize  $\sum_{q \in \mathcal{K}_h} d(q)^2$ . Finally, we remove  $\mathcal{K}_h$  from  $U$ .

*Initialisation:* We can start our procedure with an empty collection of hubs, and every site uncovered, or with an arbitrary collection of hubs, with the sites already covered computed accordingly. In the geometric setting, starting from an empty set of hubs, the initial moves introduce two tangential circles, then a third circle centred at their meeting point. To produce a uniform cover we could use a few iterations to produce a pair with near-perfect separation, and only then start to produce our cover. In practice, we have found that this does not give perceptibly improved results on noisy data.

*Completion:* Once all sites are covered, we may find that a handful (typically  $\ll 1\%$ ) of client sets in our chosen covering are double-covered: some chosen hubs have been rendered redundant by later choices. As long as any such remain, we prune one, and so eventually arrive at an irredundant cover.


## Results

The picture illustrates a covering,  $H$ , of a sample of 40,000 points drawn from a 2D Gaussian, with  $k = 200$ , by showing the convex hull of each  $\mathcal{K}_h$  for  $h \in H$ .

First, we used synthetic data to develop and tune our heuristic. The uniform hexagonal cover of the plane by discs has a multiplicity of  $\frac{2\pi}{3\sqrt{3}} \approx 1.21$ . We take this as a baseline against which we compare the efficiency of our coverings.

We used random samples of 40,000 points, from normal and uniform distributions, to compare the use of different exponents,  $w$ , in our measure,  $d(q)^w$ , of the value of covering  $q$ . The table below shows the multiplicities of typical covers.

$w$	0	1	2	3	4
2D Gaussian	1.48	1.45	1.36	1.35	1.4
Uniform	1.45	1.375	1.315	1.35	1.35



<sup>5</sup> We have no principled justification for the quadratic exponent used here. It was chosen following the experimentation described in our results.

As discussed earlier our population data can be modelled informally as a mix of synthetic data from these two distributions. This justifies the choice of a quadratic exponent for our hardness measure. Pragmatically, we find that it gives lower multiplicities (smaller covers), on our data, than other small integer values.

Applying the same algorithm to Buneman’s circles for Scotland, we obtain a cover with 4,224 hubs. However, a further relaxation of the Buneman criterion, to guarantee that every b-circle includes some site with a hub within 500 metres of it, dramatically reduces the number of hubs we require. We implement this relaxation by simply adding any site within 500m of a site in  $\mathcal{N}_x$  to  $\mathcal{N}_x$ . This clearly has the greatest effect in the most densely populated areas. The diagram on our title page of this article depicts the convex hull of  $\mathcal{K}_h$  for each of the resulting collection of 1,652 hubs.

A naïve greedy algorithm, applied to this relaxed problem, produces a cover with 1,803 hubs. For the sake of comparison, we have run a less naïve greedy algorithm on the same data. For this, we modified the code to omit the first part of each step (restricting our choice to hubs that cover the hardest sites), so at each step we consider every site not yet used as a potential hubs, then make a greedy choice favouring hubs that cover difficult sites. With a quadratic weighting, this gave a cover with 1,734 hubs.

We have shown that a greedy algorithm does not lead to an optimal cover. Market forces, however, are greedy. An unfettered market will exploit the most profitable communities and, just like the greedy algorithm, leave gaps that it is uneconomic to fill. Our next experiment suggests that market-led greed, which will focus initial investment on the most compact communities, is even less efficient than naïve greed, which looks for hubs that serve as many premises as possible, even if these premises are thinly spread.

Scotland’s current broadband policy, developed in consultation with the incumbent supplier, is to provide next generation access to over 80% of the population by 2015. Roughly 72% of Scotland’s postcodes have b-radii  $\leq 700m$ . We call these *compact* communities. They account for just over 80% of Scotland’s residential addresses.

We have run our algorithm to find a set of hubs that will serve just these compact communities, to try to model the likely impact of government policy. Our algorithm produces 1,092 hubs that cover these communities. If we then rerun the algorithm to cover all the communities not served by these 1,092 hubs, we find we need a further 707 hubs to complete the job. So, *in toto* this two-step approach would require over 10% more than the 1,652 hubs in our one-step plan for universal provision; and public support would be required for 707 hubs — over 25% more than the gap of 560 between our 1,652 requirement and the 1,092 financed by the market.

Clearly, what we have just sketched is a very simplistic, indeed over-simplistic, model of the complex planning and investment decisions negotiated between government and incumbent. Nevertheless, we believe that even this simple model captures a key challenge that policy-makers face in trying to ensure universal provision while being committed to a gap-funding policy that waits for gaps in market-driven provision to appear, before intervening to fill them.

Market-led provision is, by definition and design, greedy, and thus becomes inefficient beyond a certain point. Where it will not deliver universal provision, *post hoc* gap-filling will be inefficient.

## Related work

The idea of using distance to nearest neighbours to study the spatial distribution of populations is not new. Clark and Evans [2] use distance to nearest neighbour to study



how the distributions exhibited by populations of living organisms differ from a Poisson ideal. Loftsgaarden and Quesenberry ([6]) introduce the key idea of using distance to  $k^{\text{th}}$ -nearest event for density estimation. In Smoothed Particle Hydrodynamics (SPH) [7] “the mass of each point is distributed according to a smoothing function  $W$  whose size adapts to the local value of the density of points.” For the visualisation of galaxies Colberg ([3]) uses a grid-based adaptation of SPH. The Delaunay Triangulation Field Estimator (DTFE) introduced by Schaap and van der Weygaert ([9]), takes a Delaunay triangulation, and then considers the  $n$ -simplices incident at a point as a natural neighbourhood of that point. These examples are not exhaustive, but this brief list gives some idea of the range of applications where related ideas have been considered.

Buneman’s insight that we could use this estimator of local density as a basis for the provisioning of a rate-limited resource appears to be novel. We hope his ideas may find further application in some of these other domains.

## References

1. Giacomo Bernardi, Peter Buneman, and Mahesh K. Marina. Tegola tiered mesh network testbed in rural Scotland. In *WiNS-DR '08: Proceedings of the 2008 ACM workshop on Wireless networks and systems for developing regions*, pages 9–16. ACM, September 2008. <http://dl.acm.org/citation.cfm?id=1410067>.
2. Philip J. Clark and Francis C. Evans. Distance to nearest neighbor as a measure of spatial relationships in populations. *Ecology*, 35(4):445–453, October 1954. <http://www.jstor.org/stable/10.2307/1931034>.
3. Jorg Colberg. *Parallel Supercomputer Simulations of Cosmic Evolution*. PhD thesis, LudwigMaximiliansUniversität München, 1999. [http://edoc.ub.uni-muenchen.de/272/1/Colberg\\_Joerg.pdf](http://edoc.ub.uni-muenchen.de/272/1/Colberg_Joerg.pdf).
4. Michael Connor and Piyush Kumar. Fast construction of  $k$ -nearest neighbor graphs for point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):599–608, 2009. [http://compgeom.com/~piyush/papers/tvcg\\_stann.pdf](http://compgeom.com/~piyush/papers/tvcg_stann.pdf).
5. Nassan Lev-Tov and David Peleg. Exact algorithms for base station placement problems. In Martti Penttonen and E. Meineche Schmidt, editors, *Algorithm Theory - SWAT 2002*, volume 2368 of *LNCS*, pages 90–99. Springer-Verlag, July 2002. <http://dl.acm.org/citation.cfm?id=645901.672620>.
6. D. O. Loftsgaarden and C. P. Quesenberry. A nonparametric estimate of a multivariate density function. *Annals of Mathematical Statistics*, 36(3):1049–1051, June 1965. [http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf\\_1&handle=euclid.aoms/1177700079](http://projecteuclid.org/DPubS/Repository/1.0/Disseminate?view=body&id=pdf_1&handle=euclid.aoms/1177700079).
7. J. J. Monaghan. Smoothed particle hydrodynamics. *Annual review of astronomy and astrophysics*, 30:543–574, 1992. <http://www.annualreviews.org/doi/pdf/10.1146/annurev.aa.30.090192.002551>.
8. The Royal Society of Edinburgh. Digital Scotland, October 2010. [http://www.royalsoced.org.uk/cms/files/advice-papers/inquiry/digitalscotland/Digital%20Scotland%20\(med%20res\).pdf](http://www.royalsoced.org.uk/cms/files/advice-papers/inquiry/digitalscotland/Digital%20Scotland%20(med%20res).pdf).
9. W. E. Schaap and R. van der Weygaert. Continuous fields and discrete samples: reconstruction through delaunay tessellations. *Astron. Astrophys.*, 363:L29–L32, 2000. <http://www.astro.rug.nl/~weygaert/tim1publication/dtfeaaletter.pdf>.
10. David P. Williamson and David B. Shmoys. *The design of approximation algorithms*. Cambridge University Press, June 2011. <http://www.designofapproxalgs.com/book.pdf>.