Energy-Aware Stage Illumination

Friedrich Eisenbrand
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany
eisen@mpi-sb.mpg.de

Stefan Funke *
Computer Science Department
Gates Bldg. 375
Stanford University, CA 94305, U.S.A.
sfunke@stanford.edu

Andreas Karrenbauer
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany
karrenba@mpi-sb.mpg.de

Domagoj Matijevic
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany
dmatijev@mpi-sb.mpg.de

ABSTRACT

Consider the following illumination problem: given a stage represented by a line segment L and a set of lightsources represented by a set of points S in the plane, assign powers to the lightsources such that every point on the stage receives a sufficient amount – let's say one unit – of light while minimizing the overall power consumption. By assuming that the amount of light arriving from a fixed lightsource decreases rapidly with the distance from the lightsource, this becomes an interesting optimization problem.

We propose to reconsider the classical illumination problems as known from computational geometry literature (e.g. [12]) under this light attenuation model. This paper examines the simple problem introduced above and presents different solutions, based on convex optimization, discretization and linear programming, as well as a purely combinatorial approximation algorithm. Some experimental results are also provided.

Categories and Subject Descriptors

F.2.2 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms

*supported by the Max Planck Center for Visual Computing and Communication (MPC-VCC) funded by the German Federal Ministry of Education and Research (FKZ 01IMC01).

Keywords

illumination problem, art gallery problems, optimization

1. INTRODUCTION

Illumination and guarding problems have been a popular topic of study in Mathematics and Computer Science for several decades. One instance in this class of problems is the classical one posed by Victor Klee : *How many guards are necessary, and how many are sufficient to patrol the paintings and works of art in an art gallery with n walls ?* While this particular problem has been solved shortly after by Chvatal proving a tight $\lfloor \frac{n}{3} \rfloor$ bound, many other variants in this problem class have appeared in the literature, e.g. [3, 4, 7, 5, 8, 10, 11]; also see [12] for a general survey of the topic.

On one hand, people have restriced the allowable 'floor plans', i.e. special classes of polygons like orthogonal polygons, or looked at the problem of guarding a set of buildings from the outside. Kahn et al. have shown for example [7], that any orthogonal polygon with n vertices can be guarded with $\lceil \frac{n}{4} \rceil$ guards, and $\lceil \frac{n}{4} \rceil$ are sometimes necessary. Fejes Toth [5] has shown that for any family $\{S_1, \ldots, S_n\}$ of n disjoint compact convex sets in the plane, one can illuminate the boundaries of the sets by $4n-7$ lightsources in the complement of $S_1 \cup \cdots \cup S_n$ and sometimes that many lightsources are necessary. Common to these results is the fact that they assume that guards/lightsources cover a 360^o field of view, and distance does not affect guarding/illumination abilities.

So other people have come up with models for less powerful guards and lightsources, for example by requiring the guards to be placed at the vertices or edges of the polygon. Another restriction is to limit the field of view of the guards to an angle of 180^o, or incorporate the used field of view/ illumination angle of the lights/guards into the objective to be optimized. For example Lee and Lin [8] have shown that finding the *minimum* number of vertex guards for a polygon is NP-hard. Toth ([11]) has shown that $\lfloor \frac{n}{3} \rfloor$ lightsources with illumination angle π suffice to illuminate any polygon with n vertices (lightsources need not be placed on vertices of the polygon). Given a line segment L (the stage) and a set of n points $p_1, \ldots, p_n$ (lightsources), Czyzowicz et al. ([4]) have proved that it is possible to find in $O(n \log n)$ time a set

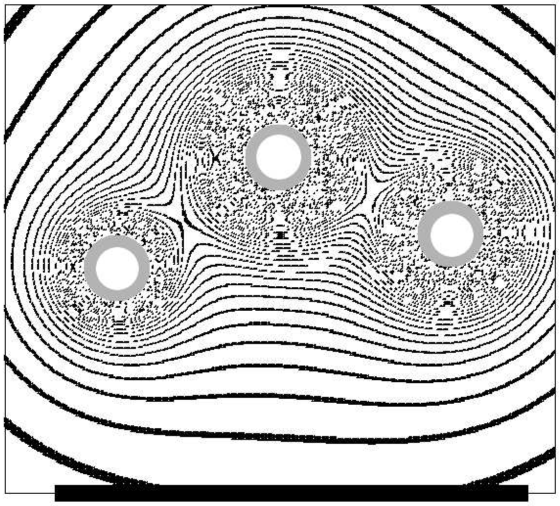

Figure 1: Isolines of light intensities induced by 3 light sources.

of floodlights $f_1, \dots f_n$ with apexes in the set $\{p_1, \dots, p_n\}$ and angles of illumination $\alpha_1, \dots \alpha_n$ such that the stage L is illuminated and the sum $\alpha_1 + \dots + \alpha_n$ is minimized.

The model we propose in this paper is in the spirit of the angle restriction employed by Czyzowicz et al., since we also aim to allow only less powerful lightsources. But while Czyzowicz et al. disallow omnidirectional lightsources (modelling floodlights in the real world), we take into account that the light emitted from a lightsource spreads with increasing distance, so the amount of light arriving at a fixed area patch decreases with the distance from the lightsource (everybody can observe this behavior on a simple light bulb, see also Figure 1, where we have sketched the isolines of light intensities induced by 3 light sources). The rationale behind our model is that it seems rather unrealistic for a guard or a lightsource to monitor/illuminate things that are arbitrarily far away, even in the absence of obstructions.

Our contribution

This paper proposes to reconsider the large collection of classical illumination problems under a light attenuation model, where the amount of light arriving from a particular lightsource decreases rapidly with the distance. As a first example, we consider the simple problem of illuminating a stage using a fixed set of lightsources, where the goal is to minimize the total amount of power assigned to the lightsources while ensuring a sufficient illumination of the stage. Several approaches in decreasing order of weight of the employed machinery are presented, namely

- a polynomial-time solution based on a convex programming formulation
- a $(1+\epsilon)$ approximate solution based on a discretization and linear programming
- a purely combinatorial $O(1)$ approximate solution with running time $O(n^2)$

We also present some experimental results suggesting that the performance analysis of the combinatorial algorithm is overly pessimistic, leaving an improved analysis and the consideration of other illumination problems in this model as open problems.

2. PRELIMINARIES

Consider the following setting: We are given a closed line segment $L \subset \mathbb{R}^2$ and a set of points $S \subset \mathbb{R}^2$, $|S| = n$. L denotes the *stage*, S a set of lightsources. Our goal is to assign powers x_s to each lightsource $s \in S$ such that any point of the stages receives a 'sufficient' amount of light – we will be more precise about that after quickly introducing the physical light model.

2.1 The physical model

For the physical model we consider the setting in three dimensions, treating the lightsources as points that emit their energy isotropically. Thereby, the energy that hits concentric spheres around the lightsource is always the same but its density decreases with growing radius. Since the energy is homogeneously distributed over the surface of such a sphere we get

$$E = \int_0^{2\pi} \int_0^{\pi} \frac{E}{4\pi} \sin\theta d\theta d\varphi$$
$$= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \underbrace{\frac{E}{2\pi} \cdot \frac{z}{\sqrt{x^2+y^2+z^2}^3}}_{F(x,y,z)} dxdy$$

where the integrand is the flux F through an infinitesimal patch on a plane at distance z from the lightsource. This is not a contradiction to the commonly known $\frac{1}{d^2}$ dependence for the intensity of a point lightsource since the latter counts for beams perpendicular to the patch. If we rotate the patch by an angle α orthogonally to the incident beam, we have to multiply the intensity by $\cos\alpha = z/\sqrt{x^2+y^2+z^2}$ yielding the same result as above.

Note that F is an additive quantity, i.e. its value can be expressed by a sum over all lightsources. We shall use reduced units such that for a point p on the stage at distances $d(p,s)$ from each lightsource $s \in S$ we have a requirement of 1 while the supply is expressed in the form $F(p) = \sum_{s\in S} \frac{x_s}{d^\sigma(p,s)}$. We may choose $\sigma = 2$, if the size of the stage is small with respect to the distance of the lightsources to the plane in which the stage is embedded, i.e. if $\cos\alpha$ is nearly 1 for all lightsources. Otherwise, we set $\sigma = 3$ and implement the distance z_s in the variable for the power x_s of each lightsource. Moreover, we scale all distances such that the minimal value of all z_s is 1.

Remark: In case of an illumination problem, it is intuitive to actually add up the arriving light/energy from all lightsources when considering some point $p \in L$. Unfortunately this cannot easily interpreted in the context of a guarding problem. If there are two guarding cameras watching for example an expensive painting in a museum, but due to their distance and limited resolution, each of the cameras can only tell with 50% confidence whether there is someone near the painting, this does not mean that by using both cameras, one can tell with 100% confidence what is happening near the painting. So, individual 'confidence ratings' do not simply add up; a reasonable model could in this concrete example assign a confidence rate of 75%, since some information might be gained by using both cameras instead of only one. Note though, that the attenuation model introduced does make sense also in this interpretation. If an

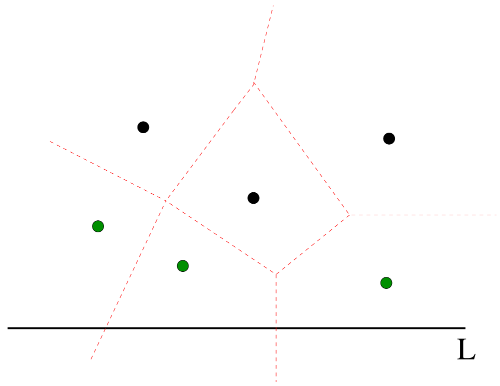

Figure 2: ***Only lights whose Voronoi cells intersect the stage are useful***

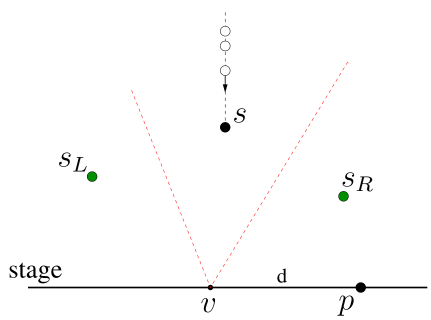

Figure 3: ***Lightsource moved towards the stage until*** $|sv| = |s_L v| = |s_R v|$

object doubles its distance to the camera, it covers only a quarter of a digital camera's CCD pixels or film emulsion. So in the above attenuation model, we would have an exponent of $\sigma = 2$.

3. ALGORITHMS IN $\mathbb{R}^2$

In this section we propose several ways to solve the problem in $\mathbb{R}^2$. The approaches will employ less and less heavy machinery, starting with a convex programming formulation, going over a combination of discretization and linear programming, to finally presenting a very simple combinatorial algorithm. In spite of the derivation of an attenuation exponent of $\sigma = 3$ in the previous section, we will assume in the following any exponent $\sigma \geq 2$, i.e. a point p receives a $\frac{1}{|ps|^\sigma}$ fraction of the light emitted from lightsource s. To allow for a simpler presentation, most calculations and proofs will be in terms of $\sigma = 2$, though generalization for larger (but constant) values of σ are straightforward.

Before presenting these algorithms we first make a simple observation which allows us to reduce the number of lightsources that have to be considered for the following steps.

3.1 Pruning lightsources

In this part we show that under the assumption that lightsources can be assigned arbitrarily high powers, only certain lightsources are of interest for our problem. Namely, we show that all lightsources whose Voronoi cell does not intersect the stage L can be replaced by lightsources whose Voronoi cells do, without incurring a larger cost in terms of the overall power used (see Figure 2). Let us state this claim more formally in the following lemma.

LEMMA 3.1. *Consider the order of the lightsources in S induced by the vertical projection on the line supporting the stage L. Let $s \in S$ be some lightsource whose Voronoi cell does not intersect the stage, $s_L, s_R \in S$ be the first neighbors to the left and right in the ordering above whose Voronoi cells intersect the stage.*

Then there always exists a power assignment x_{s_L} and x_{s_R} such that for any $p \in L$

$$\frac{x_{s_L}}{|s_L p|^2} + \frac{x_{s_R}}{|s_R p|^2} \geq \frac{x_s}{|sp|^2} \tag{1}$$

and $x_{s_L} + x_{s_R} \leq x_s$.

PROOF. In the following we will exhibit a power assignment x_{s_L}, x_{s_R} with $x_{s_L} + x_{s_R} = x_s$. Thus, we can express x_{s_L} and x_{s_R} as $x_{s_L} = \alpha \cdot x_s$ and $x_{s_R} = (1-\alpha) \cdot x_s$ for some nonnegative $\alpha \leq 1$. So we can rewrite (1) as

$$\frac{\alpha}{|s_L p|^2} + \frac{(1-\alpha)}{|s_R p|^2} \geq \frac{1}{|sp|^2}, \quad \alpha \leq 1. \tag{2}$$

Now the goal is to show that there exists an $\alpha \leq 1$ independent of the position of the point $p \in L$ and such that inequality (2) holds.

Let v denote the intersection of the Voronoi edge between s_L and s_R and the stage L. Note that we can always move the light-source s perpendicularly toward the stage until $|s_L v| = |sv| = |s_R v|$ (see Figure 3) since this is the worst case scenario for the claim of the lemma (it's easier to replace far away lightsources). Then one can observe that in the case when p lies to the left of v, $|s_L p| < |sp| < |s_R p|$ and analogously when p lies to the right of v, $|s_L p| > |sp| > |s_R p|$. For the sake of simplicity suppose $|s_L v| = |sv| = |s_R v| = 1$ and let $d = |vp|$, $\phi_1 = \angle pvs_R$, $\phi_2 = \angle pvs$ and $\phi_3 = \angle pvs_L$. We can express the distances $|s_L p|$, $|sp|$ and $|s_R p|$ with help of the law of cosines and obtain:

$$\frac{\alpha}{1 + d^2 \pm 2d\cos\phi_3} + \frac{(1-\alpha)}{1 + d^2 \pm 2d\cos\phi_1}$$

$$\geq \frac{1}{1 + d^2 \pm 2d\cos\phi_2}$$

and hence

$$\alpha \cdot \left(\frac{\pm(\cos\phi_1 - \cos\phi_3)}{|s_L p|^2}\right)$$

$$\geq \frac{\pm(\cos\phi_1 - \cos\phi_2)}{|sp|^2} \tag{3}$$

were '+' holds if p lies to the left of v and '-' if p lies to the right of v. Choosing $\alpha = \frac{\cos\phi_1 - \cos\phi_2}{\cos\phi_1 - \cos\phi_3} \leq 1$ for $0 \leq \phi_1 < \phi_2 < \phi_3 \leq \pi$ and keeping in mind that $|s_L p| > |sp|$ if p lies to the right of v and $|s_L p| < |sp|$ if p lies to the left of v, one can easily verify inequality (3) and therefore conclude the proof of the lemma. □

3.2 A convex programming formulation

The following convex program clearly solves our problem:

$$\begin{array}{llll} \min & \sum_{s \in S} x_s & & \\ s.t. & \forall p \in l: & \sum_{s \in S} x_s / d^\sigma(p,s) & \geq 1 \\ & & x_s & \geq 0 \end{array} \quad (4)$$

Here the second line exactly expresses the constraint that for every point p on the stage, when summed over all lightsources, 'enough' light should arrive at p. If lightsource s is powered up with x_s, the fraction of light arriving at p is proportional to $1/d^\sigma(p,s)$. Here σ is the attenuation exponent as derived to be $\sigma = 3$ in the previous section or $\sigma = 2$ as commonly used. Note, that in this formulation one could also incorporate *upper bounds* on the light intensity. Later we will refer to this convex program when considering only a finite number of constraints (and hence being a linear program) as *lighting LP*.

This formulation is not a linear program since the number of constraints is (uncountably) infinite, so in fact our setting is an optimization problem over a convex body rather than a simple linear program. There are numerous algorithms for optimizing over a convex body, most of which rely on an efficient method of determining whether a point $x \in \mathbb{R}^n$ is contained in the convex body, which in our case basically reduces to determining whether some degree n polynomial has a root. In the following we describe the method in detail for $\sigma = 2$. The following notions can be found in [6]. Let $K \subseteq \mathbb{R}^n$ be a convex body, $\epsilon > 0$. The set $S(K, \epsilon)$ is the set of points which have at most distance ϵ from K, $S(K, \epsilon) = \{x \in \mathbb{R}^n \mid \|x - y\| \leq \epsilon \text{ for some } y \in K\}$. The set $S(K, -\epsilon)$ is the set of points in K whose ϵ-environment is completely contained in K, $S(K, -\epsilon) = \{x \in K \mid \|x - y\| \leq \epsilon \text{ implies that } y \in K \text{ for all } y \in \mathbb{R}^n\}$.

We now recall the definitions of the weak membership- and optimization problem over convex bodies, see [6].

Definition 3.1. *The* weak membership problem *for K is the following:*

> *Given a vector $x \in \mathbb{Q}^n$ and a rational number $\delta > 0$, either*
>
> 1. *assert that $x \in S(K, \delta)$, or*
> 2. *assert that $x \notin S(K, -\delta)$.*

Definition 3.2. *The* weak optimization problem *for K is the following: Given a vector $c \in \mathbb{Q}^n$ and a rational number $\epsilon > 0$, either*

1. *find a vector $x^* \in \mathbb{Q}^n$ such that $x^* \in S(K, \epsilon)$ and $c^T x \leq c^T x^* + \epsilon$ for all $x \in S(K, -\epsilon)$, or*
2. *assert that $S(K, -\epsilon)$ is empty.*

Grötschel, Lovász and Schrijver [6, Corollary (4.3.12)] prove the following theorem.

Theorem 3.1. *There exists an oracle polynomial time algorithm that solves the weak optimization problem for every convex body K, given by a weak membership oracle, where the convex body contains a ball of radius r around a point a_0 and is contained in a ball or radius R around 0.*

Polynomial time here means, polynomial in the dimension n and the binary encoding lengths of c, ϵ, a_0, r and R.

Observe that the feasible region of the system (4) is not bounded. However, we can easily compute a bounding parameter M, such that an optimal solution is contained in $0 \leq x \leq M$. We simply let M be the largest power value, which has to be assigned to a single lightsource in order to lighten the stage by itself. If we then impose the additional constraint $0 \leq x \leq 2M$ to the system (4), the convex set is bounded and contained in the ball around 0 with radius $2Mn$ and contains the ball around $M\mathbf{1}$ with radius M. In the following we denote the set of feasible solutions by K.

Next we show that the weak membership problem for a power assignment x' can be solved in polynomial time. For this we assume that the lightsources are located on the Euclidean plane with nonnegative component in the y-axis and that the stage is the interval $[-L, L]$ on the x-axis. Suppose we are given a power assignment x'. The exact membership problem is to decide whether there exists a point $(p, 0)$ on the stage such that $\sum_{s \in S} x_s / \left((p - X_s)^2 + Y_s^2\right) < 1$ holds.

We solve the weak membership problem for any $\epsilon > 0$ in the following way. We decide whether there exists a $p \in [-L, L]$ such that $\sum_{s \in S} x_s / \left((p - X_s)^2 + Y_s^2\right) = 1$ holds. If yes, we can assert that $x' \notin S(K, -\delta)$ for each $\delta > 0$. Otherwise, we determine whether $\sum_{s \in S} x_s / \left((L - X_s)^2 + Y_s^2\right) > 1$ holds. If yes, we can assert that $x' \in K$. If not, we can assert that $x' \notin K$.

Thus we can solve the weak membership problem for K if we can determine in polynomial time, whether there exists a $p \in [-L, L]$, such that the following holds:

$$\sum_{s \in S} x_s / \left((p - X_s)^2 + Y_s^2\right) = 1 \quad (5)$$

Equation (5) can be written as $f(p) = 0$, where $f(p)$ is a rational polynomial, whose binary encoding length is polynomial in the encoding length of the positions of the lightsources. The problem now reads as follows. Given a polynomial $f(p) \in \mathbb{Q}(p)$ and an integer L, determine, whether $f(p)$ has a root in $[-L, L]$. This can be done in polynomial time, after $f(p)$ is decomposed into squarefree factors, with the method of *Sturm*, see, e.g.,[13, p. 87]. So the weak optimization problem (4) can be solved in polynomial time in the *encoding length* of the lightsource placements and the *encoding length* of the error parameter ϵ.

Theorem 3.2. *Given a set S of lightsources in the Euclidean plane and an $\epsilon > 0$, one can compute a feasible point x^* for the optimization problem (4) in polynomial time such that $\sum_{s \in S} x_s^* \leq \sum_{s \in S} \bar{x}_s + \epsilon$ for any feasible $\bar{x}$.*

Thus the most energy efficient illumination can be approximated with an *additive error* $\epsilon > 0$ in polynomial time.

3.3 A $(1 + \epsilon)$ approximation scheme

One obvious approach to obtain an approximation to our problem is to discretize the stage using a finite number of *guards* [1], solve the linear program with constraints only for the guards and then power up all lightsources sufficiently

[1]Note that in the following we use the term *guard* as a point on the stage that ensures sufficient lighting at that point. That notion differs from the use of *guard* in other work in that area, where the guard is a point which covers/watches the scene.

such that all points on the stage which were not 'guarded' by a constraint for sure also get enough light. The efficiency of this approach depends on the choice of a suitable discretization which allows for few guards but still requires only a moderate 'power up' of the lightsources to guarantee sufficient overall coverage.

DEFINITION 3.3. *For every point $p \in L$, we define ens(p) – the* empty neighborhood size *– to be the distance to the closest lightsource, i.e. ens*$(p) = \min_{s \in S} d(p, s)$.

The following observation is not hard to see since ens(.) is defined to be the minimum over some distance functions.

OBSERVATION 3.1. *ens(.) is* 1-Lipschitz*, that is ens*$(p) \leq$ *ens*$(q) + 1 \cdot |pq|$.

Our discretization will now be based upon the empty neighborhood size, in particular we will have more guards in areas where ens(.) is small and fewer guards in areas where ens(.) is large. Similar discretization approaches occur in several other places in literature: For example, Amenta et al. [2] use the so-called *local feature size* to classify discrete samples from a contiuous surface. Papadimitriou and Aleksandrov et al. also use a related discretization for the purpose of shortest path computations [9, 1].

The crucial property for our set G of chosen guards is the following:

DEFINITION 3.4. *A set $G \subset L$ of points satisfying*

$$\forall p \in L\ \exists g \in G : d(p, g) \leq \epsilon \cdot ens(p)$$

is called a ϵ-good set of guards.

Note that assuming a minimum distance of 1 of each lightsource to the stage, it is trivial to obtain an ϵ-good set of guards of size D/ϵ by placing guards at equal distance ϵ all along the stage. Here D denotes the length of the stage L. In the following we will show that one can do considerably better.

Before we show that using an ϵ-good set G we can obtain a $(1 + \epsilon)$ approximation to our original problem, let us first convince ourselves that a reasonably small set G of guards exists. For that consider one lightsource s. Assuming that we have pruned the set of lightsources according to the previous section, there is a point $p_0 \in L$ for which s is the closest lightsource. We start constructing a set G_s by first adding p_0 to G_s. We then extend G_s by adding guards $p_{-1} \in L$ (for first guard left of p_0) and $p_{+1} \in L$ (first guard right of p_o) at distance $2\epsilon\text{ens}(p_0)$ from p_0. Iteratively we place the next guard p_{i+1} at distance $2\epsilon\text{ens}(p_i)$ from p_i (and accordingly to the left). We now claim the following:

LEMMA 3.2. *The set G_s constructed above is a ϵ-good set of guards for the single lightsource s and furthermore $|G_s| = O(\frac{\log D}{\epsilon})$ where D denotes the length of the stage.*

PROOF. Assume the contrary, i.e. there exists a point $p \in L$ s.t. $\nexists p_i \in G_s$ with $d(p_i, p) \leq \epsilon\text{ens}(p)$. W.l.o.g. assume p lies between p_i and p_{i+1} (the same argumentation holds when it lies between p_{-i} and p_{-i-1}). We have $\forall p'$ between p_i and p_{i+1}: $\text{ens}(p') \geq \text{ens}(p_i)$, since we are moving away from the lightsource. Furthermore we have clearly $\min\{d(p, p_i), d(p, p_{i+1})\} \leq |p_i p_{i+1}|/2$. But since $|p_i p_{i+1}| \leq 2\epsilon\text{ens}(p_i)$ by construction we get $\min\{d(p, p_i), d(p, p_{i+1})\} \leq \epsilon\text{ens}(p_i) \leq \epsilon\text{ens}(p)$ which contradicts our assumption.

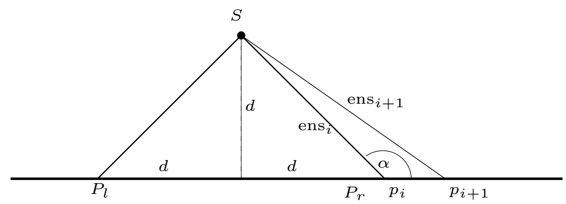

Figure 4: *Bounding the number of guards for S*

Let us now turn to the size of G_s. Look at the situation in figure 4. Clearly for all $p \in \overline{P_l P_r}$ we have $\text{ens}(p) \geq d$, hence the distance between two adjacent guards between P_l and P_r is at least $2\epsilon d$, hence there are at most $O(1/\epsilon)$ many guards placed at that part of the stage.

We are now interested in the guards outside $\overline{P_l P_r}$. We claim that for consecutive guards p_i, p_{i+1} we have $\text{ens}(p_{i+1}) \geq \text{ens}(p_i) \cdot (1 + \epsilon)$. This follows easily from the law of cosine since we have $\text{ens}(p_{i+1})^2 = \text{ens}(p_i)^2 + (2\epsilon\text{ens}(p_i))^2 - 2\text{ens}(p_i) \cdot 2\epsilon\text{ens}(p_i)\cos(\alpha) \geq \text{ens}(p_i)^2[1 + 4\epsilon^2 + 2\epsilon] \geq \text{ens}(p_i)^2(1 + \epsilon)^2$ where the last two inequalities follow from the fact that $\alpha \geq \frac{3}{4}\pi$. Hence the distance between adjacent guards outside $\overline{P_l P_r}$ grows at least by a factor of $a = (1 + \epsilon)$ in each iteration. We now establish an upper bound on the number of guards in terms of this factor a.

$$\begin{aligned} D &\geq 2\epsilon \sum_{i=1}^{|G_s|} a^i = 2\epsilon a \frac{a^{|G_s|} - 1}{a - 1} \\ \Rightarrow \quad a^{|G_s|} &\leq \frac{D}{2\epsilon} \cdot \frac{a-1}{a} + 1 \\ \Rightarrow \quad |G_s| &\leq \frac{\log(\frac{D}{2\epsilon} \cdot \frac{a-1}{a} + 1)}{\log a} \end{aligned}$$

Since the number of guards contained in $\overline{P_l P_r}$ is also $O(1/\epsilon)$ we can conclude that the total number of guards generated by our procedure is $O(\frac{\log D}{\epsilon})$ □

For a fixed length D of the stage, this estimation is tight as can be seen from the definition 3.4. It establishes an upper bound on the distance of two guards. Having a stage of length D and a lightsource at distance 1 from the stage, ens($\cdot$) is at most $\sqrt{1 + D^2}$. Therefore, we have to partition the stage into at least $\Omega(\frac{1}{\epsilon})$ parts.

Obtaining a set of ϵ-good guards G could be easily achieved by computing G_s for all lightsources s and taking the union of those sets. It is clear that the resulting set is ϵ-good for the set of all lightsources, since for any $p \in L$ there is a guard within distance $\epsilon \cdot \text{ens}(p)$ in the set G_s where s is the lightsources closest to p. The resulting union then contains $O(\frac{n}{\epsilon} \log D)$ guards. We can do better though:

LEMMA 3.3. *There exists a ϵ-good set of guards G of size $O(\frac{n}{\epsilon} \log[1 + \frac{D}{n}])$.*

PROOF. The idea is that we may consider each Voronoi cell of the lightsources on its own since the ens($\cdot$) of its points on the stage are determined by their distance to this particular lightsource. The stage is partitioned into n pieces $D_1, \ldots, D_n$ corresponding to the respective Voronoi regions of the light sources. Let d_i denote the length of piece D_i, i.e. $\sum_{i=1}^{n} d_i = D$. For each piece D_i we construct the sample set as before and get overall $O((\sum_{i=1}^{n} \log[1 + d_i])/\epsilon)$ many guards. This sum is maximized when all parts have equal length, i.e. $d_i = \frac{D}{n}$. □

Note that this bound considered asymptotically with respect to n tends to $O(\frac{D}{\epsilon})$, since $n \log[1 + \frac{D}{n}] = \log[1 + \frac{D}{n}]^n \longrightarrow \log e^D$, i.e. in the limit the number of guards does not depend on n anymore, but only on D (but linearly, not in the logarithm).

The last lemma in this section shows that given a ϵ-good set of guards, we can use this to obtain a $(1+\epsilon)$-approximate solution to the lighting problem without having to worry about the infinite number of constraints. But before proving that, we show a small auxiliary Lemma which gives an upper bound on the distance between two consecutive guards in a ϵ-good set of guards.

LEMMA 3.4. *Let G be an ϵ-good set of guards, p and q two guards in G that appear consecutively on the stage. Then $|pq| \leq \frac{2\epsilon}{1-\epsilon} ens(p)$.*

PROOF. Let $z \in L$ such that $|pz| = |zq|$. By the definition 3.4 we know that $|zq| \leq \epsilon \cdot \text{ens}(z)$. Using the 1-Lipschitz property of the local feature size $\text{ens}(z)$ we can write down $|pq| = 2|zq| \leq 2\epsilon \cdot (\text{ens}(q) + |pq|/2)$ which in turn implies the claim of the lemma. □

LEMMA 3.5. *Let $\{x_s\}$ be an optimal solution of the lighting LP (4) with respect to an ϵ-good set of guards G. Then powering up every lightsource by a factor $(1 + 6\epsilon)$ ensures that* every *point on the stage receives enough light.*

PROOF. For the power assignment $\{x_s\}$ we know that for all $p' \in G$, the lighting constraints are fulfilled, i.e. they receive enough light. Consider some point $p \in L$, $p \notin G$. Let $p' \in G$ be the closest guard to p, hence $|pp'| \leq \frac{\epsilon}{1-\epsilon}\text{ens}(p')$ according to Lemma 3.4. We want to show that after powering up all lightsources by a sufficiently large factor $\psi = 1+O(\epsilon)$, p also receives enough light. Namely, we are looking for ψ such that

$$\sum_{s \in S} \frac{\psi \cdot x_s}{(|sp'| + \text{ens}(p') \cdot \epsilon/(1-\epsilon))^2} \geq 1 \qquad (6)$$

Observe that all lightsources have distance at least $\text{ens}(p')$ to p' just by definition of $\text{ens}(\cdot)$ and keeping in mind that p' receives enough light, inequality (6) holds if ψ is chosen such that:

$$\sum_{s \in S} \frac{\psi \cdot x_s}{|sp'|^2 (1 + \epsilon/(1-\epsilon) \cdot \underbrace{\text{ens}(p')/|sp'|}_{\leq 1})^2} \geq \frac{\psi}{(1 + \epsilon/(1-\epsilon))^2} \geq 1$$

Therefore, for $\epsilon < 1/2$, powering up all lightsources by a factor of $\psi = (1 + 2\epsilon)^2 < 1 + 6\epsilon$ makes sure that p receives at least as much light as p' received before powering up all lightsources. □

We summarize by stating the main theorem of this part:

THEOREM 3.3. *Given a stage of length D and a set of lightsources S where each lightsource has at least unit distance from the stage, one can compute a power assignment $\{x_s\}_{s \in S}$ such that each point on the stage receives at least 1 unit of light and $\sum x_s \leq (1 + \epsilon) \sum x_s^{\text{opt}}$ where x_s^{opt} denotes an optimal power assignment. $\{x_s\}_{s \in S}$ can be found in polynomial time by solving a linear program with $O(\frac{n}{\epsilon} \log \frac{D}{n})$ constraints and n variables.*

3.4 Pruning guards – a simple $O(1)$ approximation algorithm

Even though the previous section provided a rather simple $(1 + \epsilon)$ approximation algorithm for our problem, it relied on solving a linear program which – in spite of being polynomial-time – is still quite time-consuming (at least in theory). Furthermore there was still a – even though only logarithmic – dependence on the length D of the stage. In the following we will propose a very simple $O(1)$-approximation algorithm that can be easily implemented to run in $O(n^2)$ time.

Similar to the previous section we will first relax our problem by restricting to a small – here $O(n)$ size – set of guards. This set is chosen such that any solution for this reduced set transfers to a solution for the original problem incurring only a $O(1)$ overhead in terms of the quality of the solution.

Consider the function ens(.) on the stage L. This continuous function consists of several arcs, each corresponding to one lightsource / their respective Voronoi cell. ens(.) is differentiable except for the positions where two adjacent arcs are joined, that is at the boundary between two Voronoi cells. ens has local maxima at all the intersection points between Voronoi edges of $\mathcal{V}(S)$ and L, and potentially at the endpoints of L – depending on the location of the left-most and right-most lightsource.

LEMMA 3.6. *Let G_V be the set of guards consisting of all points p of the stage L where ens(.) has a local maximum. Furthermore let x_V^* be a feasible power assignment to the lightsources S w.r.t. the set G_V of guards.*

Then $x^* = 4 \cdot x_V^*$ is a valid power assignment w.r.t. to all points on the stage.

PROOF. Consider any point $p \in L$, $p \notin G_V$. Let $p_V \in G_V$ be a guard such that p is contained in the circle centered at p_V with radius $\text{ens}(p_V)$. Such a guard p_V always exists since each point on the stage is contained at least in one of the Voronoi circles around its left and right neighbors in the set G_V. Obviously all lightsources have distance at least $\text{ens}(p_V)$ to p_V. But on the other hand we have $|pp_V| \leq \text{ens}(p_V)$ by choice of p_V. Therefore all lightsources satisfying p_V's demand are at most a factor 2 further away from p, hence powering up all lightsources by a factor of 4 ensures that p receives a sufficient amount of light. □

An immediate consequence is the following corollary:

COROLLARY 3.1. *A 4-approximation to the lighting problem can be obtained by solving the lighting LP consisting of $n + 1$ constraints.*

In other words, if we are only aiming for a $O(1)$ approximation, we can obtain a solution in time independent of the length of our stage D (remember in case of the $(1 + \epsilon)$ approximation we had a $\log D$ dependence on the length of the stage).

In the following we will work on this set of guards G_V as defined above. Essentially we order them according to decreasing ens(.) and one-by-one increase the power of their respective nearest lightsource such that they all get satisfied. By a primal-dual fitting argument we then show that the used amount of power does not exceed a constant times the optimum.

Our analysis relies on a special property of the set of guards, namely we want that the density of the guards is proportional to the local value of ens(.), in particular we want the distance between two adjacent guards g_i, g_j on the stage L to be lower bounded by $|g_ig_j| \geq C \cdot \max\{\text{ens}(g_i), \text{ens}(g_j)\}$ for some constant $C > 0$. This need not be the case in general, e.g. consider a set of lightsources on a line parallel to but far away from the stage. Hence we need to prune the set of guards beforehand to ensure this property.

3.4.1 Pruning Guards

Let $\alpha > 0$ be some constant. Then the following algorithm prunes a set of guards G_V to a set G_P:

1. Compute for each guards p_i its empty neighborhood size $\text{ens}(p_i)$
2. Sort the guards in decreasing order of $\text{ens}(p_i)$, i.e. $\text{ens}(p_1) \geq \text{ens}(p_2) \geq \ldots \text{ens}(p_n)$
3. for $i = 1 \ldots n$
   - if p_i has not been removed yet, remove all guards p_j at distance $\leq \alpha \cdot \text{ens}(p_i)$ (but not p_i itself)
4. return the set of guards that have not been removed as G_P

In the following we will show that for constant α, even after this pruning step, we can obtain a $O(1)$ approximation using the pruned set of guards.

LEMMA 3.7. *Let x_P^* be a feasible assignment of powers to the lightsources such that all guards in G_P are satisfied. Then $x_V^* = (1+\alpha)^2 \cdot x_P^*$ is a valid power assignment for the set of guards $G_V \supseteq G_P$.*

PROOF. Let $p_i \in G_V, \notin G_P$ be a guard that has been removed during the pruning step, $p_j \in G_P$ the guard responsible for the removal. Then we have $|p_ip_j| \leq \alpha \text{ens}(p_j)$, and hence powering up all lightsources by a factor of $(1+\alpha)^2$ ensures that p_i receives enough light due to the same reasoning as in Lemma 3.6. □

Furthermore, our desired property is obviously fulfilled:

LEMMA 3.8. *For any two guards in the pruned set $p_i, p_j \in G_P$, we have $|p_ip_j| \geq \alpha \max\{ens(p_i), ens(p_j)\}$.*

PROOF. Assume otherwise, then either p_i or p_j would have been pruned away when considering the other guard. □

An immediate corollary of Lemma 3.7 is the following:

COROLLARY 3.2. *A $4(1+\alpha)^2$ approximation to the lighting problem can be obtained by solving the lighting LP w.r.t. to the pruned set of guards G_P.*

It is now time to describe the algorithm which we will use to derive a power assignment to the lightsources. For that let us denote by s_i the lightsource that is closest to guard p_i, x_i its assigned power for all guards $p_i \in G_P$. Without loss of generality we assume that no lightsource is the closest for more than one guard (our derived bounds only get better if we remove this assumption). The algorithm works as follows:

1. Compute the set of guards G_V (via the Voronoi diagram of S)
2. Prune the set of guards G_V with pruning constant α to obtain G_P, $|G_P| = m$.
3. Let G_P be ordered such that $\text{ens}(p_1) \geq \text{ens}(p_2) \geq \cdots \geq \text{ens}(p_m)$
4. for all $i = 1 \ldots m$
   - $x_i = \max\{0, |p_is_i|^2 \cdot \left(1 - \sum_{j=1}^{i-1} \frac{x_j}{|s_jp_i|^2}\right)\}$

Informally speaking this algorithm takes the guards one-by-one in decreasing order of their ens(.) value and increases the power of their closest lightsource just sufficiently such that they receive enough light. It can be trivially implemented to run in $O(n^2)$ time.

The crux of the analysis will be to show that no guard receives more than a constant amount of excessive light. This property will then allow us to use a primal-dual fitting argument bounding the quality of our solution.

Let P_i be the amount of light experienced by some guard p_i after the execution of the algorithm. Let us write this as $P_i = P_i^< + P_i^= + P_i^>$ where $P_i^<$ denotes the power received from lightsources s_j, $j < i$, $P_i^=$ the power received from lightsource s_i and $P_i^>$ the power received from lightsources s_j, $j > i$. Clearly $P_i^= > 0 \Leftrightarrow P_i^< < 1$, that is s_i will only be used if p_i did not already receive enough light from lightsources which were switched on before in the course of the algorithm. In the following we will bound $P_i^>$ and $P_i^<$ and show that they are at most some constant ($P_i^= \leq 1$ is obvious).

LEMMA 3.9. $P_i^> \leq 4$

PROOF. Assume w.l.o.g. that all guards p_j, $j > i$ lie to the right of p_i (at the end we simply multiply the obtained bound by 2 to obtain a bound for all p_j). We have

$$P_i^> = \sum_{j=i+1}^{m} \frac{x_j}{|p_is_j|^2} \leq \sum_{j=i+1}^{m} \frac{\text{ens}(p_j)^2}{((\sum_{l=i+1}^{j} \alpha \cdot \text{ens}(p_{l-1})) - \text{ens}(p_j))^2}$$

following from Lemma 3.8 and since each lightsource s_j is at most powered up to $\text{ens}(p_j)^2$. Assuming $\alpha \geq 2$ we can continue with

$$\leq \sum_{j=i+1}^{m} \frac{\text{ens}(p_{j-1})^2}{(\sum_{l=i}^{j-1} \text{ens}(p_l))^2}$$

But this sum is of the form $\sum_{i=1}^{n} \frac{\delta_i^2}{(\sum_{j=1}^{i} \delta_j)^2}$ with $\delta_i \geq \delta_{i+1}$ (new indices here !). We then get

$$\sum_{i=1}^{n} \frac{\delta_i^2}{(\sum_{j=1}^{i} \delta_j)^2} \leq \sum_{i=1}^{n} \frac{\delta_i^2}{(i\delta_i)^2} \leq \frac{\pi^2}{6} \leq 2$$

by our decreasing odering of the δ_i.

□

Furthermore we have for the energy collected from the lightsources assigned previously in the course of the algorithm:

LEMMA 3.10. $P_i^< \leq 6$.

PROOF. Consider the first guard p_j, $j < i$ to the left of p_i whose lightsource is switched on, i.e. $x_j > 0$. Clearly we have $P_j^= + P_j^< = 1$ by definition of the algorithm. Furthermore due to the previous Lemma we know that p_j receives at most 2 units of light from the left from lightsources $k > j$. Hence at most 3 units of light can arrive at p_i from the left, which makes 6 units of light overall considering both the contributions from the left and the right. □

From these two Lemmas and the observation that $P_i^= \leq 1$ we can derive the following

Corollary 3.3. *If powers are assigned to the lightsources according to our algorithm, we have for every guard $p_i \in G_p$: $1 \leq P_i = P_i^< + P_i^= + P_i^> \leq 4 + 1 + 6 = 11$.*

which says that any guard receives between 1 and 11 units of light.

3.4.2 Bounding the Quality of the Solution

In the following we will argue that the solution x^* obtained with respect to the pruned set of guards G_P is almost optimal, i.e. only a constant factor away from the optimal solution (w.r.t. G_P). Since x^* can be easily extended to a feasible solution for the whole stage incurring an additional cost factor of at most $4 \cdot (1+\alpha)^2$ according to Corollary 3.2 and since an optimum solution with respect to some finite set of guards is always a lower bound on the optimum solution for the whole stage, we obtain the desired $O(1)$ approximation guarantee.

Let us rewrite the linear program w.r.t. the pruned set of guards G_P and lightsources S_P:

$$\begin{array}{llll} \min & \sum_{s \in S_P} x_s & & \\ s.t. & \forall p \in G_P: & \sum_{s \in S_P} x_s / d^2(p,s) & \geq 1 \\ & & x_s & \geq 0 \end{array} \tag{7}$$

The dual of this program looks as follows:

$$\begin{array}{llll} \max & \sum_{p \in G_P} y_p & & \\ s.t. & \forall s \in S_P: & \sum_{p \in G_P} y_p / d^2(p,s) & \leq 1 \\ & & y_p & \geq 0 \end{array} \tag{8}$$

The interpretation of the dual is the following: You want to assign weights y_p to each guard $p \in G_p$ such that for each lightsource, the 'influence' of the guards does not exceed 1.

To show that the power assignment constructed by our algorithm is not too far off the optimum, it suffices to exhibit a feasible solution to the dual program which has about the same objective function value. Weak duality then tells us that the optimum solution to the primal program is sandwiched between the solution of our algorithm and the feasible dual solution. Making use of the fact that the distance between a lightsource s_i and a guard g_j is essentially the same as the distance between lightsource s_j and p_i (after the α-pruning), we can use the amount of light arriving at each guard p_i as value for y_i and after scaling by a constant factor obtain a feasible solution to the dual program.

Lemma 3.11. *For any lightsource s_i and guard p_j in the pruned set of guards and lightsources G_P and S_P we have:*

$$|s_j p_i| \cdot \left(1 - \frac{2}{\alpha - 1}\right) \leq |s_i p_j| \leq |s_j p_i| \cdot \left(1 + \frac{2}{\alpha - 1}\right)$$

Proof. We show the right inequality, the left works analogously. We have by triangle inequality $|s_i p_j| \leq |p_i p_j| + \text{ens}(p_i) \leq |p_i s_j| + \text{ens}(p_i) + \text{ens}(p_j) \leq |p_i s_j| + 2E$ for $E = \max\{\text{ens}(p_i)\text{ens}(p_j)\}$. But since $E \leq \frac{|p_i s_j| + E}{\alpha}$ we get after rearranging $E \leq \frac{|p_i s_j|}{\alpha - 1}$ which yields the desired bound. □

In other words, for $\alpha \geq 3$ the distances $|s_i p_j|$ and $|p_i s_j|$ can differ by at most a factor of two.

Lemma 3.12. *Let x^* be the solution to the primal LP (7) as computed by our algorithm, c_p^* its objective function value. Then there exists a feasible solution y^* to the dual LP (8) with function value $c_d^* \geq c_p^*/22$.*

Proof. Let us set for every $p_i \in G_P$: $y_i^* = x_i^*/22$. We need to verify that $\forall s \in S_P : \sum_{p \in G_P} y_p^*/|ps|^2 \leq 1$. But according to Corollary 3.3 and together with Lemma 3.11 we have for $\alpha \geq 3$:

$$\sum_{j=1}^{m} \frac{y_j^*}{|p_j s_i|^2} \leq \frac{2^2}{22} \left(\sum_{j=1}^{m} \frac{x_j^*}{|p_i s_j|^2} \right) \leq 1$$

□

So we have established a dual feasible solution with function value at least $c_p^*/22$, i.e. the optimum value c_{opt} must lie between $c_p^*/22$ and c_p^* which implies that x_p^* is no worse than a 22-approximation for the LP w.r.t. the pruned set of guards. And because an optimum solution w.r.t. the pruned set of guards can be extended to a feasible solution for the original problem at a cost of an additional $4 \cdot (1+\alpha)$ factor, we conclude with the following main theorem of this section:

Theorem 3.4. *Given a stage L and n lightsources, one can compute in $O(n^2)$ time a power assignment to the lightsources such that any point on the stage receives at least one unit of light. The solution produced requires at most $O(1)$ times the optimal amount of energy.*

3.4.3 Questions

It is not clear whether the pruning is indeed necessary for the analysis of the algorithm. As the experiments later on show, even without pruning the algorithm achieves a rather good approximation ratio, so it might be possible that the pruning was only necessary due to our inability to give a more precise analysis.

4. GENERALIZATIONS AND OPEN PROBLEMS

In the following we will list some extensions and other open illumination problems that are worth analysing within our light attenuation model.

4.1 Generalization to higher dimensions ($\mathbb{R}^3$)

There is a straightforward way of extending the described model to a 3-dimensional setting by assuming the stage L to be some bounded two-dimensional surface patch in $\mathbb{R}^3$. The definition of ϵ-good sampling 3.4 can still be used, and ϵ-good sample sets can be derived in a similar manner as described for the 2-dimensional case. Their sizes then depend on the area of the two-dimensional surface to be sampled (again, a logarithmic instead of linear dependence is achievable). The LP-based solution strategy can still be applied, whereas the constant approximation probably requires some more work.

4.2 Open Problems

There is a vast number of variations of the basic illumination/guarding problems that have been considered in the past. Many of them can also be considered in our light attenuation model, for example:

	$LP_{1+\epsilon}$	LP_V^*	C_0^*
$D = 70$, $\lvert G_V\rvert = 6$	390.97	1547.96	39776
$D = 70$, $\lvert G_V\rvert = 15$	106.03	416.04	11671.44
$D = 70$, $\lvert G_V\rvert = 22$	82.28	324.12	8883.60
$D = 140$, $\lvert G_V\rvert = 8$	690.88	2507	66584.32
$D = 140$, $\lvert G_V\rvert = 19$	314.59	1227.44	33846.56
$D = 140$, $\lvert G_V\rvert = 41$	172.40	675.96	20138.80

Table 1: Energy costs for different stage lengths (D), number of guards ($|G_V|$), and algorithms (LP-based $(1 + \epsilon)$-approximation, LP-based 4-approximation, combinatorial $O(1)$-approximation).

Art Gallery Illumination with k Lightsources

Given some fixed number k, determine position and power assignments of k lightsources such that any point on the boundary (or also in the interior) of a polygon with n vertices receives at least 1 unit of light. Minimize the sum of assigned powers.

Floodlight Illumination

Given a stage L and a set $F = \{f_1, \ldots, f_n\}$ of floodlights of angle sizes $\alpha_1, \ldots, \alpha_n$ such that their apexes are located at some fixed points on the plane, all on the same side of L. Decide if it possible to rotate them such that every point on L is illuminated, and if yes, determine the rotations and power assignments, such that the stage is sufficiently illuminated at every point and the overall power assigned is minimized.

Stage Illumination with Obstacles

The same problem as considered in this paper can be examined also in the presence of obstacles. In this case, neither the pruning of lightsources nor the discretization can be applied immediately, though some similar approach seems doable.

5. EXPERIMENTAL RESULTS

In this experimental section we want to investigate the actual behaviour of the proposed algorithms, since we believe that our analysis of the approximation ratios is overly pessimistic. We have implemented the LP-based approximation algorithms using a ϵ-good set of guards and the local maxima of the ens(.) function, as well as the combinatorial $O(1)$ approximation algorithm based on the pruned set of Voronoi vertices. We have run the algorithms for different lengths D of the stage also varying the number of relevant (in the sense of Section 3.1) lightsources (which were randomly generated around the stage).

Performance according to the Analysis

In Table 1 we have listed the sum of the power assigments made by our algorithms for varying values of D and $|G_V|$ (the number of light sources whose Voronoi cells actually intersect the stage). Columns $LP_{1+\epsilon}$ ($\epsilon = 0.01$), LP_V^*, and C_0^* contain the results for the $(1 + \epsilon)$-, 4-, and $O(1)$-approximation algorithms, *including* the power-up of all lightsources to satisfy all points on the stage as proved in the previous sections.

Not surprisingly, the differences in the approximation ratios of the latter two algorithms are dominated by the 'power-up' factor. Note though, that this completely disregards to which degree some point on the stage gets insufficient light. This will be taken into account in the next measurement.

D, $\lvert G_V\rvert$	$LP_{1+\epsilon}$	LP'_V	C'_0	APX$_V$	APX$_{\text{COMB}}$
70, 6	414.43	561.13	1071.24	1.43	2.73
70, 15	112.40	139.37	167.11	1.31	1.57
70, 22	87.22	102.10	162.52	1.24	1.97
140, 8	732.34	814.775	1112.26	1.11	1.51
140, 19	333.47	527.79	615.39	1.58	1.95
140, 41	182.75	239.96	299.79	1.39	1.74

Table 2: Energy costs for different stage lengths, number of guards, and algorithms, but with adaptive power-up.

D, $\lvert G_V\rvert$	$\alpha = 0$	$\alpha = 1$	$\alpha = 3$	$\alpha = 5$
70, 14	3.39	3.31	3.19	3.16
70, 20	2.75	2.76	1.54	1.51
70, 31	5.09	5.09	4.98	3.92
140, 21	6.81	7.03	1.75	1.74
140, 30	10.15	10.06	9.98	9.88
140, 65	5.45	5.40	2.64	2.54

Table 3: Maximum excess of light for different stage lengths, number of guards, and values of α.

Actual Performance

To assess the real amount of 'power-up' required to satisfy all points on the stage, we evaluated the power assignments on a ϵ-good set of guards with $\epsilon = 0.01$ and determined the 'worst' guard, i.e. the guard that received the least amount of light. We use this to compute an appropriate power-up factor (for all lightsources; one could improve here by powering up lightsources locally only). The results can be found in Table 2; here LP'_V and C'_0 denote the power assignments resulting from this more careful power-up strategy for the 4- and the $O(1)$-approximation scheme.

We have also included in the last two columns the resulting approximation factor of the solutions (taking the $(1+\epsilon)$ solution with $\epsilon = 0.01$ as a lower bound). It turns out that in fact, by using this more refined power-up strategy, the combinatorial $O(1)$ as well as the 4-approximation get much closer to the optimum solution than guaranteed by the pessimistic theoretical analysis.

Further Observations

In the proof of the approximation ratio of the $O(1)$ algorithm, we employed a pruning procedure with some parameter α to actually be able to bound the amount of excess light at any guard. But already when proving this, we suspected that this was only necessary because of our inability to come up with a better analyis. We want to substantiate this suspicion by an experiment. In Table 3, the worst amount of excess light at any guard is stated for different values of α. It seems that even without pruning (i.e. $\alpha = 0$) the amount of excessive light is bounded by a small constant; in our proof we were only able to bound it for $\alpha \geq 3$.

D, $\lvert G_V\rvert$	$LP_{1+\epsilon}$	C_0^*	$C_{\text{SIMPLE}}^{1+\epsilon}$	$\text{APX}_{\text{SIMPLE}}^{1+\epsilon}$
70, 6	565.28	54175.6	613.63	1.08
70, 16	130.94	14386.3	180.66	1.38
70, 32	55.77	6106.18	72.90	1.31
140, 15	581.27	63846.1	684.78	1.18
140, 38	235.43	37361.8	322.78	1.37
140, 54	117.95	13528	154.83	1.31

Table 4: Energy costs of the LP-based $(1+\epsilon)$-approximation, the combinatorial algorithm as described, and the combinatorial algorithm using a ϵ-good set of guards.

In a last experiment we have run the simple combinatorial algorithm on a ϵ-good set of sample points. Even though we cannot prove any better approximation ratio than for the original algorithm, the results look quite promising as can be seen in Table 4. Here we denote by $C_{\text{SIMPLE}}^{1+\epsilon}$ the outcome of running the $O(1)$ algorithm on an ϵ-good set of guards, including the required power-up.

6. CONCLUSIONS

In this paper we have introduced a light attenuation model under which the large class of illumination problems can be considered. Our model also takes into account the decrease of light intensity with distance – something that had not been regarded in classical models for illumination problems. As a concrete example we have examined the problem of illuminating a stage using a fixed set of lightsources minimizing the overall energy.

Looking at our solutions, we believe this new model creates quite a number of new and interesting open questions in the context of illumination problems for which only a combination of geometric reasoning and techniques from combinatorial optimization will lead to provable results.

Acknoledgements

We want to thank Arno Eigenwillig for helpful comments and suggestions.

7. REFERENCES

[1] L. Aleksandrov, A. Maheshwari, and J.-R. Sack. Approximation algorithms for geometric shortest path problems. In *Proc. 32nd Annu. ACM Sympos. Theory Comput.*, pages 286–295, 2000.

[2] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical models and image processing: GMIP*, 60(2):125–135, 1998.

[3] G. Csizmadia and G. Tóth. Note on an art gallery problem. *Comput. Geom. Theory Appl.*, 10(1):47–55, 1998.

[4] J. Czyzowicz, E. Rivera-Campo, and J. Urrutia. Optimal floodlight illumination of stages. In *Proc. 5th Canad. Conf. Comput. Geom.*, pages 393–398, 1993.

[5] L. Fejes-Tóth. Illumination of convex disks. In *Acta Math. Acad. Sci. Hungar. 29*, pages 355–360, 1977.

[6] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. 1988.

[7] J. Kahn, M. M. Klawe, and D. Kleitman. Traditional galleries require fewer watchmen. *SIAM J. Algebraic Discrete Methods*, 4:194–206, 1983.

[8] D. Lee and A. Lin. Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32(2):276–282, 1986.

[9] C. H. Papadimitriou. An algorithm for shortest-path motion in three dimensions. *Inform. Process. Lett.*, 20:259–263, 1985.

[10] B. Speckmann and C. D. Tóth. Allocating vertex π-guards in simple polygons via pseudo-triangulations. In *Proc. 14th ann. ACM-SIAM symp. on Discr. Alg.*, pages 109–118. SIAM, 2003.

[11] C. D. Toth. Art gallery problems with guards whose range of vision is 180^o. *Comput. Geom. Theory Appl.*, 17:121–134, 2000.

[12] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North-Holland, 2000.

[13] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.