Procedural Design of Exterior Lighting for Buildings with Complex Constraints

Supplemental Material

MICHAEL SCHWARZ and PETER WONKA

1. On projecting arbitrary configurations into the space of valid configurations

To generate a random valid configuration $\Lambda \in \mathcal{V}$, it may be tempting to start with an arbitrary configuration $\Lambda^0 \in \mathcal{D}$ (for instance by determining a random value for each variable of Λ^0) and then project it into the space of valid configurations $\mathcal{V}$. However, due to the complex shape of $\mathcal{V}$ with the intricate dependencies among its variables, such a projection is generally hard and there is also no obvious notion of what the closest valid configuration Λ^* is, let alone finding such a Λ^*. Pivotally, it is not feasible to simply project each variable separately because these are interrelated, with the value chosen for one affecting the permissible domains for others.

Consequently, settling on a certain value for one variable may require changes to other variables, which then may necessitate further changes. And unless these values are selected carefully and fixed in the correct order, which actually would require some accounting for all constraints like in our approach, later changes may require changing previously settled values again, potentially triggering another chain of changes. Such a scheme may hence require many iterations to arrive at a valid configuration.

Finding not just any but the closest valid configuration Λ^* is even harder. In particular, when adjusting one variable, it is often not possible to assess which value would eventually yield Λ^* without considering the whole chain of changes entailed by this specific value. Another issue is that while a distance between two configurations can be defined, it is unclear how to meaningfully weight the differences in the individual variables relative to each other. For instance, would it be better to select luminaire A and strongly decrease the flux or to select luminaire B, which requires only a moderate increase in flux but also a reduction in the number of luminaires?

2. Implementation-specific choices

In the following, we provide some additional implementation details for our optimization approach. The underlying choices are of empirical nature.

Interleaved exploration of chains. For each chain, we keep track of how many iterations n_{acc} have passed since a proposal Λ' has been accepted, how many iterations n_{rdc} have passed since a proposal Λ' led to a reduction in cost (i.e., $C(\Lambda') < C(\Lambda)$), and how many iterations n_{bst} have passed since a proposal Λ' improved on the best solution Λ^* encountered so far. If either $n_{\text{acc}} > 30$, $n_{\text{rdc}} > 50$, or $n_{\text{bst}} > 100$ (and $n_{\text{rdc}} > 20$), we consider switching to another chain with a probability of 0.8 every 100 iterations.

Annealing schedule. For each chain, the temperature is chosen as $T = C(\Lambda^{(0)})/15 \cdot T'$, where $\Lambda^{(0)}$ denotes the chain's initial configuration and $T' = 100/\max\{n_{\text{acc}}, 100\}$.

Elementary mutations. When determining a relative change $\Delta\tau$, the standard deviation of the employed truncated normal distribution is chosen as $1000 \cdot T'$ for fluxes (e.g., Φ), $0.5\ell \cdot T'$ for translations (e.g., u_k, Δv_k), with ℓ denoting the length along which translation may happen (e.g., $\|L_p\|$), $15° \cdot T'$ for rotations (e.g., ϕ), and $1.5 \cdot T'$ for counts (e.g., n).

Multiple mutations per iteration. The number of mutations m executed per iteration varies between one and five; these are chosen with probabilities $0.75 - 0.25w$, $0.14 + 0.08w$, $0.07 + 0.07w$, $0.03 + 0.06w$, $0.01 + 0.04w$, respectively, where $w = \max\{0, (T' - 0.2)/0.8\}$.

Change proposals. A random mutation is applied with a probability of 0.3. Otherwise, an unsatisfied goal or a subeffective installation site is targeted (randomly chosen according to their contribution to the overall cost C).