

Spatial Regionalization: A Hybrid Quantum Computing Approach

Yunhan Chang^{1,2}

¹Department of Computer Science and Engineering

²Center for Geospatial Sciences
University of California, Riverside
ychan268@ucr.edu

Federico M. Spedalieri³

³Information Sciences Institute
University of Southern California
fspedali@isi.edu

Amr Magdy^{1,2}

¹Department of Computer Science and Engineering

²Center for Geospatial Sciences
University of California, Riverside
amr@cs.ucr.edu

Ibrahim Sabek^{4,5}

⁴Department of Computer Science

⁵Spatial Sciences Institute
University of Southern California
sabek@usc.edu

Abstract

Quantum computing has shown significant potential to address complex optimization problems; however, its application remains confined to specific problems at limited scales. Spatial regionalization remains largely unexplored in quantum computing due to its complexity and large number of variables. In this paper, we introduce the first hybrid quantum-classical method to spatial regionalization by decomposing the problem into manageable sub-problems, leveraging the strengths of both classical and quantum computation. This study establishes a foundational framework for effectively integrating quantum computing methods into realistic and complex spatial optimization tasks. Our initial results show a promising quantum performance advantage for a broad range of spatial regionalization problems and their variants.

ACM Reference Format:

Yunhan Chang^{1,2}, Amr Magdy^{1,2}, Federico M. Spedalieri³, and Ibrahim Sabek^{4,5}. 2025. Spatial Regionalization: A Hybrid Quantum Computing Approach. In *Proceedings of Minneapolis, USA (SIGSPATIAL '25)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Quantum annealing [11] is an optimization technique inspired by quantum mechanics, designed to solve complex optimization problems by exploring a vast solution space through quantum fluctuations. Unlike its classical counterpart, simulated annealing [14] that is based on thermal fluctuations, quantum annealing uses quantum tunneling to escape local minima and potentially converge on the global optimum. In recent years, the quantum annealing hardware has been significantly advanced to support thousands of qubits, e.g., *D-Wave Quantum Inc.* annealer currently offers approximately **5000 qubits** [13]. However, current quantum systems remain constrained by notable technical limitations that limit solving problems

of large sizes. To address the scale challenge, quantum annealers are supplemented by hybrid solvers that integrate classical CPUs with Quantum Processing Units (QPUs) at different phases. The classical computer handles problem decomposition, setting initial conditions, and refining results using well-established optimization methods, while the QPU leverages quantum annealing with its native QUBO-like formulations [10]. These solvers, such as DQM, CQM, and NL solvers [7], can partition large-scale problems into smaller sub-problems that are then solved with the QPU. Such advances have enabled quantum computers to handle complex optimization problems and enterprise-level applications [6].

Spatial regionalization is an essential and expensive spatial operation that serves many real-world applications in economics, urban planning, and other domains requiring efficient spatial resource allocation [1, 12, 16]. Regionalization aggregates spatial areas into contiguous regions, aiming to minimize heterogeneity within each region. Constrained by explicit spatial contiguity requirements, spatial regionalization is proved to be NP-hard [12], and finding optimal solutions remains computationally infeasible due to combinatorial complexity [2, 12, 16].

This paper introduces the first hybrid quantum-classical approach designed for seeding-based spatial regionalization. We decompose regionalization into smaller, quantum-processable sub-problems and explore the solution space more effectively than the traditional heuristic methods. Specifically, we decompose seeding-based regionalization into two phases, spanning five stages, adapting a variation of our framework proposed in [3] for classical computers. Then, the quantum annealing is leveraged to improve the performance of two stages: **seed selection** and **local optimization**, which are used by all seeding-based regionalization techniques. Therefore, our proposed improvements are general and applicable in various applications. Our initial empirical results on the D-Wave quantum annealer show promising improvements in seed selection quality and local optimization runtime. Ongoing and future explorations might exploit our quantum model in [5] for enforcing spatial contiguity constraints, which are imposed by all regionalization techniques, to build a quantum-optimized *region growing* stage. The rest of this paper introduces a brief background, discusses the related work, and presents our proposed approach and initial empirical results.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGSPATIAL '25, 978-1-4503-XXXX-X/2025/06

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/XXXXXXX.XXXXXXX>

2 Background

Quantum annealing (QA) has been proposed as a novel heuristic to solve many of the well-known NP-hard problems [11], and it is similar to Simulated Annealing (SA). In SA, a solution space is explored by randomly proposing changes to the current solution. If the new value of the objective function is decreased, the proposed change is accepted. If the objective function increases, the change is accepted with some probability according to the temperature. In QA, these changes are applied on a quantum superposition of the different solutions, allowing for a massively parallel search of the solution space. This quantum parallelism, combined with the quantum tunneling [5], facilitates fast exploration and potentially avoids being stuck in local minima.

Quadratic Unconstrained Binary Optimization (QUBO). In QA, the problem is represented as a **QUBO formulation** that consists of a quadratic objective function over binary variables, while any constraints are represented as penalty terms [10]. QUBO is an optimization technique that is inherently similar to other operations research techniques and requires no or limited background knowledge of the internals of the quantum annealing process from a user perspective. Under the hood, the quantum annealer encodes the QUBO formulation in the interactions between *qubits*, which are the quantum counterpart of the classical binary bits. More details about the quantum annealing process can be found in [5].

Hybrid quantum-classical computing. Although many optimization problems can be formulated as QUBOs, the limitations in the current quantum hardware, e.g., small size and low degree of qubit connectivity, restrict the size and complexity of problems that can be practically solved. This challenge is further amplified in problems with many constraints, where encoding constraints as penalty terms [10] often requires additional slack variables, consuming even more qubits. A widely adopted and practical approach is integrating quantum annealers within a hybrid framework, combining both classical CPU and quantum computing powers to solve large problems efficiently using **hybrid solvers**. The core idea of a *hybrid solver* is to decompose large or dense problems into smaller subproblems that fit within the connectivity and size constraints of the quantum hardware. In this framework, a quantum processor is paired with a classical heuristic in an iterative loop: the quantum processor solves selected subproblems, while the classical component manages orchestration, constraint handling, and post-processing. By offloading constraint enforcement and much of the solution space exploration to the classical side, hybrid solvers effectively mitigate the challenges posed by sparse qubit connectivity. *D-Wave Quantum Inc.*, a leading provider of quantum annealing hardware, currently offers four hybrid solvers tailored to different problem types and variable domains: the Binary Quadratic Model (BQM), the Discrete Quadratic Model (DQM), the Constrained Quadratic Model (CQM), and the Nonlinear-Program Hybrid Solver (NL-solver) [7]. Notably, both CQM and NL-solver support the explicit expression of constraints in a declarative form.

3 Related Work

Classical spatial regionalization clusters a collection of spatial areas, i.e., polygons, into regions that are spatially contiguous and homogeneous with respect to one or more attributes [1–3, 9, 12, 16]. Numerous variants exist, for example, the *p*-regions problem takes

the number of regions *p* as user input, the max-*p*-regions problem automatically finds the maximum number of regions, whereas the *p*-compact-regions problem emphasizes spatial compactness of regions. All regionalization variants are NP-hard, so practical solutions rely on heuristics to find approximate *good* solutions for large data. The two major classes of classical solutions are: (a) **seeding-based regionalization** and (b) **graph-based regionalization**. Seeding-based regionalization techniques select seed areas and build regions, i.e., clusters, around these seeds [3, 9]. Graph-based regionalization builds connected regions by working directly on spatial adjacency graphs [9]. It includes both top-down, e.g., SKATER, and bottom-up agglomerative techniques, e.g., REDCAP [9]. Over the past decade, it has been empirically shown that seeding-based techniques have better scalability than graph-based techniques, so they empower spatial data analysts to support large datasets [1, 2, 12, 16]. However, classical heuristics are computationally extensive and prone to getting trapped in local optima, making the search for high-quality solutions expensive. By using quantum tunneling, quantum annealing can achieve high-quality solutions faster. None of the existing techniques leverages quantum annealing to address this NP-hard problem.

Quantum-based spatial clustering has emerged as a promising tool for NP-hard clustering and partitioning problems. Mapping spatial clustering to QUBO form is non-trivial, yet several studies show feasibility for related graph-partitioning tasks. Ushijima-Mwesigwa [18] encodes modularity-maximizing community detection as a QUBO, with each node's binary spin indicating community membership. Negre [17] extends this to *k* communities on a D-Wave 2000Q, obtaining modularity scores comparable to leading classical algorithms. Arthur and Date [4] formulate balanced *k*-means clustering as a QUBO, achieving clustering quality on par with classical *k*-means and highlighting potential scalability benefits. These proof-of-concept results illustrate that quantum annealing can effectively handle clustering objectives, optimizing over assignments of points to clusters in a manner that, in some cases, escapes local minima that trap classical heuristics.

Quantum clustering studies omitted explicit enforcement for spatial contiguity, as it requires a prohibitive number of qubits and scales poorly with the problem size. They relied on adjacency graph structures that implicitly encoded connectivity. Very recent work addresses spatial contiguity directly [5] based on the discrete quadratic model (DQM) [7]. However, none of the existing quantum-based clustering approaches have addressed seeding spatial regionalization, which is shown over the past decade to be the most scalable class of spatial regionalization algorithms. **Our work is the first to explore seeding spatial regionalization on quantum computers.**

4 Quantum Spatial Regionalization

Spatial regionalization is an NP-hard clustering problem, so all existing algorithms use heuristic-based search to find *good approximate solutions* rather than optimal solutions. Most existing seeding-based techniques work in two phases: (1) finding an initial solution, and (2) optimizing the initial solution to output a final solution. Building on our previous work in [3], the first phase has four stages: *seed selection*, *region growing*, *enclave assignment*, and *constraint adjustments*. The second phase is one stage: *local optimization* that

employs heuristic search techniques to improve the approximate solution quality, measured in regional heterogeneity.

We propose a hybrid approach that uses quantum annealing to improve the **seed selection** and the **local optimization** stages, while other stages are handled using efficient classical algorithms as in [16]. These two stages have the highest potential to utilize the quantum annealing advantages to improve the runtime or solution quality. The seed selection is purely quantum-encoded. The local optimization is an iterative hybrid loop: the CPU prepares a set of potential region border areas, which the QPU then shuffles to refine the solution. This ongoing work shows a preliminary, promising exploration that will be further strengthened in future explorations. The rest of this section describes each quantum-encoded stage.

4.1 Quantum Seed Selection

The seed selection stage takes n spatial areas, i.e., polygons, as input and finds p spatially scattered seed areas to build regions around. Scatteredness ensures seeds remain well-separated so that subsequent regional growth is not squeezed together. We measure scatteredness as the minimum pairwise distance between geometric centroids of seed areas, as in [16], with the objective of maximizing this scatteredness. Therefore, seed selection maps to a well-known problem called the max-min dispersion [19]. To address the max-min dispersion, we adopt a quantum-friendly formulation based on the Maximum Independent Set (MIS) [8]. Given a set of 2D points V , we first define a complete graph $G = (V, E)$ where each edge e_{ij} is assigned a weight equal to the Euclidean distance d_{ij} between points i and j . The MIS is a set of vertices where no two vertices in the set are connected by an edge, and the set has the largest possible size among all such sets in the graph. For a candidate minimum (dispersion) distance threshold d_m , we construct a subgraph $G' = (V, E')$ by removing all edges with $d_{ij} > d_m$. Solving the MIS on G' yields a subset of points where all pairwise distances are at least d_m . If the MIS size is less than the required number of seeds p , d_m is reduced; otherwise, it can be increased. A binary search over feasible values of d_m identifies the largest achievable dispersion. We associate a binary variable x_i with each spatial area, where $x_i = 1$ if area i is selected as a seed and 0 otherwise. The primary objective is to maximize the number of seeds, which is equivalent to minimizing the term $-\sum_i x_i$. To enforce the MIS constraint, we introduce a penalty λ_{MIS} for any pair of connected seeds being selected simultaneously. This results in the following objective function:

$$X = -\sum_{i \in V} x_i + \lambda_{\text{MIS}} \sum_{(i,j) \in E'} x_i x_j$$

Minimizing X thus yields the largest possible set of seeds that are all separated by at least the distance threshold d_m . This technique is encoded as a BQM model [7]. For space limitations, our complete BQM formulation is described in Appendix A.

4.2 Quantum Local Optimization

The local optimization stage takes an initial solution of p spatially-contiguous regions as input, and outputs an optimized solution. The optimized solution still consists of p spatially-contiguous regions but with lower *regional heterogeneity*. This stage is by far the most time-consuming stage for all seeding-based regionalization techniques, and it also determines the final objective function value,

i.e., heterogeneity. This makes it the best candidate to leverage the quantum advantage for faster runtime and better objective value.

We employ quantum annealing to model a “swap-or-move” local search technique that improves the initial solution by shuffling border movable areas between neighboring regions. The goal is minimizing an objective function representing the overall **regional heterogeneity** while maintaining the **spatial contiguity** of each region. By finding articulation areas using the classic method [16], we can find all movable areas in linear time. The heterogeneity of a region R is defined as $H_R = \sum_{a_i, a_j \in R, i < j} |A(a_i) - A(a_j)|$, for the dissimilarity attribute A of all pairs of areas $a_i, a_j \in R$. The overall heterogeneity H is the summation of heterogeneity of all the p regions, i.e., $H = \sum_{k=1}^p H_{R_k}$, $\forall 1 \leq k \leq p$. If an area a_i moves from a donor region R_d to a receiver region R_r , both H_{R_d} and H_{R_r} change, and hence H changes.

We employ a Constraint Quadratic Model (CQM) [7] to optimize choosing area moves that improve the objective function H . It is computationally expensive to aggregate the attribute A over all area pairs directly on a QPU due to an architectural limitation of the qubits. To overcome this, we adopt an estimation strategy that avoids full recalculation of H by directly estimating the *change* in heterogeneity (Δ_H) resulting from a move. The objective function is modeled using key statistical attributes of a region R , its total number of areas (N_R), mean attribute A value (M_R), and attribute A variance (V_R), along with the value of attribute A_m of a movable border area $a_m \in R$. These statistics are chosen for their computational efficiency, as they can be updated in linear or constant time on a classic computer. The potential destinations for movable areas serve as the model’s variables. By assuming the attributes of large datasets are normally distributed, when a region adds or removes an area, this change is estimated using its statistical properties:

$$\Delta \hat{H}_{\text{add}} = N_R \cdot \sqrt{V_R + (M_R - A_m)^2}$$


$$\Delta \hat{H}_{\text{remove}} = -(N_R - 1) \cdot \sqrt{V_R + (M_R - A_m)^2}$$

$\Delta_H = \Delta \hat{H}_{\text{add}} + \Delta \hat{H}_{\text{remove}}$. By estimating Δ_H for every potential move, we can leverage quantum optimization to simultaneously evaluate multiple relocations and select the combination that most significantly reduces total system heterogeneity. To preserve spatial contiguity, we impose a critical constraint: each region may either donate or receive at most one area per iteration, but not both. This constraint is necessary because moving an area can alter the set of remaining movable areas and change the receivable areas of the donor region.


The hybrid quantum local optimization stage executes multiple iterations. Each iteration computes the list of movable areas using the classical CPU. Then, the list is dispatched to the QPU to explore multiple area moves using a CQM model to find a better solution. Based on the separator theorem [15], the number of potential movable border areas ranges from $2\sqrt{np}$ to $2p\sqrt{n}$ in *every single iteration*. For space limitations, our complete CQM formulation is described in Appendix B.

5 Empirical Results

The proposed BQM and CQM models for *seed selection* and *location optimization*, respectively, are implemented using Python 3.12 and constructed using the D-Wave Ocean toolbox ‘dimod’ library



(a) Quantum Improvements (%) (b) Absolute Runtime
Figure 1: Performance in local optimization vs. dataset size



(a) Quantum Improvements (%) (b) Absolute Quality
Figure 2: Performance in seed selection vs. p

(version 0.12.14). We run our algorithms on an on-premise D-Wave Advantage annealing quantum computer (housed at the USC Quantum Computing Center) accessed via D-Wave’s Leap cloud APIs. The current QPU has Pegasus P16 topology and provides ~5000 superconducting flux-qubits, each is 15-way connected. We compare our quantum algorithms versus the classical counterparts in PRUC [16], setting its user threshold to zero and using the same evaluation datasets. The default value of dataset size is 50 and p is 10. All results are collected after 10 iterations of local optimization. The relative improvement of runtime is calculated as $(\frac{T_{classic} - T_{quantum}}{T_{classic}} \times 100)\%$, and similarly for quality improvements.

Figure 1 shows **initial results** for the local optimization performance. Figure 1a depicts the relative improvements, as percentages, of our quantum local optimization versus the classical counterpart in both runtime and quality, measured as the objective function value (overall regional heterogeneity H). With increasing dataset size (DS), the runtime improvement increases while the quality improvement decreases. The quality improvement starts at 22.4% at DS=50 and continuously decreases to 6.3% at DS=5k. This suggests that quantum annealing keeps marginal improvements in solution quality with increasing data sizes, yet it will not degrade the quality. On the contrary, the runtime improvement continuously increases from -191% (a slowdown) at DS=50 to 62.3% (a speedup) at DS=5k. These results are confirmed in Figure 1b, which shows relatively stable quantum runtime as data size increases. This shows a promising quantum advantage in terms of runtime speedup with larger datasets. It worth noting that a dataset of all US counties has 3.2k polygons, so 5k datasets are used in many practical applications.

Figure 2 shows **initial results** for the seed selection performance. Figure 2a depicts the relative improvements, as percentages, showing better runtime and quality improvements with increasing p value. However, the quantum runtime ranges from 300- 500 milliseconds compared to 3-7 milliseconds for the classical algorithm for this small dataset size, so the increasing improvement is not significant and must be tested on larger data sizes that our

preliminary model does not currently support. On the contrary, the quality improvement, measured as the min seed distance, is significantly improved to reach 46.5%. This is confirmed by the absolute quality values in Figure 2b as larger p makes it harder for classical heuristics to optimize for many seeds, while the quantum annealer can still explore a larger search space. This suggests a potential quantum advantage in improving the seed selection quality.

6 Conclusion and Future Explorations

This paper has presented the first hybrid quantum approach to address the spatial regionalization problem. The approach decomposes the regionalization process into stages, and exploits the quantum computing power for the two most promising stages, seed selection and local optimization. Two algorithms have been proposed and realized using D-Wave’s BQM and CQM hybrid solvers. The presented initial results suggest a potential quantum advantage to improve the objective function quality of seed selection and the runtime efficiency of the local optimization. More explorations are ongoing to improve the models in terms of scalability, more effective heuristics, and support richer variants of spatial regionalization problems. Future extensions can use our quantum model to enforce spatial contiguity [5] towards a quantum *region growing* stage.

References

- [1] H. Alrashid, Y. Liu, and A. Magdy. SMP: Scalable Max-P Regionalization. In *SIGSPATIAL*, 2022.
- [2] H. Alrashid, Y. Liu, and A. Magdy. PAGE: parallel scalable regionalization framework. *ACM Transactions on Spatial Algorithms and Systems*, 9(3):1–26, 2023.
- [3] H. Alrashid and A. Magdy. A Scalable Unified System for Seeding Regionalization Queries. In *Proceedings of the 18th International Symposium on Spatial and Temporal Data, SSTD*, pages 96–105, 2023.
- [4] D. Arthur and P. Date. Balanced k -means clustering on an adiabatic quantum computer. *Quantum Information Processing*, 20(9):294, 2021.
- [5] Y. Chang, A. Magdy, and F. M. Spedalieri. Quantum Modeling of Spatial Contiguity Constraints, 2025. <https://arxiv.org/abs/2505.12608>. To appear at the ACM SIGMOD Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data 2025).
- [6] D-Wave. Real-World Quantum Applications at Business Scale. <https://www.dwavequantum.com/learn/customer-success-stories/>, 2025.
- [7] D-Wave Systems Inc. *Leap Service’s Hybrid Solvers*. D-Wave Systems Inc., 2025. Accessed 2025-06-12.
- [8] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, K. Pichler, H. Bernien, et al. Quantum optimization of maximum independent set using rydberg atom arrays. *Nature*, 617(7958):1209–1215, 2022.
- [9] GeoDa Center for Geospatial Analysis and Computation. Spatial Clustering – pygeoda 1.0.0 documentation. https://geodacenter.github.io/pygeoda/spatial_clustering.html, 2025. Accessed: 2025-06-12.
- [10] F. Glover, G. Kochenberger, R. Hennig, and Y. Du. Quantum bridge analytics i: a tutorial on formulating and using qubo models. *Annals of Operations Research*, 314(1):141–183, 2022.
- [11] T. Kadowaki and H. Nishimori. Quantum annealing in the transverse ising model. *Phys. Rev. E*, 58:5355–5363, Nov 1998.
- [12] Y. Kang and A. Magdy. EMP: Max-P Regionalization with Enriched Constraints. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, 2022.
- [13] A. D. King et al. Quantum critical dynamics in a 5,000-qubit programmable spin glass. *Nature*, 617(7959):61–66, 2023.
- [14] Y. Koshka and M. A. Novotny. Comparison of D-Wave Quantum Annealing and Classical Simulated Annealing for Local Minima Determination. In *Proceedings of the 2020 IEEE International Conference on Rebooting Computing (ICRC)*, 2020.
- [15] R. J. Lipton and R. E. Tarjan. A Separator Theorem for Planar Graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [16] Y. Liu, A. R. Mahmood, A. Magdy, and S. Rey. PRUC: P-regions with User-Defined Constraint. *Proc. VLDB Endow.*, 15(3), 2021.
- [17] C. Negre, H. Ushijima-Mwesigwa, and S. Mniszewski. Detecting multiple communities using quantum annealing on the d-wave system. *PLOS ONE*, 15, 2020.
- [18] H. Ushijima-Mwesigwa, C. F. A. Negre, and S. M. Mniszewski. Graph partitioning using quantum annealing on the d-wave system. In *Proceedings of the Second International Workshop on Post Moores Era Supercomputing, PMES’17*, 2017.
- [19] K. Yukiyoishi, T. Mikuriya, H. S. Rou, G. T. F. de Abreu, and N. Ishikawa. Quantum speedup of the dispersion and codebook design problems. *IEEE Transactions on Quantum Engineering*, 2024.

Appendix

A BQM in Seed Selection

The seed selection stage is modeled as a max-min dispersion problem, which we solve by repeatedly finding the Maximum Independent Set (MIS) in a graph. We solve this by employing a quantum-assisted binary search to find the largest possible minimum distance (d_m) that allows for a valid selection of p seeds.

Variables. Let $V = \{a_1, a_2, \dots, a_n\}$ be the set of spatial areas. We define a binary variable for each area:

$$x_i = \begin{cases} 1 & \text{if area } a_i \text{ is selected as a seed} \\ 0 & \text{otherwise} \end{cases}$$

Objective Function. For a given minimum distance threshold d_m , we construct a graph $G' = (V, E')$ where edge $(i, j) \in E'$ if $d'_{ij} \leq d_m$. The MIS objective maximizes the number of selected areas:

$$\text{Maximize } \sum_{i=1}^n x_i$$

By keep finding MIS qualified for a minimum distance, we can find a more precise range of d_m by iteration.

Constraints: To ensure we select exact p seeds from dataset and try to make them scattered, we enforce the following constraints:

MIS Constraint. No two adjacent areas in G' can both be selected as seeds. For each edge $(i, j) \in E'$:

$$x_i \cdot x_j = 0$$

This is enforced by adding a quadratic penalty term $\lambda_{\text{MIS}} \cdot x_i \cdot x_j$ to the objective, where λ_{MIS} is a sufficiently large penalty coefficient. The overall objective function with constraints penalties is:

$$X = - \sum_{i \in V} x_i + \lambda_{\text{MIS}} \sum_{(i,j) \in E'} x_i x_j$$

The requirement of selecting exactly p seeds is handled by a classical binary search that repeatedly calls this BQM. After the BQM returns a maximum independent set of size k , if $k \geq p$, the distance d_m is feasible; we increase it to seek better dispersion. If $k < p$, the distance d_m is too high; we must decrease it. This process converges on the maximum d_m that supports a solution of at least p seeds. If it converges to more than p , the final output p seeds are selected at random.

B CQM in Location Optimization

To run the local optimization phase on a D-Wave system, we designed a Constrained Quadratic Model (CQM). This model is defined by three key components: binary variables representing all areas' potential moves calculated by CPU [16], an objective function to estimate the change in heterogeneity, and constraints to preserve spatial contiguity.

Variables. Let S be the set of all possible moves, where a move is a triplet (i, d, r) representing the transfer of area a_i from a donor region R_d to a receiver region R_r . We define a binary variable for

each possible move:

$$x_{i,d,r} = \begin{cases} 1 & \text{if area } a_i \text{ is moved from } R_d \text{ to } R_r \\ 0 & \text{otherwise} \end{cases}$$

Based on the separator theorem [15], the number of potential movable border areas ranges from $2\sqrt{np}$ to $2p\sqrt{n}$ in *every single iteration*. This is a major number for large values of n , which is typically associated with larger p values as well.

Objective Function. The objective is to minimize the total estimated change in heterogeneity across all selected moves. This is formulated as:

$$\text{Minimize } \sum_{(i,d,r) \in S} (\Delta \hat{H}_{\text{add}} + \Delta \hat{H}_{\text{remove}}) \cdot x_{i,d,r}$$

The terms $\Delta \hat{H}_{\text{add}}$ and $\Delta \hat{H}_{\text{remove}}$ are calculated for each potential move based on the statistical properties of the involved regions, as defined in Section 4.2.

Constraints. To ensure the solution is connectivity-stable during areas exchanging, we impose the following constraints:

1. *One Move per Area.* Each movable area a_i can be transferred at most once. This is enforced for every area a_i in the set of movable areas. Given that for a specific area a_i , its donor region R_d is fixed, we sum over all possible receiver regions R_r :

$$\forall a_i \in \text{MovableAreas}, \sum_{r | (i,d,r) \in S} x_{i,d,r} \leq 1$$

2. *Region Stability.* Each region R_k can either act as a donor for one move or as a receiver for one move, but not both. This is enforced for every region k from 1 to p :

$$\forall k \in \{1, \dots, p\}, \sum_{(i,k,r) \in S} x_{i,k,r} + \sum_{(i,d,k) \in S} x_{i,d,k} \leq 1$$

This constraint ensures that all selected moves are mutually compatible within a single iteration, preventing cascading changes that could invalidate the premises of other concurrent moves.