Information-driven Gas Distribution Mapping

Emile Meyer

Professor : Alcherio Martinoli
Assistant(s) : Nicolaj Bösel-Schmid

Mapping the spatial distribution of gases (GDM) using autonomous drones is a critical capability for early-warning industrial leak detection and environmental monitoring. While recent physics-based algorithms, such as ADApprox, offer superior accuracy by locally approximating the advection-diffusion equation, they suffer from a major computational bottleneck. Their reliance on dense, static uniform grids creates a dual inefficiency: wasting immense computational power evaluating empty space while under-sampling sharp, complex plume boundaries. This results in an $\mathcal{O}(NM^2C)$ complexity that severely limits real-time robotic deployment.

To overcome this limitation, this project structurally optimizes the computational grid by explicitly integrating Adaptive Mesh Refinement (AMR). The adaptation is actively driven by the spatial gradient of the preliminary gas concentration, allowing the system to naturally allocate resources exactly where they are most valuable: the highly informative outer edges of the plume. Two distinct dynamic architectures were developed and evaluated. The first is an r-refinement approach utilizing Artificial Potential Fields (APF) to continuously relocate a fixed number of computational nodes via simulated attractive and repulsive forces. The second is an h-refinement strategy based on Quadtree recursive spatial decomposition to dynamically alter the mesh topology.

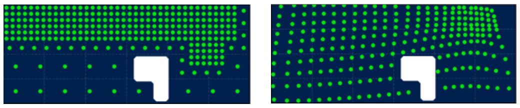

Furthermore, an additional methodological contribution of this research lies in breaking the standard assumption of isotropy inherent in traditional AMR methods. Because gas dispersion is advection-dominated, plumes stretch smoothly downwind but form abrupt cross-wind boundaries. To prevent the algorithms from wasting computational nodes along the advective axis, explicit wind-driven anisotropic constraints were implemented. The APF method integrates a stretch factor to elliptically deform its repulsive field along the wind vector, while the Quadtree incorporates an angle limit to perform highly efficient binary splits perpendicular to the advective flow.

These architectures were rigorously benchmarked on the Wanting dataset, a complex simulated environment featuring numerous obstacles. The wind-driven Quadtree matches the fastest baseline grid in both spatial accuracy and geometric fidelity while significantly slashing execution time. Furthermore, incorporating wind information cuts the standard isotropic Quadtree's computation time by more than half without degrading its shape preservation. Conversely, adding wind direction to the APF degrades all metrics on this cluttered dataset. This rigid stretching along a single macro wind vector over-fits to global flow, failing to capture the localized turbulence and unpredictable gas scattering caused by obstacles.

Method	Time [ms] ↓	MAE ↓	Shape [%] ↑
ADApprox (0.2m)	**6807**	**309.7**	93.6
ADApprox (0.4m)	**984**	**350.6**	**93.3**
Quadtree	1256	328.6	**93.1**
Quadtree w/Wind	**572**	**350.5**	**93.1**
APF	**1314**	**323.3**	**92.9**
APF w/Wind	1379	336.3	92.7

Ultimately, while the APF method excels in obstacle-free simulations, its rigidly stretched potential field struggles to capture localized turbulence in cluttered physical environments. Consequently, the Quadtree architecture emerges as the definitive and most robust solution for real-world deployment. It successfully reconciles high-fidelity mapping with strict latency constraints, almost matching the absolute spatial precision of a high-resolution 0.1m grid while operating over five times faster, thereby fully supporting autonomous robotic exploration.

Method	Time [ms] ↓	MAE ↓	Shape [%] ↑
ADApprox (0.1m)	**234**	10.05	91.8
ADApprox (0.2m)	55	10.47	91.2
ADApprox (0.4m)	15	14.19	91.6
Quadtree	**42**	**10.16**	**91.3**
Quadtree w/Wind	**16**	11.29	90.3
APF	65	11.13	91.1
APF w/Wind	**62**	**10.97**	**90.6**