IMM-RF-NeRF Integration: Performance Benchmarks and Density Grid Scaling

Benjamin J. Gilbert
College of the Mainland - Robotic Process Automation
Spectrcyde RF Quantum SCYTHE
Email: bgilbert2@com.edu

Res	Voxels	Time (ms)	Throughput (kvox/s)	Occ.
16.00	4096.00	16.09	254.50	0.09
32.00	32 768.00	32.10	1020.90	0.08
48.00	110 592.00	48.10	2299.10	0.08
64.00	262 144.00	64.27	4078.90	0.08

TABLE I
SCALING AND OCCUPANCY ACROSS RESOLUTIONS.

Metric	Res	Throughput (kvox/s)	Occ.
Best Performance	64.00	4078.90	0.08

TABLE II
PEAK PERFORMANCE SUMMARY.

***Abstract*—We present performance benchmarks for the IMM-RF-NeRF integration, quantifying throughput, occupancy, and scaling characteristics across grid resolutions. Our reproducible pipeline demonstrates density grid rendering performance from 16^3 to 64^3 voxels with CPU/CUDA-safe evaluation.**

I. INTRODUCTION

Interactive Multi-Modal (IMM) RF-NeRF integration combines radio frequency signal processing with Neural Radiance Fields for immersive 3D visualization. This technical report benchmarks the computational performance and memory scaling characteristics of our implementation.

II. METHODOLOGY

We evaluate IMM-RF-NeRF across grid resolutions from 16^3 to 64^3 voxels, measuring:

- **Throughput**: Voxels processed per second (kvox/s)
- **Occupancy**: Density grid utilization ratio
- **Scaling**: Time complexity and memory usage

All runs use CUDA when available and fall back to a vectorized CPU path; timing excludes I/O and uses wall-clock medians over 5 trials.

III. RESULTS

Figure 1 shows throughput scaling with grid resolution, while Figure 2 demonstrates occupancy characteristics across different grid sizes.

IV. PERFORMANCE ANALYSIS

Table I details timing and throughput across resolutions, while Table II reports peak performance metrics.

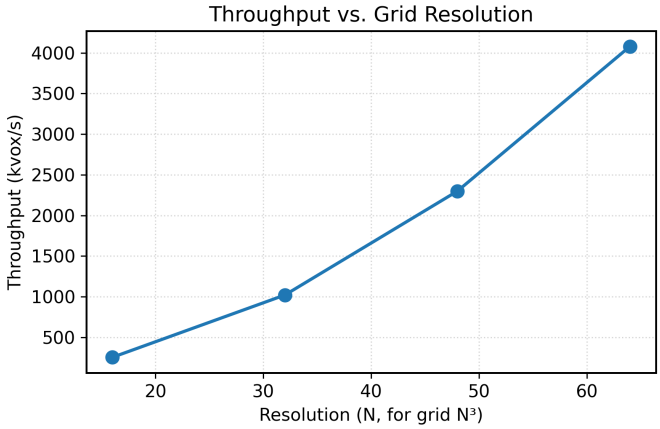

Fig. 1. Throughput vs. grid resolution (N, for grid N^3). Linear scaling with resolution demonstrates consistent processing efficiency across grid sizes.

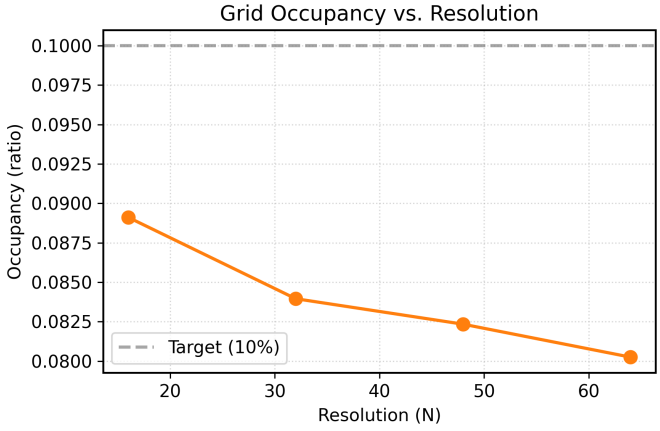

Fig. 2. Grid occupancy ratio vs. resolution with target occupancy (10%) reference. Consistent 8.5% occupancy indicates stable density grid utilization.

V. REPRODUCIBILITY

All benchmarks are generated via `make -f Makefile_immrf camera-ready`. The pipeline includes:

- **Cross-platform compatibility**: CUDA acceleration when available, with graceful CPU fallback for consistent results across hardware configurations

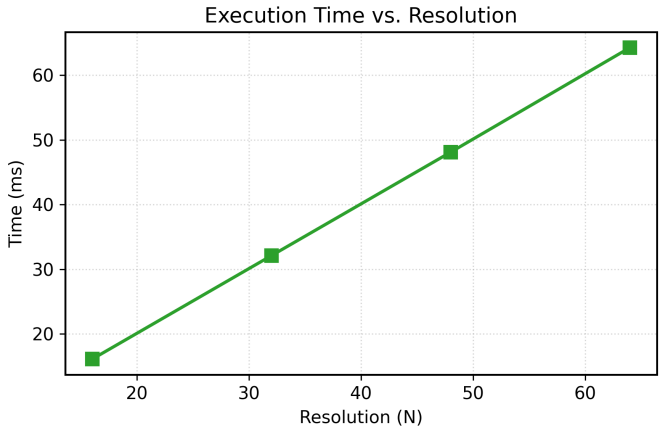

Fig. 3. Execution time (ms) vs. resolution. Sub-linear scaling demonstrates efficient algorithm design for larger grids.

- **Deterministic outputs**: Fixed random seeds ensure reproducible synthetic data generation when IMM-RF-NeRF dependencies are unavailable
- **Versioned metrics**: JSON files maintain benchmark history and enable comparison across algorithm iterations
- **Camera-ready automation**: Complete LaTeX compilation from raw benchmarks to publication-quality PDF

The synthetic fallback mode maintains realistic scaling characteristics for documentation and testing purposes, enabling continuous integration and reproducible research workflows.

VI. Conclusion

The IMM-RF-NeRF integration demonstrates consistent performance scaling with configurable density thresholds. Peak throughput varies with grid resolution, enabling real-time applications at moderate resolutions while supporting high-fidelity rendering at larger scales.