

MOPAC Performance Optimization Center for Large Molecules

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: Mopac

Cat. No.: B1609853

[Get Quote](#)

Welcome to the technical support center for optimizing MOPAC performance, specifically tailored for researchers, scientists, and drug development professionals working with large molecules. This guide provides answers to frequently asked questions and troubleshooting steps for common issues encountered during computational experiments.

Section 1: Frequently Asked Questions (FAQs)

This section addresses common questions regarding slow calculations and the fundamental tools available in MOPAC to accelerate performance for large systems.

Q1: My MOPAC calculation on a large molecule is running very slowly. What are the first things I should check?

A: When dealing with large molecules, several factors can drastically affect calculation time. The primary areas to investigate are the choice of SCF solver, the hardware utilization, and the semi-empirical method.


- Use the MOZYME Solver: For large systems like proteins (up to 15,000 atoms), the linear-scaling MOZYME algorithm is essential.^{[1][2]} Conventional MOPAC is typically limited to about 1,500 atoms.^[1] MOZYME replaces the standard SCF procedure with a localized molecular orbital (LMO) method, where the time required scales linearly with the system size.^{[2][3]}

- **Enable Parallel Processing:** By default, MOPAC utilizes multi-threading to accelerate jobs. Ensure you are using a version of MOPAC compiled to take advantage of multi-core CPUs or GPUs. Significant speedups can be achieved, especially on modern hardware.
- **Select an Appropriate Method:** While newer methods like PM7 are generally more accurate, they may have different performance characteristics. Ensure the chosen method is appropriate for your system and computational goals.
- **Check Your Geometry Input:** For systems with large rings, using internal coordinates can lead to instability and slow convergence. Switching to Cartesian coordinates with the XYZ keyword can often resolve this.

Q2: What is MOZYME and when should I use it?

A: MOZYME is a unique feature in MOPAC designed specifically for calculations on very large, closed-shell systems like enzymes and proteins. It employs a linear-scaling algorithm based on localized molecular orbitals (LMOs), which dramatically reduces the computational cost compared to traditional methods.

You should use MOZYME whenever you are working with systems that are too large for conventional SCF methods, typically those exceeding 1,500 atoms. It is the standard for geometry optimizations on biomolecules. However, a known limitation is that MOZYME requires a correctly identified Lewis structure to initialize its calculations.



[Click to download full resolution via product page](#)

Q3: How can I use multiple cores or GPUs to speed up my calculation?

A: MOPAC can be accelerated on high-performance computers by using shared-memory parallel strategies and GPU computing.

- **Multi-Core CPUs:** By default, MOPAC will use multi-threading to run a single job faster. For running a large number of separate jobs, it is more efficient to limit each job to a single thread by adding the keyword `THREADS=1` and running multiple jobs simultaneously.
- **GPUs:** Specialized versions of MOPAC can leverage NVIDIA GPU chips via the CUDA toolkit. This can dramatically accelerate calculations, with reported speedups of over 50 times compared to serial versions for certain calculations. Using a GPU approximately doubles the speed for large systems compared to a CPU-only job.

Experimental Protocol: Enabling Hardware Acceleration

- **Verify MOPAC Version:** Ensure you have a version of MOPAC compiled for multi-core or GPU usage. Versions labeled "CPU+GPU" will automatically detect and use a compatible GPU.

- **Install Dependencies:** For GPU acceleration, the NVIDIA CUDA Toolkit must be installed on your system.
- **Keyword Specification:**
 - For multi-core CPUs, MOPAC's default multi-threading is often sufficient for a single large job.
 - For GPU usage, no specific keyword is needed if you have a GPU-enabled version; it will be used automatically. To prevent the GPU from being used, add the NOGPU keyword.

Table 1: Reported Performance Gains with Hardware Acceleration

Hardware Component	Strategy	Reported Speedup	Applicable System Size
Multi-Core CPU	Multi-Threading	~3x with 4 threads	Large Systems
NVIDIA GPU	GPU Offloading	~2x (GPU vs. CPU)	Large Systems
Combined	MKL + Multi-Threading + GPU	Up to 200x (vs. legacy serial)	Large Systems (e.g., Bacteriorhodopsin)

[Data sourced from references 1 and 15]

Q4: Which semi-empirical method (PM6, PM7, etc.) should I choose for large systems?

A: The choice of method involves a trade-off between accuracy and computational cost. PM7 is the latest major reparameterization and is generally considered the most accurate for a wide range of systems. For large biological systems where non-covalent interactions are critical, dispersion and hydrogen-bond corrections are vital. Methods like PM6-D3H4 have shown excellent accuracy for interaction energies.

Table 2: Comparison of Average Unsigned Errors in Heats of Formation (kcal/mol)

Method	Average Unsigned Error	Root Mean Square Error
PM7	4.01	5.89
PM6	4.42	6.16
B3LYP 6-31G(d)	5.14	7.36
PM3	6.23	9.44
AM1	10.00	14.65

[For a set of 1,366 compounds containing C, H, O, N, F, Cl, S, P, and Br. Sourced from reference 3.]


For drug development professionals focusing on protein-ligand binding, achieving high accuracy is paramount. PM6-D3H4 has shown a lower average unsigned error for interaction energies (1.72 kcal/mol) compared to PM7 (2.91 kcal/mol) in some studies.

Section 2: Troubleshooting Geometry Optimization

This section provides guidance on resolving common problems encountered during the geometry optimization of large molecules.

Q5: My geometry optimization fails to converge. What are the common causes?

A: Convergence failure in large molecules can stem from several issues, including poor starting geometries, the complexity of the potential energy surface (PES), and issues with coordinate systems.



[Click to download full resolution via product page](#)

Q6: The geometry optimization terminates, but the gradient is still high. How can I achieve a tighter convergence?

A: The default geometry optimizers for large systems, like L-BFGS, may terminate prematurely several kcal/mol above the true energy minimum. To achieve a more reliable and tightly converged geometry, you can use the following keywords:

- GNORM=n.n: This keyword sets the termination criterion for the gradient norm. For high-quality geometries, a smaller value like GNORM=0.25 or lower is recommended, though this will increase calculation time.
- LET(nnn) or BIGCYCLES=n: For difficult optimizations, the default procedure may struggle. Using LET with a large value (e.g., LET(250)) can force the optimizer to take more steps and overcome small energy barriers, leading to a more complete optimization. This significantly increases the run time but improves the reliability of the final geometry.

Q7: My molecule has large rings and the optimization is unstable. What should I do?

A: This is a common issue known as the "big ring" problem. When using internal coordinates (the default), small changes in angles or dihedrals within a large ring can cause large, unrealistic changes in the distance between atoms that should be bonded. This leads to instability and optimization failure.

Solution: Add the keyword XYZ to your input file. This forces MOPAC to perform all geometric operations in Cartesian coordinates, which avoids the "big ring" problem.


Caution: Do not use the XYZ keyword if you are also using the SYMMETRY keyword or if some atomic coordinates are intentionally not being optimized (i.e., frozen).

Q8: I'm working with a protein from a PDB file. What are the essential preparation steps?

A: Using a raw PDB file directly in MOPAC will lead to incorrect results. PDB files must be carefully prepared to represent a chemically reasonable system.

Experimental Protocol: PDB File Preparation for MOPAC

- **Resolve Disorder:** Address any positional or structural disorder reported in the PDB file. This often involves choosing one of the modeled conformers for a given residue.
- **Add Hydrogen Atoms:** PDB files typically do not include hydrogen atoms, which are essential for any quantum chemical calculation. Use a reliable molecular modeling program to add hydrogens.
- **Verify Protonation States:** Carefully check the protonation states of titratable residues (e.g., His, Asp, Glu, Lys) and the ligand, as these are pH-dependent and critical for calculating accurate interaction energies.
- **Check Valency and Connectivity:** When a protein is first run, MOPAC will print errors related to the number and position of hydrogen atoms. Carefully read the output file, identify faulty residues, and correct them in your input structure. Pay special attention to non-residue hetero groups, which are not automatically checked.
- **Define the Active Site (Globule):** For very large proteins, it is often more efficient to model a "globule" – a smaller sphere of atoms around the active site where the chemistry of interest occurs. This significantly reduces computational cost while maintaining accuracy for the local environment.



[Click to download full resolution via product page](#)

Section 3: Advanced Performance Tuning

This section covers more specific keywords and techniques for expert users to further optimize calculations.

Q9: How do I choose the best SCF or geometry optimization algorithm?

A: MOPAC offers several algorithms for both the Self-Consistent Field (SCF) procedure and the geometry optimization. While the defaults are generally robust, certain keywords can be beneficial for difficult cases.

Table 3: Selected Keywords for Advanced Algorithm Control

Keyword	Type	Description	When to Use
CAMP	SCF Converger	A powerful but CPU-intensive SCF converger based on the Camp-King method.	Use when the default SCF procedure fails to converge.
BFGS	Optimizer	Uses the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm for geometry optimization.	An alternative to the default Eigenvector Following (EF) optimizer, can be more efficient for some systems.
EF	Optimizer	The default Eigenvector Following routine for finding minimum energy structures.	The standard choice for most geometry optimizations.
TS	Optimizer	Uses the EF routine specifically to locate transition states.	For reaction mechanism studies, not for ground-state geometry optimization.

Q10: What is the RAPID keyword and how does it work?

A: The RAPID keyword is a specialized tool used with MOZYME to accelerate geometry optimizations where only a small part of a very large molecule is being modified. For example,

if you are optimizing the side-chain of a single residue in a large protein, RAPID can provide a significant speedup.

Methodology: You must specify which atoms are being modified. Atoms being optimized are flagged with a "1", while all other static atoms are flagged with a "0". RAPID then focuses the computational effort on the changing part of the system.

Important Considerations:

- RAPID should not be used if the atoms being optimized are scattered throughout the molecule.
- Due to its approximations, the termination criterion should be loosened (e.g., GNORM=5).
- After a RAPID optimization, it is recommended to run a final, conventional MOZYME calculation on the entire system to obtain a more accurate final geometry and heat of formation.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. openmopac.net [openmopac.net]
- 2. MOPAC – MolSSI [molssi.org]
- 3. scm.com [scm.com]
- To cite this document: BenchChem. [MOPAC Performance Optimization Center for Large Molecules]. BenchChem, [2025]. [Online PDF]. Available at: [\[https://www.benchchem.com/product/b1609853#optimizing-mopac-performance-for-large-molecules\]](https://www.benchchem.com/product/b1609853#optimizing-mopac-performance-for-large-molecules)

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com