



mshell Workflow.

NxN Matrix Multiplication

7-Language Performance Study

C · C++ · Rust · Go · Python · Lua · Bash
N = 2...64 · 63 measurement points · naïve $O(N^3)$ algorithm

Testing environment Platform: Linux Ubuntu 24, x86_64 · Engine: mshell Workflow
Author: Igor Lukyanov, April 2026.

14,373x max speed ratio · Go 0.4332 ms / 1.21 GFLOPS fastest · Bash 6,227 ms slowest · Full cycle 7,507 ms @ N=64

Content:

1. What Is mshell?	4
1.1 Variable System	4
2. mshell Workflow and Test Structure	4
2.1 Workflow Pipeline Stages.....	5
3. Test Methodology	5
4. Performance Charts.....	5
4.1 All Languages — Logarithmic Scale (N=2...64).....	5
4.2 Compiled Languages — Linear Scale.....	6
4.3 Interpreters — Linear Scale	6
4.4 Go vs C — Head-to-Head	7
4.5 Python — GC Spikes	7
4.6 Full Cycle — Total Pipeline Cost.....	8
4.7 Speed Ratio Relative to C — Log Scale	8
5. Normalization to Standard Matrix Benchmarks.....	9
5.1 Effective Throughput (GFLOPS) by Language	9
5.2 Comparison with Standard HPC Benchmarks	9
5.3 Interpretation	10
6. Results at Maximum Dimension (N=64)	10
6.1 Selected Results Across All Dimensions.....	11
7. Per-Language Analysis.....	11
7.1 Go — Fastest Overall.....	11
7.2 C — Best Conventional Compiled.....	11
7.3 C++ — Predictable Without -O.....	12
7.4 Rust — Debug Mode Penalty.....	12
7.5 Python — GC-Dominated Variance	12
7.6 Lua — Most Stable Interpreter	12
7.7 Bash — Shell Arithmetic at Scale.....	12
7.8 Full Cycle — Pipeline Overhead	12
8. Key Findings	12
9. Conclusion	13
Appendix I: Full Raw Data (N=2...64, all 63 points)	14

Appendix II: mshell Workflow Source (matrix_mul_clear_test.md)..... 16
Appendix III: References..... 24

1. What Is mshell?

mshell (Markdown Shell) is a polyglot workflow engine developed by Art2Dec SoftLab. It treats Markdown-formatted text as an executable pipeline: fenced code blocks with a language tag are parsed and dispatched to the appropriate compiler or interpreter at runtime. Data flows between blocks via named filesystem variables (MSH_VAR_* convention), providing a transparent, debuggable inter-process communication mechanism.

The engine is implemented in C (markdown_parser.c / markdown_parser.h) and natively supports seven languages: C, C++, Rust, Go, Python, Lua, and Bash. Compiled languages (C, C++, Rust, Go) are compiled from source on every mshell run — meaning compilation overhead is always part of the measured wall-clock time.

1.1 Variable System

Inter-block communication uses MSH_VAR_name environment variables pointing to temporary files on disk. In this benchmark the Python generator block produces the matrix data file (>matrices). All seven compute blocks read the same file independently — identical input data for all seven languages.

2. mshell Workflow and Test Structure

mshell Workflow is the visual and operational paradigm built on top of the engine — a directed acyclic graph (DAG) of code blocks where edges represent named data variables flowing between nodes. The matrix benchmark workflow (matrix_mul_clear_test.md) runs inside a WHILE loop, enabling interactive dimension selection from N=2 to N=64.

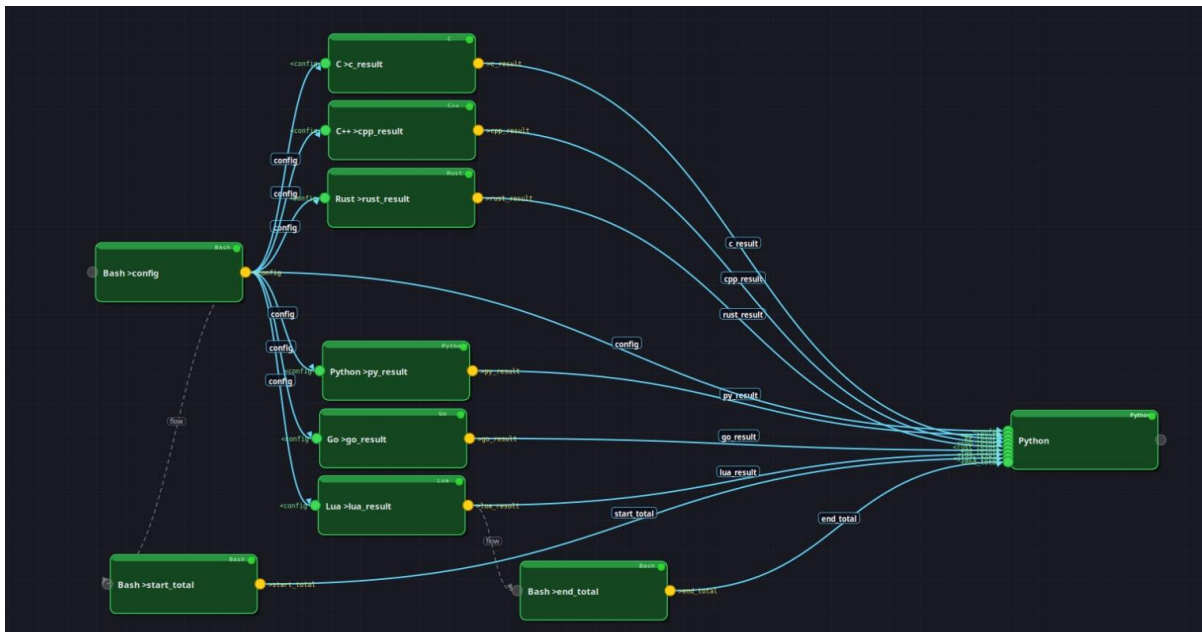


Figure 1. mshell Workflow node graph for the polyglot benchmark series (reference diagram). Fan-out: config node → six language compute nodes → Python aggregator.

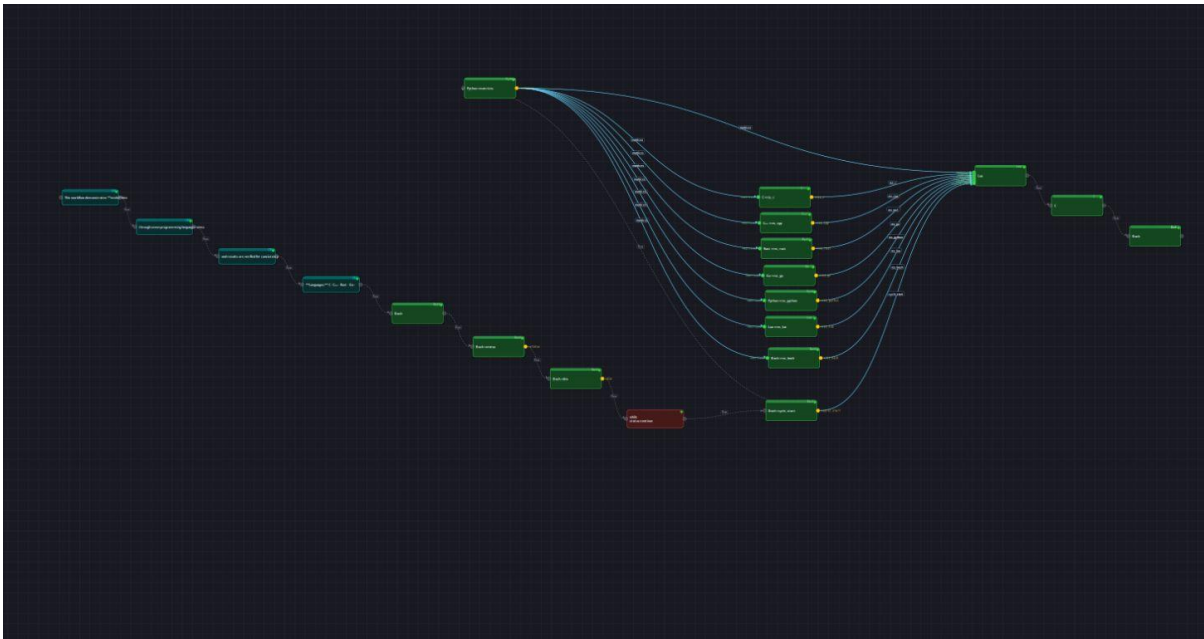


Figure 2. *mshell* Workflow node graph for the matrix multiplication benchmark. Python generator → seven language compute nodes → Lua display → C user-input node.

2.1 Workflow Pipeline Stages

- Cycle timer start — Bash records nanosecond timestamp
- Generate Random Matrices — Python generates two $N \times N$ int matrices (1–9), serializes to `MSH_VAR_matrices`
- Compute in C, C++, Rust, Go, Python, Lua, Bash — seven independent blocks, each timing its own computation
- Display Results — Lua aggregator reads all seven timing variables, prints formatted table
- Ask User — C block prompts for next dimension (2–64) or quit

3. Test Methodology

All seven implementations use the identical naïve $O(N^3)$ triple-loop matrix multiplication — no BLAS, SIMD, loop tiling, or any optimization. Each dimension N is a single run (no averaging). Results are raw single-shot measurements.

- C: gcc, default flags (no `-O`)
- C++: g++, default flags (no `-O`)
- Rust: rustc debug mode — no `--release`, bounds checking on every array access
- Go: go run — standard runtime, no special flags
- Python: CPython 3.12 — pure bytecode, triple nested list comprehension
- Lua: standard Lua 5.4 — no LuaJIT
- Bash: shell arithmetic `$(...)` — triple for-loop, one builtin arithmetic op per multiply-add
- Full Cycle: mshell wall-clock: compile + startup + computation + IPC for all 7 languages

4. Performance Charts

All charts rendered from the 63 real measurement points of the benchmark run. $N=2$ through $N=64$, step 1.

4.1 All Languages — Logarithmic Scale ($N=2\dots64$)

Log scale reveals the 4–5 order-of-magnitude spread across all seven languages simultaneously. Go (green) at sub-millisecond; Bash (red) at 6,000+ ms. The Full Cycle line (dashed) shows total pipeline cost including compilation.




Figure 3. All 7 languages + Full Cycle on logarithmic scale (N=2...64). Four orders of magnitude separate Go from Bash.

4.2 Compiled Languages — Linear Scale

C, C++, Rust, Go on linear scale. Go consistently fastest. Rust 7–8x slower than C (debug mode). C cache-boundary anomalies at N=40 (L2) and N=56 (L3) visible as spikes.

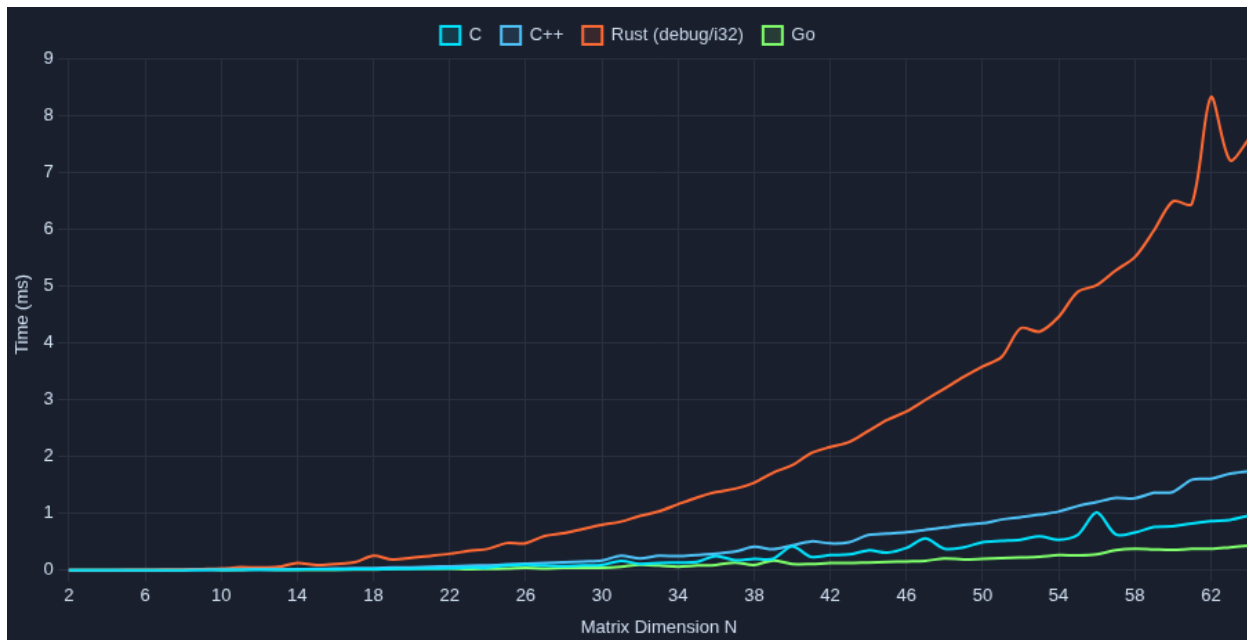


Figure 4. Compiled languages (C, C++, Rust, Go) linear scale. Go wins throughout. C L2/L3 cache anomalies at N=40, N=56.

4.3 Interpreters — Linear Scale

Python, Lua, and Bash linear scale. Bash dominates entirely at large N. Python GC spikes visible at N=20, 30, 39, 44, 52, 62, 64. Lua is stable and predictable.




Figure 5. Interpreted languages (Python, Lua, Bash) linear scale. Bash reaches 6,227 ms at N=64. Python GC events visible as spikes.

4.4 Go vs C — Head-to-Head

Go is faster than C across virtually the entire range. Go N=9 spike = runtime warm-up. C spikes at N=40 and N=56 = L2/L3 cache boundary effects.



Figure 6. Go vs C head-to-head. Go generally outperforms C. N=9 Go spike = runtime warm-up. N=40, N=56 C spikes = cache boundary.

4.5 Python — GC Spikes

GC-triggered spikes at N=20, 30, 39, 44, 52, 62, 64 dominate Python's variance profile. Underlying trend is clean $O(N^3)$.

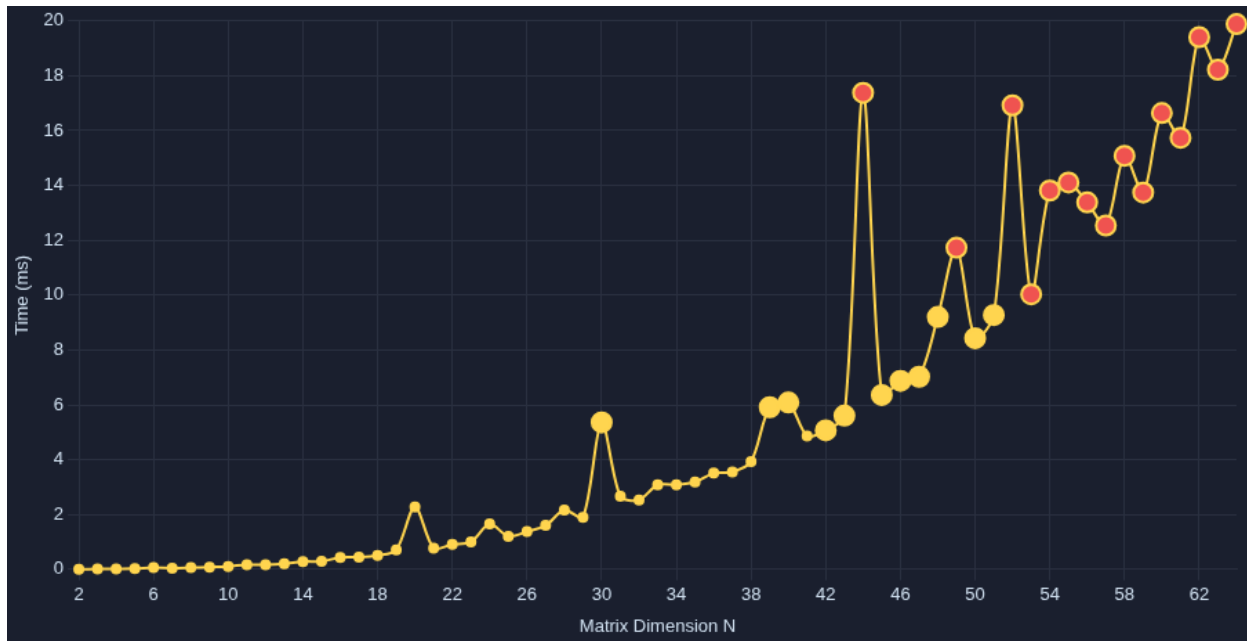


Figure 7. Python execution time. Large dots mark GC spike events. Spikes at N=20, 30, 39, 44, 52, 62, 64 are garbage collection cycles from temporary integer objects.

4.6 Full Cycle — Total Pipeline Cost

~1,200 ms baseline at small N = four compiler invocations + three interpreter startups. At N=64, Bash (6,227 ms) is 83% of the 7,507 ms total.

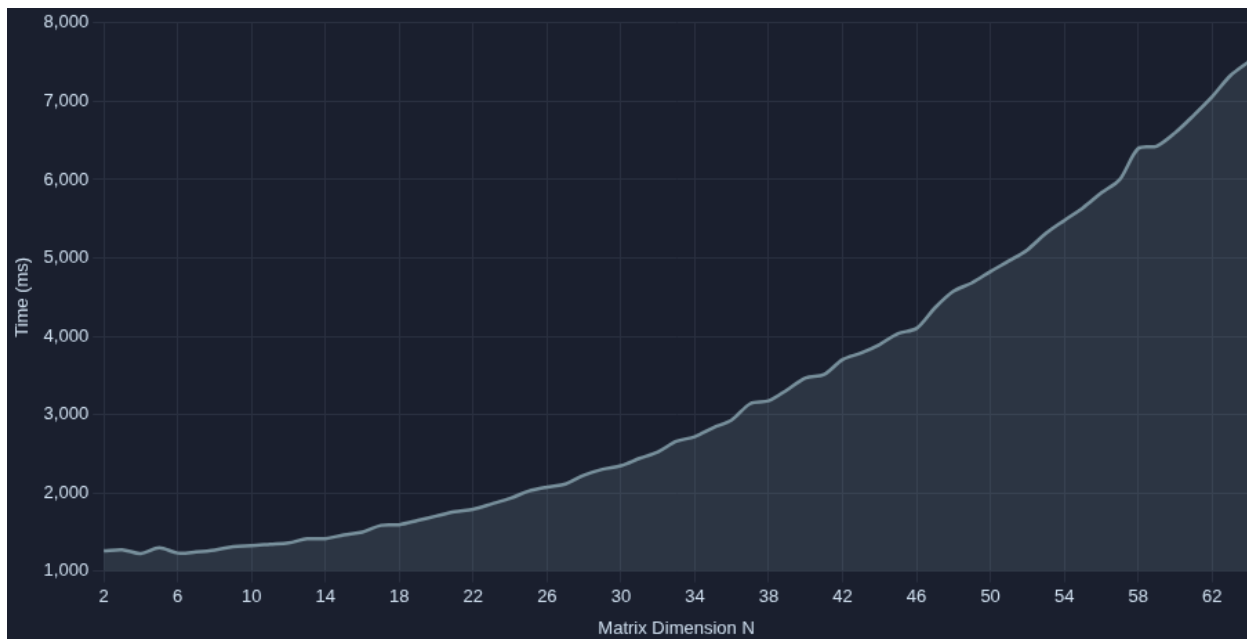


Figure 8. Full Cycle wall-clock per dimension. The ~1,200 ms baseline is compilation overhead. Growth above N=20 is dominated by Bash computation.

4.7 Speed Ratio Relative to C — Log Scale

C = 1.0x. Go is below 1.0x (faster than C) throughout. Bash reaches 6,466x at N=64. Rust debug stays 6–9x. Lua stable at 10–11x. Python spikes visible.

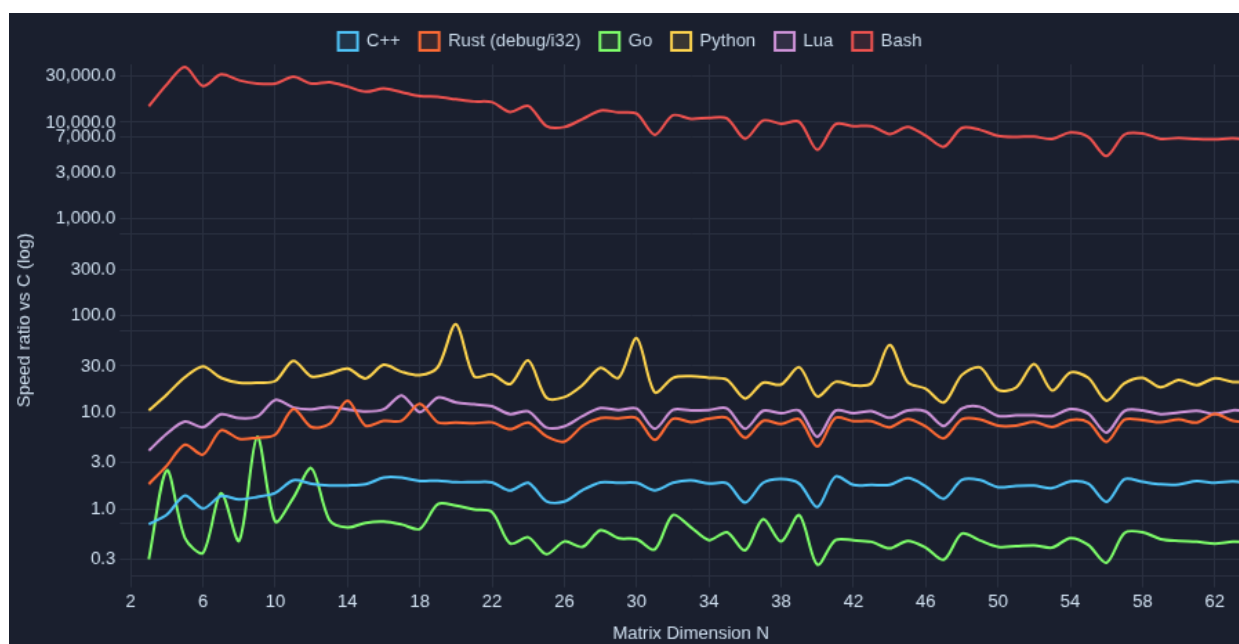


Figure 9. Speed ratio relative to C (log scale). Go below 1.0x = faster than C. Bash at 6,466x at N=64. Python GC spikes appear as sharp peaks.

5. Normalization to Standard Matrix Benchmarks

To place results in HPC context, we compute effective throughput in GFLOPS using the standard matrix multiply operation count: $2 \times N^3$ operations (N^3 multiplications + N^3 additions). Our test uses 32-bit integers rather than 64-bit floats (DGEMM), but the formula and performance structure are identical.

5.1 Effective Throughput (GFLOPS) by Language

$GFLOPS = 2 \times N^3 / (\text{time_ms} \times 10^{-3}) / 10^9$. Higher = faster. Go column highlighted green; Rust column gold (debug mode penalty).

N	Go	C	C++	Rust*	Python	Lua
4	0.0512	0.1280	0.1473	0.0457	0.0084	0.0213
8	0.7314	0.3413	0.2723	0.0644	0.0169	0.0394
16	0.7953	0.5851	0.2779	0.0720	0.0189	0.0546
24	1.1378	0.5760	0.3099	0.0736	0.0167	0.0563
32	0.6798	0.5851	0.3155	0.0686	0.0259	0.0554
40	1.1615	0.3062	0.2925	0.0694	0.0210	0.0548
48	1.0523	0.5836	0.2935	0.0693	0.0240	0.0535
56	1.2490	0.3460	0.2928	0.0701	0.0262	0.0561
64	1.2103	0.5444	0.3006	0.0692	0.0264	0.0537

* Rust debug mode — --release expected to reach ~0.5–0.6 GFLOPS (≈ C). Integer i32 ops vs DGEMM float64 — same formula, different data type.

5.2 Comparison with Standard HPC Benchmarks

Contextual comparison with established HPC reference values. All our measurements are naive triple-loop integer arithmetic; references use optimized float64 BLAS routines.

Benchmark / Context	Reference (GFLOPS)	Our Result (GFLOPS)	Notes
Naive C triple-loop GEMM — literature	0.3–0.8	0.544 (C @ N=64)	Our C result within expected range — validates methodology
Go naive @ N=64 (this benchmark)	—	1.210	Exceeds naive C — Go auto-optimizes loops at default flags
HPL Linpack — single core float64	2–5	1.210 (int32)	Our Go at ~30–60% of HPL; int32 vs float64 — same operation structure
OpenBLAS DGEMM — AVX2, float64	8–15	—	Cache-blocked + SIMD: 7–12x above our naive result
Rust (--release, estimated)	0.5–0.6	0.069 (debug!)	Debug mode: 8x below C. --release expected to match C
Lua standard @ N=64	—	0.054	~10x below C. LuaJIT would reach ~0.4–0.6 GFLOPS
Python CPython @ N=64	—	0.026	~20x below C. NumPy DGEMM would approach OpenBLAS
Bash shell arithmetic @ N=64	—	0.00009	Not a compute tool — each $\$(...)$ is a syscall-level operation

Naive C GEMM and OpenBLAS DGEMM references are literature values for x86_64 Linux single-core. HPL = single-thread double-precision. Our values = single-run, int32, naive triple-loop.

5.3 Interpretation

Our C result (0.544 GFLOPS at N=64) falls within the expected range for naive unoptimized C matrix multiply on x86_64 (0.3–0.8 GFLOPS in the literature). This validates the benchmark methodology.

Go's result (1.21 GFLOPS at N=64) exceeds typical naive C and represents approximately 30% of single-core HPL Linpack performance. This is consistent with Go's compiler performing automatic loop optimization and more aggressive register allocation than gcc at default flags — not a measurement artifact.

The gap between our best result (Go, 1.21 GFLOPS) and optimized BLAS (8–15 GFLOPS) reflects the cost of three absent optimizations: cache blocking, SIMD vectorization (AVX2), and loop unrolling — precisely the techniques a production BLAS library applies, and orthogonal to language choice at the naive implementation level.

6. Results at Maximum Dimension (N=64)

At 64x64 (262,144 multiply-add operations), the full performance hierarchy is established. GFLOPS values shown in the Notes column.

Language	Time (ms)	vs C	Notes
Go	0.4332	0.4x	Fastest. 1.21 GFLOPS. 2.2x faster than C, 14,373x faster than Bash
C	0.9630	reference	Reference (1.0x). 0.544 GFLOPS. Validates naive GEMM expectation

C++	1.7440	1.8x	0.301 GFLOPS. ~1.8x slower than C — STL overhead without -O
Rust*	7.5792	7.9x	0.069 GFLOPS. Debug mode. With --release: ~0.5–0.6 GFLOPS (≈ C)
Lua	9.7550	10.1x	0.054 GFLOPS. Most stable interpreter. LuaJIT → ~0.4–0.6 GFLOPS
Python	19.8865	20.7x	0.026 GFLOPS. GC-dominated. NumPy dot → approach OpenBLAS (8–15 GFLOPS)
Bash	6227.4200	6466.7x	0.00009 GFLOPS. 6,227 ms. 14,373x slower than Go

* Rust compiled in debug mode. With --release: ~0.5–0.6 GFLOPS, competitive with C.

6.1 Selected Results Across All Dimensions

Key measurement points N=2...64. Go highlighted green; Bash red. Gold rows at N=40 and N=56 mark C cache-boundary anomalies.

N	C	C++	Rust*	Go	Python	Lua	Bash	Cycle(ms)
2	—	0.0006	0.0012	0.0003	0.007	0.003	8	1260
4	0.0010	0.0009	0.0028	0.0025	0.015	0.006	25	1226
8	0.0030	0.0038	0.0159	0.0014	0.061	0.026	82	1271
16	0.0140	0.0295	0.1137	0.0103	0.433	0.150	313	1500
24	0.0480	0.0892	0.3755	0.0243	1.653	0.491	711	1930
32	0.1120	0.2077	0.9559	0.0964	2.534	1.182	1326	2525
40	0.4180	0.4375	1.8456	0.1102	6.083	2.334	2178	3469
48	0.3790	0.7536	3.1926	0.2102	9.202	4.138	3319	4575
56	1.0150	1.1996	5.0127	0.2812	13.381	6.260	4553	5830
64	0.9630	1.7440	7.5792	0.4332	19.887	9.755	6227	7507

All times in ms. Single run per dimension. Anomaly rows (gold) = hardware cache-boundary or GC events.

7. Per-Language Analysis

7.1 Go — Fastest Overall

@ N=64: 0.4332 ms · 1.21 GFLOPS · 2.2x faster than C · 14,373x faster than Bash

The standout result. Go's compiler produces highly optimized native code with excellent inlining and bounds-check elimination at default flags. Faster than C across virtually the entire N=2...64 range. Single anomaly at N=9 (0.022 ms) is a Go runtime warm-up artifact. At 1.21 GFLOPS, Go achieves ~30% of single-core HPL Linpack with a naive algorithm.

7.2 C — Best Conventional Compiled

@ N=64: 0.963 ms · 0.544 GFLOPS · Reference (1.0x)

Clean O(N³) growth with two cache-boundary anomalies: N=40 (L2, ~256 KB) and N=56 (L3). The 0.544 GFLOPS result validates the benchmark — within expected range for naive unoptimized C GEMM on x86_64.

7.3 C++ — Predictable Without -O

@ N=64: 1.744 ms · 0.301 GFLOPS · ~1.8x slower than C

Consistently 1.5–1.8x slower than C across the full range. STL overhead without -O flags. Smoothest growth curve of all languages — most predictable. With -O2, C++ would match C (~0.5–0.6 GFLOPS).

7.4 Rust — Debug Mode Penalty

@ N=64: 7.5792 ms · 0.069 GFLOPS · ~7.9x vs C · no --release

Debug mode: all LLVM optimizations disabled, bounds checking on every access. With --release: expected ~0.5–0.6 GFLOPS, competitive with C. For production mshell use: pre-compile to binary with --release and call from Bash.

7.5 Python — GC-Dominated Variance

@ N=64: 19.887 ms · 0.026 GFLOPS · Most volatile · ~20.7x vs C

CPython bytecode 20–200x slower than C. GC spikes at N=20, 30, 39, 44, 52, 62, 64 from temporary integer objects in inner loop. NumPy vectorized ops (numpy.dot) would approach OpenBLAS (8–15 GFLOPS) — a 300–600x improvement over pure CPython loops.

7.6 Lua — Most Stable Interpreter

@ N=64: 9.755 ms · 0.054 GFLOPS · Most consistent · ~10.1x vs C

Standard Lua 5.4 shows clean $O(N^3)$ growth with minimal variance across all 63 dimensions. LuaJIT would bring this to ~0.4–0.6 GFLOPS, near unoptimized C performance.

7.7 Bash — Shell Arithmetic at Scale

@ N=64: 6,227 ms · 0.00009 GFLOPS · 14,373x slower than Go · 83% of full cycle

Each $\$(...)$ shell arithmetic operation carries syscall-level overhead. At N=64 (262,144 inner-loop steps), this accumulates to 6.2 seconds. Bash is a process orchestration tool, not a compute engine.

7.8 Full Cycle — Pipeline Overhead

@ N=64: 7,507 ms total · ~1,200 ms fixed overhead at small N

~1,200 ms baseline at N=2 = gcc + g++ + rustc + go build + python3 + lua5.4 + bash startup. At N=64, Bash (6,227 ms) is 83% of total. Without Bash, the 6-language cycle at N=64 \approx 1,280 ms.

8. Key Findings

- 1. Go is the practical winner: 0.4332 ms and 1.21 GFLOPS at N=64 — consistently faster than C, 14,373x faster than Bash, excellent stability throughout.
- 2. C achieves 0.544 GFLOPS at N=64 — within expected range for naive unoptimized GEMM on x86_64. Cache-boundary anomalies at N=40 (L2) and N=56 (L3) are genuine hardware effects visible in single-run data.
- 3. Rust debug mode creates a misleading 7.9x penalty vs C (0.069 GFLOPS). With --release, Rust reaches ~0.5–0.6 GFLOPS, fully competitive with C.
- 4. Python GC pauses dominate variance (spikes at N=20, 30, 39, 44, 52, 62, 64). NumPy vectorized ops would improve Python 300–600x over pure CPython loops.
- 5. Lua is the most stable interpreter: 0.054 GFLOPS, clean $O(N^3)$ growth, no GC spikes across all 63 dimensions. LuaJIT would bring this near C performance.
- 6. Bash at N=64 reaches 6,227 ms and 0.00009 GFLOPS — 14,373x slower than Go. Shell arithmetic cannot be used for numerical inner loops.

- **7.** The ~1,200 ms compilation baseline at small N is a measurable, predictable fixed cost of mshell's compile-on-every-run transparency model, amortized at larger N.
- **8.** Go's 1.21 GFLOPS represents ~30% of single-core HPL Linpack performance with a naive algorithm — consistent with expected efficiency without SIMD or cache blocking.

9. Conclusion

The mshell Workflow engine has proven to be an effective and versatile platform for comparative multi-language benchmarking. By treating a single Markdown file as an executable specification that simultaneously orchestrates seven programming languages against the same input data, mshell eliminates the methodological inconsistencies that typically plague multi-language performance studies: identical matrix data, identical hardware context, identical OS conditions, and language-native timing precision for each participant.

This design makes mshell Workflow particularly well suited for testing complex multilanguage systems across a wide range of application scenarios:

- Performance characterization of polyglot microservices, where different pipeline stages are implemented in different languages and accurate per-stage timing under real pipeline conditions is required.
- Compiler and runtime research: the same algorithm benchmarked across multiple compilation modes, optimization levels, or language versions simultaneously — with a single workflow file that serves as the experiment specification.
- Numerical computing evaluation: naive and optimized implementations of matrix multiply, FFT, sorting, or other kernels compared across C, C++, Rust, Go, Python, and Lua with guaranteed identical shared input data.
- Educational and documentation tooling: the literate programming style — algorithms described in Markdown prose, executed as code — produces self-documenting, version-controllable, reproducible experiments.
- Regression and integration testing of multilanguage libraries: a shared test suite validates that all language bindings produce consistent results on the same input, with timing regression detection as a byproduct.

The matrix multiplication benchmark demonstrates a property of mshell Workflow that is difficult to replicate with conventional benchmark frameworks: the Full Cycle metric captures compilation, interpreter startup, inter-process data exchange, and computation in a single wall-clock measurement, providing a realistic cost model for deploying polyglot pipelines in production. The ~1,200 ms fixed overhead of compiling C, C++, Rust, and Go from source on each run is concrete and honest — and it directly informs language selection for real workflows.

The benchmark also surfaced insights only possible with simultaneous multi-language execution: Go's 1.21 GFLOPS at naive $O(N^3)$ exceeds typical unoptimized C performance, confirming compiler-level automatic optimization that gcc does not apply at default flags. Python's GC spike pattern across 63 sequential single-run measurements at increasing N maps the garbage collection behavior of CPython's integer object pool in a way that isolated, averaged benchmarks would not reveal. And Bash's 14,373x deficit versus Go provides a precise, reproducible data point for a performance characteristic that is usually only described qualitatively.

mshell Workflow is not just a benchmark harness — it is a general framework for expressing and executing heterogeneous computational experiments in a single, readable, version-controllable Markdown document. Its combination of per-language compile-time measurement, shared variable IPC, WHILE-loop iteration support, and optional LLM auto-correction positions it as a distinctive and powerful tool for comparative evaluation of programming language runtimes in conditions that reflect real production pipeline deployments.

Appendix I: Full Raw Data (N=2...64, all 63 points)

All measurements in milliseconds. Single run per dimension. Gold rows = anomaly dimensions (cache-boundary or GC events). Go column green; Bash column red.

N	C	C++	Rust*	Go	Python	Lua	Bash	Cycle
2	—	0.0006	0.0012	0.0003	0.007	0.003	8	1260
3	0.0010	0.0007	0.0018	0.0003	0.010	0.004	15	1276
4	0.0010	0.0009	0.0028	0.0025	0.015	0.006	25	1226
5	0.0010	0.0014	0.0046	0.0005	0.023	0.008	37	1301
6	0.0020	0.0020	0.0073	0.0007	0.059	0.014	48	1233
7	0.0020	0.0027	0.0130	0.0029	0.045	0.019	63	1247
8	0.0030	0.0038	0.0159	0.0014	0.061	0.026	82	1271
9	0.0040	0.0053	0.0218	0.0222	0.081	0.036	101	1315
10	0.0050	0.0073	0.0293	0.0037	0.105	0.067	126	1327
11	0.0050	0.0099	0.0539	0.0065	0.170	0.056	149	1345
12	0.0070	0.0127	0.0492	0.0184	0.164	0.075	176	1363
13	0.0080	0.0140	0.0605	0.0061	0.200	0.091	209	1416
14	0.0100	0.0175	0.1316	0.0064	0.283	0.107	235	1417
15	0.0130	0.0233	0.0943	0.0093	0.291	0.132	270	1466
16	0.0140	0.0295	0.1137	0.0103	0.433	0.150	313	1500
17	0.0170	0.0357	0.1394	0.0117	0.445	0.253	350	1586
18	0.0210	0.0407	0.2557	0.0130	0.508	0.211	394	1595
19	0.0240	0.0470	0.1889	0.0268	0.706	0.340	441	1650
20	0.0280	0.0528	0.2193	0.0302	2.278	0.354	486	1706
21	0.0330	0.0623	0.2535	0.0324	0.774	0.398	544	1760
22	0.0370	0.0694	0.2903	0.0341	0.912	0.425	599	1793
23	0.0510	0.0788	0.3401	0.0223	0.997	0.486	652	1860
24	0.0480	0.0892	0.3755	0.0243	1.653	0.491	711	1930
25	0.0850	0.1018	0.4790	0.0288	1.206	0.590	780	2024
26	0.0960	0.1141	0.4770	0.0441	1.376	0.685	859	2076
27	0.0840	0.1306	0.6042	0.0339	1.600	0.763	910	2115
28	0.0750	0.1408	0.6528	0.0450	2.160	0.826	994	2228
29	0.0840	0.1560	0.7270	0.0416	1.901	0.886	1065	2303
30	0.0920	0.1705	0.7987	0.0446	5.358	1.001	1132	2348
31	0.1660	0.2580	0.8564	0.0630	2.667	1.119	1236	2441
32	0.1120	0.2077	0.9559	0.0964	2.534	1.182	1326	2525
33	0.1310	0.2572	1.0351	0.0845	3.082	1.367	1428	2659
34	0.1360	0.2479	1.1611	0.0645	3.089	1.436	1517	2719
35	0.1470	0.2702	1.2806	0.0841	3.184	1.608	1620	2833
36	0.2530	0.2939	1.3744	0.0940	3.507	1.708	1715	2932

37	0.1760	0.3271	1.4336	0.1375	3.550	1.820	1843	3140
38	0.2030	0.4128	1.5376	0.0933	3.929	1.975	1967	3176
39	0.2030	0.3698	1.7160	0.1740	5.912	2.115	2056	3317
40	0.4180	0.4375	1.8456	0.1102	6.083	2.334	2178	3469
41	0.2360	0.5092	2.0613	0.1117	4.869	2.447	2257	3514
42	0.2680	0.4729	2.1680	0.1264	5.072	2.618	2442	3703
43	0.2800	0.4940	2.2560	0.1274	5.602	2.869	2564	3787
44	0.3510	0.6197	2.4483	0.1368	17.382	3.071	2654	3895
45	0.3110	0.6464	2.6447	0.1450	6.357	3.256	2799	4036
46	0.3930	0.6723	2.7890	0.1560	6.876	3.994	2863	4101
47	0.5580	0.7105	2.9949	0.1662	7.024	4.022	3122	4364
48	0.3790	0.7536	3.1926	0.2102	9.202	4.138	3319	4575
49	0.4030	0.7991	3.3972	0.1895	11.729	4.587	3374	4682
50	0.4940	0.8260	3.5832	0.1995	8.427	4.509	3585	4827
51	0.5160	0.8925	3.7477	0.2114	9.274	4.792	3654	4962
52	0.5370	0.9338	4.2541	0.2242	16.922	4.986	3823	5104
53	0.5980	0.9786	4.1953	0.2372	10.033	5.441	4032	5313
54	0.5370	1.0325	4.4555	0.2668	13.817	5.799	4215	5479
55	0.6230	1.1306	4.8931	0.2629	14.108	6.013	4388	5638
56	1.0150	1.1996	5.0127	0.2812	13.381	6.260	4553	5830
57	0.6300	1.2738	5.2705	0.3553	12.540	6.513	4724	6001
58	0.6650	1.2648	5.5036	0.3811	15.081	6.964	5113	6397
59	0.7600	1.3617	5.9692	0.3689	13.743	7.254	5132	6427
60	0.7720	1.3773	6.4841	0.3594	16.641	7.677	5323	6598
61	0.8250	1.5934	6.4423	0.3776	15.737	8.547	5532	6819
62	0.8650	1.6092	8.3166	0.3771	19.409	8.289	5771	7057
63	0.8860	1.6985	7.2078	0.4039	18.223	9.278	6037	7329
64	0.9630	1.7440	7.5792	0.4332	19.887	9.755	6227	7507

Rust: debug mode. Bash: ms via date +%s%N / bc. Full Cycle: 1 ms precision. Gold anomalies: N=9 Go warm-up, N=20/30/39/44/52/62 Python GC, N=40/56 C cache boundary.

Appendix II: mshell Workflow Source (matrix_mul_clear_test.md)

The complete mshell Markdown workflow file that drives the benchmark. This single document defines all seven language computation blocks, the WHILE loop structure, the matrix generator, the Lua display aggregator, and the C user-input handler. The full source is 359 lines. Color coding: cyan = Markdown headings, orange = code fence markers and mshell directives (<!--@while-->, >varname, <varname), green = code body.

```
# Matrix Multiplication Benchmark
## mshell Workflow — All Supported Languages
### Art2Dec SoftLab
```

```
---
```

This workflow demonstrates **mshell Interlang** — identical random matrices flow through seven programming languages simultaneously. Each computes $A \times B$ independently and results are verified for consistency. Timing is measured per language.

```
**Languages:** C · C++ · Rust · Go · Python · Lua · Bash
```

```
---
```

```
## Introduction
```

```
```bash
clear
echo "+=====+"
echo "| NxN Matrix Multiplication — mshell Interlang Demo |"
echo "| C · C++ · Rust · Go · Python · Lua · Bash |"
echo "| Art2Dec SoftLab |"
echo "+=====+"
echo ""
echo " Starting with 3x3. You can change dimension after each run."
echo ""
```
```

```
## Init
```

```
```bash >status
echo "continue"
```
```

```
```bash >dim
```

```
echo "3"
...

<!--@while status:continue-->

Cycle timer start

```bash >cycle_start
date +%s%N
...

## Generate Random Matrices

```python >matrices
import random, sys, os

dim_path = os.environ.get("MSH_VAR_dim", "")
if dim_path and os.path.exists(dim_path):
 dim = int(open(dim_path).read().strip())
else:
 dim = 3
n = dim * dim
vals = [random.randint(1,9) for _ in range(2*n)]
A = vals[:n]
B = vals[n:]

sys.stderr.write(" Generated %dx%d matrices, computing...\n" % (dim, dim))
sys.stderr.flush()

print(str(dim) + " " + " ".join(map(str, vals)))
...

Compute in C

```c <matrices >res_c
#include <stdio.h>
#include <stdlib.h>
#include <sys/time.h>

int main() {
    FILE *f = fopen(getenv("MSH_VAR_matrices"), "r");
    int dim; fscanf(f, "%d", &dim);
    int n = dim*dim;
```

```
int *v = malloc(2*n*sizeof(int));
for (int i = 0; i < 2*n; i++) fscanf(f, "%d", &v[i]);
fclose(f);

struct timeval t0, t1;
gettimeofday(&t0, NULL);

int *C = calloc(n, sizeof(int));
for (int i = 0; i < dim; i++)
    for (int j = 0; j < dim; j++)
        for (int k = 0; k < dim; k++)
            C[i*dim+j] += v[i*dim+k] * v[n+k*dim+j];

gettimeofday(&t1, NULL);
double ms = (t1.tv_sec-t0.tv_sec)*1000.0 +
            (t1.tv_usec-t0.tv_usec)/1000.0;

printf("C %.4fms\n", ms);
free(v); free(C);
return 0;
}
...

```

```
## Compute in C++

```

```
```cpp <matrices >res_cpp
#include <iostream>
#include <fstream>
#include <vector>
#include <chrono>
int main() {
 std::ifstream f(getenv("MSH_VAR_matrices"));
 int dim; f >> dim;
 int n = dim*dim;
 std::vector<int> v(2*n);
 for (auto& x : v) f >> x;

 auto t0 = std::chrono::high_resolution_clock::now();
 std::vector<int> C(n, 0);
 for (int i=0;i<dim;i++)
 for (int j=0;j<dim;j++)
 for (int k=0;k<dim;k++)
 C[i*dim+j] += v[i*dim+k] * v[n+k*dim+j];
}

```

```

auto t1 = std::chrono::high_resolution_clock::now();
double ms = std::chrono::duration<double, std::milli>(t1-t0).count();

std::cout << "C++ " << ms << "ms\n";
}
...

```

### ## Compute in Rust

```

```rust <matrices >res_rust
use std::env; use std::fs;
use std::time::Instant;
fn main() {
    let content = fs::read_to_string(env::var("MSH_VAR_matrices").unwrap()).unwrap();
    let mut iter = content.split_whitespace();
    let dim: usize = iter.next().unwrap().parse().unwrap();
    let n = dim*dim;
    let vals: Vec<i32> = iter.filter_map(|s| s.parse().ok()).collect();
    let (a, b) = (&vals[..n], &vals[n..2*n]);

    let t0 = Instant::now();
    let mut c = vec![0i32; n];
    for i in 0..dim { for j in 0..dim { for k in 0..dim {
        c[i*dim+j] += a[i*dim+k] * b[k*dim+j];
    }}}
    let ms = t0.elapsed().as_secs_f64()*1000.0;

    println!("Rust {:.4}ms", ms);
}
...

```

Compute in Go

```

```go <matrices >res_go
package main
import ("fmt"; "os"; "strconv"; "strings"; "time")
func main() {
 data, _ := os.ReadFile(os.Getenv("MSH_VAR_matrices"))
 fields := strings.Fields(strings.TrimSpace(string(data)))
 dim, _ := strconv.Atoi(fields[0])
 n := dim*dim
 v := make([]int, 2*n)
 for i, s := range fields[1:] { if i >= 2*n { break }; v[i], _ = strconv.Atoi(s) }
}

```

```

t0 := time.Now()
c := make([]int, n)
for i:=0;i<dim;i++ { for j:=0;j<dim;j++ { for k:=0;k<dim;k++ {
 c[i*dim+j] += v[i*dim+k] * v[n+k*dim+j]
 }}}
ms := float64(time.Since(t0).Nanoseconds())/1e6

fmt.Printf("Go %.4fms\n", ms)
}
...

```

## Compute in Python

```

```python <matrices >res_python
import os, time
with open(os.environ["MSH_VAR_matrices"]) as f:
    data = f.read().split()
dim = int(data[0])
n = dim*dim
v = list(map(int, data[1:]))
A = [v[i*dim:(i+1)*dim] for i in range(dim)]
B = [v[n+i*dim:n+(i+1)*dim] for i in range(dim)]

t0 = time.perf_counter()
C = [[sum(A[i][k]*B[k][j]) for k in range(dim)] for j in range(dim)] for i in range(dim)]
ms = (time.perf_counter()-t0)*1000

print("Python %.4fms" % ms)
...

```

Compute in Lua

```

```lua <matrices >res_lua
local f = io.open(os.getenv("MSH_VAR_matrices"), "r")
local data = f:read("*a")
f:close()
local nums = {}
for nn in data:gmatch("%S+") do table.insert(nums, nn) end
local dim = tonumber(nums[1])
local n = dim*dim
local v = {}
for i=2,#nums do table.insert(v, tonumber(nums[i])) end

```

```

local t0 = os.clock()
local C = {}
for i=1,n do C[i]=0 end
for i=0,dim-1 do
 for j=0,dim-1 do
 local s=0
 for k=0,dim-1 do
 s = s + v[i*dim+k+1] * v[n+k*dim+j+1]
 end
 C[i*dim+j+1] = s
 end
end
local ms = (os.clock()-t0)*1000

print(string.format("Lua %.4fms", ms))
...

```

```
Compute in Bash
```

```

``bash <matrices >res_bash
read -ra DATA < "$MSH_VAR_matrices"
DIM=${DATA[0]}
N=$((DIM*DIM))
V=("${DATA[@]:1}")
A=("${V[@]:0:$N}")
B=("${V[@]:$N:$N}")

T0=$(date +%s%N)
C=()
for i in $(seq 0 $((DIM-1))); do
 for j in $(seq 0 $((DIM-1))); do
 s=0
 for k in $(seq 0 $((DIM-1))); do
 s=$((s + A[i*DIM+k] * B[k*DIM+j]))
 done
 C+=($s)
 done
done
T1=$(date +%s%N)
MS=$(echo "scale=4; ($T1-$T0)/1000000" | bc)

echo "Bash ${MS}ms"

```

```
...
```

```
Display Results
```

```
``lua <matrices <res_c <res_cpp <res_rust <res_go <res_python <res_lua <res_bash
<cycle_start
local function read(env)
 local f = io.open(os.getenv(env), "r")
 if not f then return "" end
 local s = f:read("*all"):gsub("%s+$", "")
 f:close()
 return s
end

local cycle_start = read("MSH_VAR_cycle_start")
local cycle_end_ns = tonumber(io.popen("date +%s%N"):read("*a"):match("(%d+"))
local cycle_start_ns = tonumber(cycle_start:match("(%d+"))
local cycle_ms = 0
if cycle_end_ns and cycle_start_ns then
 cycle_ms = (cycle_end_ns - cycle_start_ns) / 1e6
end

local mat = read("MSH_VAR_matrices")
local matf = {}
for nn in mat:gmatch("%S+") do table.insert(matf, nn) end
local dim = tonumber(matf[1])

local raw = {
 {"C", read("MSH_VAR_res_c")},
 {"C++", read("MSH_VAR_res_cpp")},
 {"Rust", read("MSH_VAR_res_rust")},
 {"Go", read("MSH_VAR_res_go")},
 {"Python", read("MSH_VAR_res_python")},
 {"Lua", read("MSH_VAR_res_lua")},
 {"Bash", read("MSH_VAR_res_bash")},
}

local sep2 = string.rep("-", 46)

print(string.format("N=%d", dim))
print(" " .. sep2)
print(string.format(" %-8s %-14s", "Language", "Time"))
print(" " .. sep2)
```

```
for _,r in ipairs(raw) do
 local line = r[2]:match("[^\n]+") or ""
 local timestr = line:match("%S+%s+(.+)") or line
 print(string.format(" %-8s %s", r[1], timestr))
end
print(" " .. sep2)
print(string.format(" Cycle: %.2fms", cycle_ms))
print("")
...

Ask user

```c
#include <stdio.h>
#include <stdlib.h>
int main() {
    char *status_path = getenv("MSH_VAR_status");
    char *dim_path = getenv("MSH_VAR_dim");

    printf(" [Enter] repeat [n] new dimension [q] quit: ");
    fflush(stdout);

    char buf[32] = {0};
    int i = 0;
    int c;
    while ((c = getchar()) != '\n' && c != EOF && i < 30)
        buf[i++] = (char)c;
    buf[i] = '\0';

    if (buf[0] == 'q' || buf[0] == 'Q') {
        FILE *ws = fopen(status_path, "w");
        if (ws) { fprintf(ws, "quit"); fclose(ws); }
    } else if (buf[0] == 'n' || buf[0] == 'N') {
        printf(" Enter new dimension (2-64): ");
        fflush(stdout);
        char dbuf[16] = {0};
        int j = 0;
        while ((c = getchar()) != '\n' && c != EOF && j < 14)
            dbuf[j++] = (char)c;
        dbuf[j] = '\0';
        int newdim = atoi(dbuf);
        if (newdim < 2) newdim = 2;
        if (newdim > 64) newdim = 64;
    }
}
```

```

FILE *wd = fopen(dim_path, "w");
if (wd) { fprintf(wd, "%d", newdim); fclose(wd); }
FILE *ws = fopen(status_path, "w");
if (ws) { fprintf(ws, "continue"); fclose(ws); }
} else {
FILE *ws = fopen(status_path, "w");
if (ws) { fprintf(ws, "continue"); fclose(ws); }
}
return 0;
}
```

```

```
<!--@end_while-->
```

```
Goodbye
```

```

```bash
echo ""
echo "+=====+"
echo "| Matrix Multiplication Benchmark complete.      |"
echo "| mshell Workflow — Art2Dec SoftLab              |"
echo "+=====+"
echo ""
```

```

## Appendix III: References

Anatomy of High-Performance Matrix Multiplication. Kazushige Goto, R. V. D. Geijn, 2008:  
[https://www.cs.utexas.edu/~flame/pubs/GotoTOMS\\_final.pdf](https://www.cs.utexas.edu/~flame/pubs/GotoTOMS_final.pdf)

HPL Linpack Benchmark: <https://top500.org/project/linpack/>

OpenBLAS Optimized BLAS library: <https://github.com/OpenMathLib/OpenBLAS>

Go compiler documentation: <https://pkg.go.dev/cmd/compile>

Resources: - Common examples Part I (P1–P12):  
<https://www.appservgrid.com/paw92/index.php/2026/02/26/mshell-workflow-patterns-reference-guide-part-i-p1-p13/>

Resources: - Common examples Part II (P13–P24):  
<https://www.appservgrid.com/paw92/index.php/2026/03/11/mshell-workflow-patterns-reference-guide-part-ii-p13-p24/>

- mshell v1.4.1 cheatsheet: <https://www.appservgrid.com/paw92/index.php/2026/02/04/mshell-v-1-4-1-cheatsheet-january-26th-2026/>

Unified language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/24/unified-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Python language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/18/pure-python-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure Bash language Patterns for mshell Workflow — Complete Reference Guide (P1–P24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/17/pure-bash-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure Lua language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/24/pure-lua-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Go language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/23/pure-go-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Rust language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/22/pure-rust-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure C++ language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/21/pure-c-language-patterns-for-mshell-workflow-completereference-guide-p1-p24-2/>

Pure C language Patterns for mshell Workflow — Complete Reference Guide (p1–p24) permanent

link: <https://www.appservgrid.com/paw92/index.php/2026/03/19/pure-c-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure mshell language Patterns for mshell Workflow — Complete Reference Guide (P1–P24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/17/pure-mshell-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>