



HAL
open science

Performance of the extended precision VRP processor on various Krylov subspace solvers

Yves Durand, Alexandre Hoffmann, Jérôme Fereyre

► **To cite this version:**

Yves Durand, Alexandre Hoffmann, Jérôme Fereyre. Performance of the extended precision VRP processor on various Krylov subspace solvers. LACIRM24 - Numerical Linear Algebra, Sep 2024, Marseille, France. , 2024. cea-04799487

HAL Id: cea-04799487

<https://cea.hal.science/cea-04799487v1>

Submitted on 23 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Performance of the extended precision VRP processor on various Krylov subspace solvers

Yves Durand¹, Alexandre Hoffmann², Jérôme Fereyre³

¹Univ. Grenoble Alpes, CEA, List,F-38000 Grenoble, France
email: yves.durand@cea.fr

²Univ. Grenoble Alpes, CEA, List,F-38000 Grenoble, France
email: alexandre.hoffmann@cea.fr

³Univ. Grenoble Alpes, CEA, List,F-38000 Grenoble, France
email: jerome.fereyre@cea.fr

Abstract

Numerical stability is a central aspect of high performance scientific computation, since the growing scale of modern problems has led researchers to complex numerical techniques which are vulnerable to round-off and quantization errors. This is particularly true for iterative Krylov subspace projective solvers, which are the workhorse the modern scientific applications. Here, higher precision, e.g. larger significand size, can mitigate these numerical instabilities and allow for simpler and more memory efficient computations.

The VRP Processor [1] is designed to accelerate extended precision arithmetics in hardware. It supports 1/ a fast arithmetic unit for up to 512 bits of mantissa and 2/ support in memory for unaligned floating-point (FP) arrays.

The objective of this study is to assess the performance impact of using the VRP with extended precision on common solver algorithms, e.g. conjugate gradient (CG), its preconditioned variant (PCG) and biconjugate gradient (BiCG).

We consider two metrics: 1/ convergence speed, which refers to the number of iterations necessary for reaching that objective and 2/ execution time, including memory access time, measured in clock cycles. Execution time depends from the actual implementation of both hardware and low-level software, and from the input matrix structure and values.

Our sample matrices principally come from the Florida sparse Matrix Collection [2]. We restrict ourselves to real matrices, which may be symmetric (for CG and PCG) or asymmetric (for BiCG). We compare execution time between different precisions (including standard double format) on the same VRP platform

Our results confirm the benefits of extended precision for Krylov subspace solvers. For the CG solver, extending precision around 128 mantissa bits appears optimal in terms of iteration count and mostly beneficial for cycle count. The benefit for the BiCG solver is even greater. When using precisions above 256 bits, BiCG convergence becomes predictable.

References

- [1] E. Guthmuller et al., *Xvppfloat: RISC-V ISA Extension for Variable Extended Precision Floating Point Computation*, Journal **IEEE Transactions on Computers**, doi: **10.1109/TC.2024.3383964**. (2024)
- [2] T. Davis et Y. Hu, *The University of Florida Sparse Matrix Collection*, Journal **ACM Transactions on Mathematical Software**, (2011)