

Out-of-Step Pipeline for Gather/Scatter Instructions

Yi Ge*, Katsuhiro Yoda*, Makiko Ito*, Toshiyuki Ichiba*, Takahide Yoshikawa*, Ryota Shioya[†] and Masahiro Goshima[‡]

* Fujitsu Ltd.
Tokyo, Japan

Email: katsu-t@fujitsu.com

[†] The University of Tokyo
Tokyo, Japan

Email: shioya@ci.i.u-tokyo.ac.jp

[‡] National Institute of Informatics
Tokyo, Japan

Email: goshima@nii.ac.jp

Abstract—Wider SIMD units suffer from low scalability of *gather/scatter* instructions that appear in sparse matrix calculations. We address this problem with an *out-of-step* pipeline which tolerates bank conflicts of a multibank L1D by allowing element operations of SIMD instructions to proceed out of step with each other. We evaluated it with a sparse matrix-vector product kernel for matrices from *HPCG* and *SuiteSparse Matrix Collection*. The results show that, for the SIMD width of 1024 bit, it achieves 1.91 times improvement over a model of a conventional pipeline.

I. INTRODUCTION

Demands for sparse matrix calculation have been increasing. Examples are given as follows: NVIDIA implemented the Sparse Tensor Core for sparse matrix in AI workloads. Graph processing also requires large sparse matrices for large graphs. The work that won the 2021 ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research simulates aerosolized droplet with a linear system in a large sparse matrix. The *HPCG* benchmark was developed to reflect the characteristics of these programs [1]. *HPCG* solves a linear system in a large sparse matrix by the CG method. *HPCG* is now one of the three categories of TOP500 along with HPL (HP Linpack).

A wider SIMD unit contributes greatly to dense but little to sparse matrix calculations, which is typically seen in the TOP500 results of the Supercomputer Fugaku. Though it won the first ranks in HPL/*HPCG* in 2021, the efficiencies were as high as 82.3% for HPL but as low as 3.0% for *HPCG*.

We found that the cause of this low efficiency is a scalability problem of SIMD *gather* instructions that inevitably appear in sparse matrix calculation. We address this problem with a *many-bank level-1 data cache (MBL1D)*, and a newly proposed *out-of-step* pipeline which does not reduce the bank conflicts on MBL1D but tolerates them by allowing element operations of SIMD instructions to proceed out of step with each other.

II. OUT-OF-STEP BACKEND PIPELINE

Fig. 1 shows conceptual block diagrams of conventional *in-step* (*InS*) and our *out-of-step* (*OoS*) backend pipelines. Each of the pipelines has two SIMD pipes of two lanes, and the left pipe executes *gather/scatter* instructions to the MBL1D.

a) In-Step Pipeline: The *InS* pipeline can be considered as a single pipeline. Between adjacent stages is a single-entry pipeline register that spans all the lanes. Instructions issued to the lanes flow through the stages in step with each other.

At the cycle $t = t_0$, in the *InS* pipeline are SIMD instructions A , B , and C , each of which has two element operations, such as A_0 and A_1 . The instruction C depends on the *gather* instruction A , and receives the gathered results through the bypass. A_0 and

A_1 access the same bank, resulting in a bank conflict. In this situation, no SIMD instructions can proceed while keeping the *InS* principle, for the following reasons:

A_1 cannot, because it loses in the bank conflict.

B_1 cannot, or it will overwrite A_1 .

C_1 cannot, or it will miss the result of A_1 through the bypass.

This situation is equivalent to an L1D miss in the sense that a load target cannot be retrieved at the timing that the scheduler has assumed. Out-of-order cores generally re-schedule instructions on an L1D miss, because pipeline stall is difficult to implement due to heavy load of the write enable signals to deassert [2]. However, rescheduling on a bank conflict diminishes the gain of MBL1D, because the bank conflict rate is extremely higher than the L1D miss rate.

b) Out-of-Step Pipeline: In the *OoS* pipeline on the right side of Fig. 1, the single-entry pipeline register immediately before the execution stages is replaced with small FIFO buffers for the lanes. In addition, the downstream pipeline registers are divided for the lanes. As a result, the downstream part of the lanes from the buffers can independently proceed.

In the *OoS* pipeline in the figure, the positions of the element operations are at the cycle $t = t_0 + 1$. These element operations can proceed out of step with each other, as follows:

A A_1 lost in the bank conflict has remained in the buffer, and in turn accesses the bank.

B B_1 has been stored in the second entry of the buffer without overwriting A_1 . B_0 accesses another bank in parallel with A_1 .

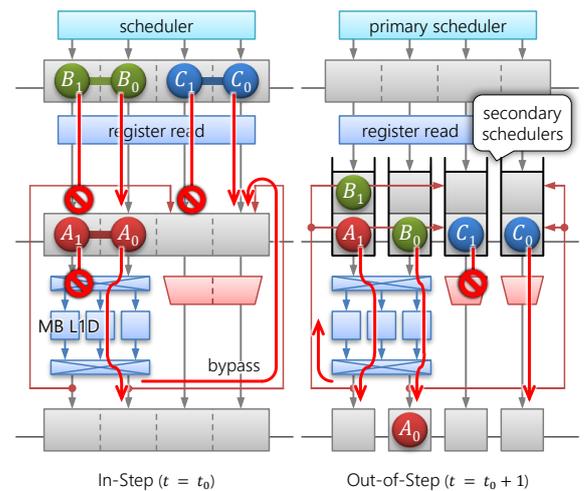


Fig. 1. In-/Out-of-Step backend pipelines.

