Proceeding Series of the Brazilian Society of Applied and Computational Mathematics, Vol. 3, N. 1, 2015.
Trabalho apresentado no XXXV CNMAC, Natal-RN, 2014.

Implementing a Parallel Sparse Matrix-Vector Multiplication Using Dataflow

Brunno F. Goldstein
PESC – COPPE – UFRJ
bfgoldstein@cos.ufrj.br

Tiago A. O. Alves
PESC – COPPE – UFRJ
tiagoaoa@cos.ufrj.br

Michael Souza
DEMA – UFC
michael@ufc.br

Luiz Mariano Carvalho
DMA – IME – UERJ
luizmc@ime.uerj.br

Leandro A. J. Marzulo
DICComp – IME – UERJ
leandro@ime.uerj.br

Felipe M. G. França
PESC – COPPE – UFRJ
felipe@cos.ufrj.br

ABSTRACT

Sparse matrix-vector multiplication (SpMV) is an important operation in sparse linear algebra problems. According to Bell [1], *"they represent the dominant cost in many iterative methods for solving large-scale linear systems"*. In this paper we propose a dataflow parallel version of the SpMV kernel using Compressed Sparse Row (CSR) storage format. This implementation is executed in a Dataflow Runtime Environment called Trebuchet [3][4][5].

The dataflow model[2] exposes parallelism by essence. Task execution is triggered as soon as the input data is available. A dataflow program can be described as a dataflow graph where nodes represent instructions (or tasks) and edges represent their input/output dependencies. Two instructions can run concurrently if there is no directed path between them and their pieces of data are available. This differs from Von Neumann model where instructions are guided by control flow and their execution is inherently sequential.

Our experimental environment consists of a machine with 4 Quad-Core Intel Core i7-3820 processors running at 3.6GHz and 64 GB RAM. We performed the SpMV kernel in three matrices from University of Florida Sparse Matrix Collection[1], listed in Table 1 . Figure 1 presents the results for Dataflow(running on Trebuchet), Intel Math Kernel Library(MKL)[2] (which uses OpenMP[3] in parallel version) implementations.

Results show that with a naive dataflow implementation it is already possible to achieve higher speedups than the ones obtained with Intel MKL. For smalled matrices the dataflow runtime environment overheads become more evident, but for larger ones, Trebuchet provides better results. In order to improve performance, we intend to implement other SpMV kernels using different storage formats and other approaches to get better use of cache locality[6]. Moreover, since dataflow allows more flexibility in describing application dependencies, we may find more interesting parallelization patterns that are easier to implement and provide smaller overheads in dataflow.

Keywords: *Dataflow, Sparse matrix-vector multiplication, Trebuchet*

Acknowlegments

We would like to thank FAPERJ and CAPES for the support given to the authors of this work.

[1] http://www.cise.ufl.edu/research/sparse/matrices/groups.html
[2] http://software.intel.com/en-us/intel-mkl
[3] http://openmp.org/

DOI: 10.5540/03.2015.003.01.0108

© 2015 SBMAC

Group	Matrix Name	Rows	Columns	Nonzeros	Data Type	Structure
Rudnyi	water_tank	60,740	60,740	2,035,281	Real	Unsymmetric
Botonakis	FEM_3D_thermal2	147,900	147,900	3,489,300	Real	Unsymmetric
PARSEC	Ga41As41H72	268,096	268,096	18,488,476	Real	Symmetric
Freescale	circuit5M_dc	3,523,317	3,523,317	14,865,409	Real	Unsymmetric

Table 1: Matrices used for speedup tests.

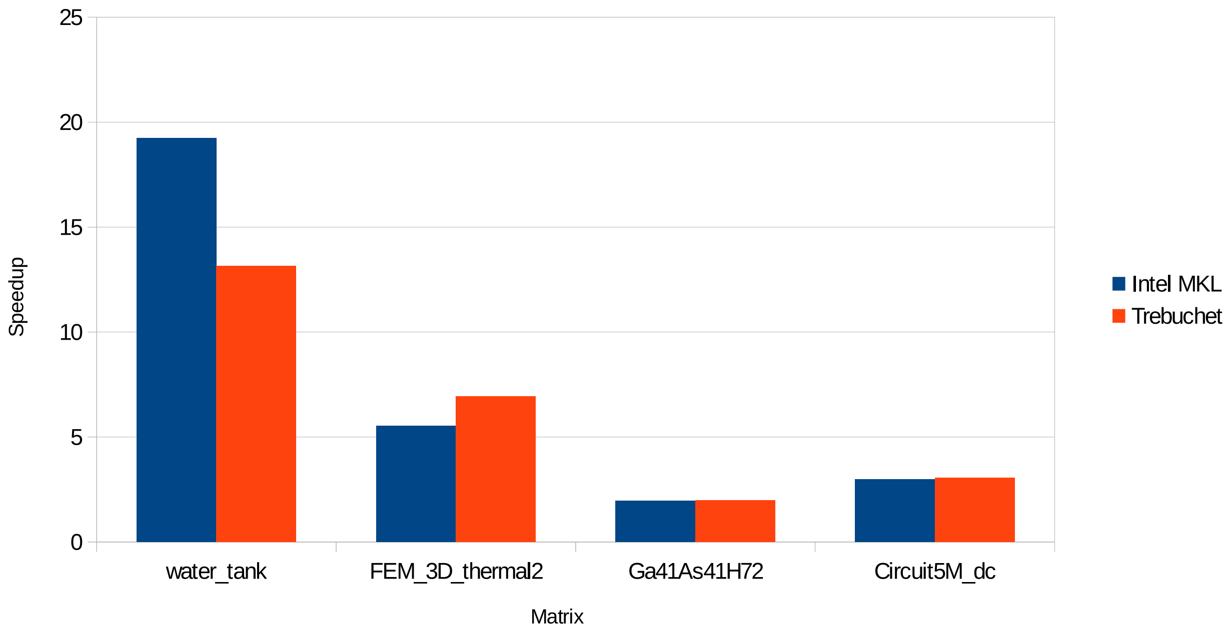

Figure 1: Speedup results for Intel MKL and Dataflow Trenbuchet implementations.

References

[1] Nathan Bell and Michael Garland. “Efficient sparse matrix-vector multiplication on CUDA”. NVIDIA Technical Report NVR-2008-004, NVIDIA Corporation, December 2008.

[2] J. B. Dennis and D. P. Misunas, “A preliminary architecture for a basic data-flow processor,” SIGARCH Comput. Archit. News, vol. 3, no. 4, pp. 126–132, 1974.

[3] L. A. J. Marzulo, T. A. Alves, F. M. G. Franca, and V. S. Costa, “TALM: A Hybrid Execution Model with Distributed Speculation Support,” Computer Architecture and High Performance Computing Workshops, International Symposium on, vol. 0, pp. 31–36, 2010.

[4] T. A. Alves, L. A. Marzulo, F. M. Franca, and V. S. Costa, “Trebuchet: exploring TLP with dataflow virtualisation,” International Journal of High Performance Systems Architecture, vol. 3, no. 2/3, p. 137, 2011.

[5] T. A. Alves, L. A. J. Marzulo, F. M. G. França, and V. S. Costa, “Trebuchet: Explorando TLP com Virtualização DataFlow,” in WSCAD- SSC’09. São Paulo: SBC, Oct. 2009, pp. 60–67.

[6] Aydin Buluc, Samuel Williams, Leonid Oliker, and James Demmel. “Reduced-bandwidth multi-threaded algorithms for sparse matrix-vector multiplication.” In Proceedings of the 2011 IEEE International Parallel & Distributed Processing Symposium, IPDPS ’11, pages 721–733, Washington, DC,USA, 2011. IEEE Computer Society.

DOI: 10.5540/03.2015.003.01.0108 © 2015 SBMAC