Delta-Stepping Synchronous Parallel Model

Weidong Zhang*, Chang Cui†, Yifeng Chen*, Qifei Zhang§

School of Electronics Engineering and Computer Science, Peking University* Yuanpei College, Peking University† School of Software Technology, Zhejiang University§

Abstract

Many synchronous parallel algorithms like PageRank require a large number of iteration steps. The overheads of global synchronizations on general-purpose cluster take substantial proportion of the execution time for lightweight computations. We propose a variant of Bulk Synchronous Parallel (BSP), Delta-Stepping Synchronous Parallel (DSP), with fewer iteration steps. It achieves faster convergence process by exploring full advantage of data locality.

Introduction

Many step-wise parallel algorithms can be intuitively expressed as BSP pattern [2]. This type of parallel pattern is described as in Figure 1.

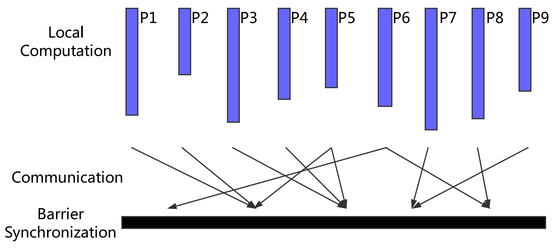

Figure 1: The BSP pattern: One round of local computation is conducted in each superstep.

As is shown in Figure 1, varying degrees of interdependence exist among these processors $\{P_1, P_2, \ldots, P_n\}$. When the dependence between P_i and P_j $(j \neq i)$ is subtle, the orientation of the convergence of P_i will be mainly decided by the data residing in itself. So we conjecture that increasing local computing steps in each superstep will speed up local convergence, sequentially advance the global convergence. The idea is sketched in Figure 2.

By further formalization and derivation, we prove that, if the algorithm converges with two local computations, then it converges with any number of local computations. For the convex optimization problems and local-optimal insensitive problems, the convergence is sufficient.

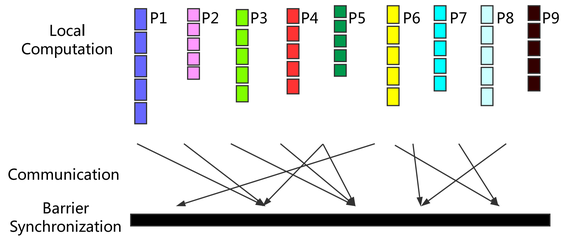

Figure 2: The DSP pattern: Δ rounds of local computation is conducted in each superstep.

Case Study

To demonstrate the applicability and performance, we apply the model on several algorithms: Max Value Propagation (MVP), Jacobi Iterative Method (JIM), Single Source Shortest Path (SSSP) and PageRank (PR).

As is shown in Figure 3, the figures show that DSP reduces the numbers of iterations and communication of MVP, JIM, SSSP and PR significantly. Figure 4 show that DSP reduces the execution time and the number of iterations of SSSP and PageRank dramatically.

Conclusion

DSP is a variant of BSP. It utilizes inaccurate global data when performs multiple computation steps in each superstep. These advanced computation steps further exploit the locality of data, and accelerate the convergence. This research was supported by the Zhejiang Engineering Research Center of Intelligent Medicine (2016E10011) and the research and application of key technologies for rapid individualized sculpturemanufacture and carving stone materials appraisal (2016C03SAB80611).

References

[1] Leskovec and Krevl and Leskovec. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. http://snap.stanford.edu/data. (June 2014).

[2] L. G. Valiant, A bridging model for parallel computation, Communications of the Acm, vol. 33, no. 8, pp. 103-111, 1990.

(a) Max Value Propagation, G($\|V\| = 40,000, p\_in = 1.25/\|V\|, p\_out = 0.005/\|V\|$)										
	BSP	DSP								
		$\Delta = 2$	$\Delta = 3$	$\Delta = 4$	$\Delta = 5$	$\Delta = 6$	$\Delta = 7$	$\Delta = 8$	$\Delta = 9$	$\Delta = 10$
iteration	68	50	42	42	42	42	42	42	42	42
Communication(KB)	9925	7298	6130	6130	6130	6130	6130	6130	6130	6130
(b) Jacobi Iterative Method. The linear system consists of 10,000 equations.										
	BSP	DSP								
		$\Delta = 2$	$\Delta = 3$	$\Delta = 4$	$\Delta = 5$	$\Delta = 6$	$\Delta = 7$	$\Delta = 8$	$\Delta = 9$	$\Delta = 10$
iteration	438	220	147	111	89	74	64	56	50	45
Communication(KB)	17109	8593	5742	4335	3476	2890	2500	2187	1953	1757
		$\Delta = 50$	$\Delta = 100$	$\Delta = 200$	$\Delta = 300$	$\Delta = 400$	$\Delta = 500$	$\Delta = 800$		
iteration		11	7	4	3	3	2	2		
Communication(Byte)		429	273	156	117	117	78	78		
(c) Single Source Shortest Path, G($\|V\| = 40,000, p\_in = 1.25/\|V\|, p\_out = 0.0025/\|V\|$)										
	BSP	DSP								
		$\Delta = 2$	$\Delta = 3$	$\Delta = 4$	$\Delta = 5$	$\Delta = 6$	$\Delta = 7$	$\Delta = 8$	$\Delta = 9$	$\Delta = 10$
iteration	53	36	35	35	35	35	35	35	35	35
Communication(KB)	7894	5362	5213	5213	5213	5213	5213	5213	5213	5213
(d) PageRank, G($\|V\| = 40,000, p\_in = 5.0/\|V\|, p\_out = 0.01/\|V\|$)										
	BSP	DSP								
		$\Delta = 2$	$\Delta = 3$	$\Delta = 4$	$\Delta = 5$	$\Delta = 6$	$\Delta = 7$	$\Delta = 8$	$\Delta = 9$	$\Delta = 10$
iteration	51	40	30	25	23	22	22	22	22	22
Communication(KB)	13388	10500	7875	6562	6037	5777	5777	5777	5777	5777

Figure 3: Performance comparison between DSP and BSP. The graphs used in (a, c, d) are random graphs, p_in, p_out indicate the possibilities of a edge existed between a pair of vertices in the same partition and different partitions respectively.

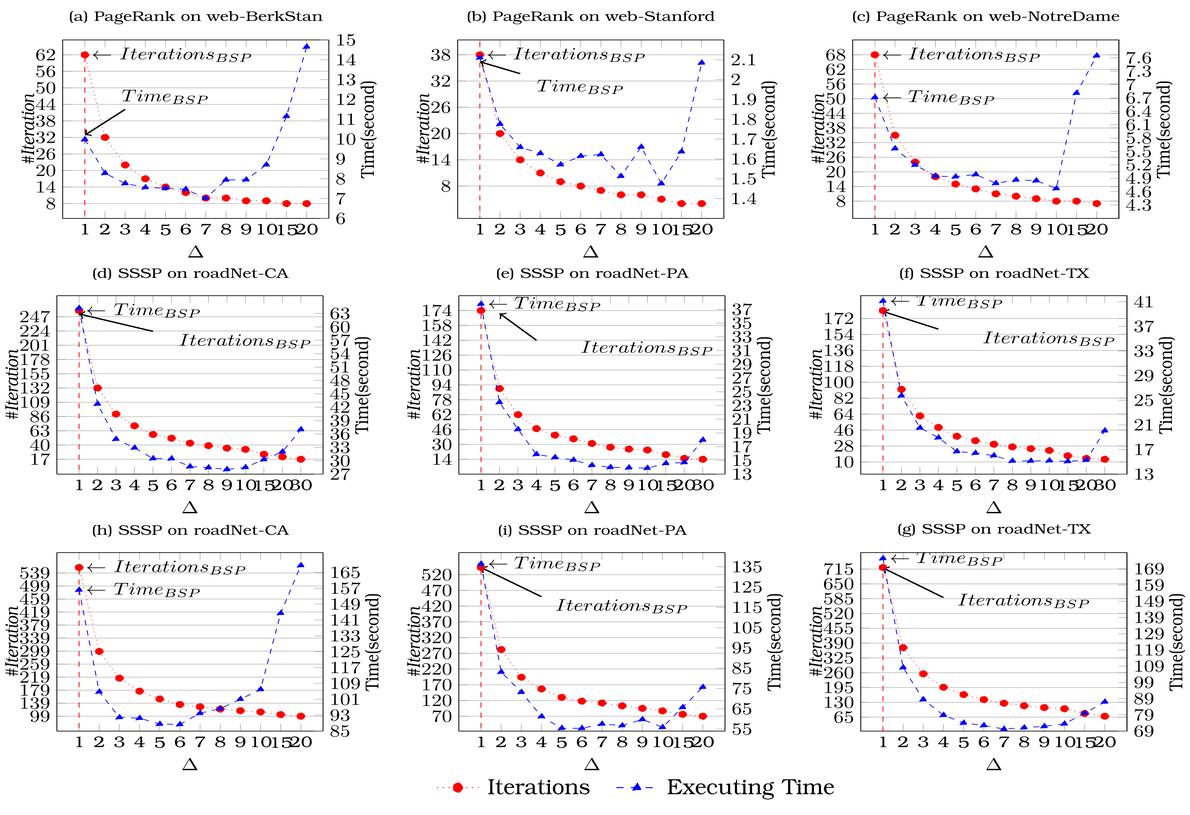

Figure 4: Performance comparison between DSP and BSP. (a-c) show the results of PageRank working on well-partitioned subgraphs, (d-f) and (h-g) show the results of SSSP working on well-partitioned and random-partitioned subgraphs respectively. The convergence accuracy of PageRank is set to 10^{-10}. The real web graphs and road networks [1] are used in this experiment.