BSF: a parallel computation model for scalability estimation of iterative numerical algorithms on cluster computing systems

Leonid B. Sokolinsky

South Ural State University (National Research University) 76, Lenin prospekt, Chelyabinsk, Russia, 454080 E-mail: leonid.sokolinsky@susu.ru Phone: +79048040330

Abstract

This paper examines a novel parallel computation model called bulk synchronous farm (BSF) that focuses on estimating the scalability of compute-intensive iterative algorithms aimed at cluster computing systems. The main advantage of the proposed model is that it allows to estimate the scalability of a parallel algorithm before its implementation. Another important feature of the BSF model is the representation of problem data in the form of lists that greatly simplifies the logic of building applications. In the BSF model, a computer is a set of processor nodes connected by a network and organized according to the master/slave paradigm. A cost metric of the BSF model is presented. This cost metric requires the algorithm to be represented in the form of operations on lists. This allows us to derive an equation that predicts the scalability boundary of a parallel program: the maximum number of processor nodes after which the speedup begins to decrease. The paper includes examples of applying the BSF model to designing and analyzing parallel numerical algorithms. The large-scale computational experiments conducted on a cluster computing system confirm the adequacy of the analytical estimations obtained using the BSF model. *Keywords:* parallel computation model, cluster computing systems, iterative numerical algorithms, performance evaluation, parallel speedup, scalability boundary, BSF model, bulk synchronous farm

Preprint submitted to Journal of LATEX Templates

January 2, 2021

1. Introduction

Currently, we are entering the era of exascale computers operating at a speed of 10 to the 18th power (10^{18}) flops [1]. The recent TOP500 list (November 2020) [2] shows that more than 90% of the most powerful supercomputers have the cluster architecture. The design of numerical algorithms for such computing clusters requires new approaches to achieve the high efficiency of parallelization. It is important to evaluate the scalability of a parallel algorithm at an early stage of its development. *Scalability* has been widely used in practice to describe how system sizes and problem sizes influence the performance of parallel computers and algorithms [3]. The main measure for evaluating the scalability of a parallel algorithm on a cluster computing system is the *speedup a*, which is defined as the ratio of the algorithm execution time T_1 on one processor node to the algorithm execution time T_K on K processor nodes:

$$a(K) = \frac{T_1}{T_K}.$$
(1)

It is well known that for a given computing cluster architecture and a fixed-size problem, the speedup of a parallel algorithm does not continue to increase with an increase in the number of processor nodes, but it tends toward saturation and culminates in a peak at a certain system size after which the speedup begins to decrease. Let us define the *scalability boundary* of the parallel algorithm as the number of processor nodes K_{max} at which the speedup peak is reached for the given problem size on the target cluster computing system. To detect the scalability boundary of a parallel algorithm, we have the following two possibilities. First, we can conduct a series of large-scale computational experiments on the target cluster system to plot the speedup curve and visually determine the scalability boundary. However, it takes time and effort to build a compilable and executable implementation of the parallel algorithm in some programming language. Moreover, we need to obtain access to a sufficiently large cluster computing system for a sufficiently long time. The second possibility is to use a suitable parallel computation model that can predict the execution time of the algorithm for the target cluster computing system.

The *computational model* is a simplified and abstract description of a computer. A computer architect, algorithm designer and program developer can use such a model as a basis to assess their work, including the suitability of one computer architecture to various applications, the computation complexity of an algorithm and the potential performance of one program on various computers, etc. A good computational model can simplify the complicated work of the architect, algorithm designer and program developer while mapping their work effectively onto real computers [4]. Thus, such a computational model is sometimes also called a "bridging model" [5]. An universal bridging model can be applied to any algorithms and any computers (see Fig. 1 a). The RAM (random access machine) model [6, 7, 8] was such a universal model bridging the sequential computers and algorithms. With the advent of parallel computers, numerous attempts were made to build a similar universal model bridging the multiprocessors and parallel algorithms [9], but these attempts failed. This is mainly due to the large variety of multiprocessor architectures that are rapidly emerging and developing in response to the demands of increasing computer performance. Under these conditions, creating a simple and accurate universal model of parallel computations is almost impossible. The approach schematically shown in Fig. 1 b was applied to overcome these difficulties [10]. According to this approach, the parallel architectures were divided into three classes: shared memory, distributed memory, and hierarchical memory multiprocessors [4].

Separate parallel computation models were created for each class of multiprocessors, but almost all these models were universal with respect to a variety of parallel numerical algorithms. This approach generated simple and reliable models with a high level of abstraction, such as PRAM [11], BSP [5], and LogP [12]. Numerous attempts have been made to refine and extend these models to adapt them to the increasing complexity of multiprocessor system architectures. This has led to the emergence of more accurate but complicated to apply models of parallel computations (see, for example, [13, 14, 15, 16, 17, 18]). Dividing the entire set of algorithms into different types allows us to correct



Figure 1: Models bridging algorithms and computers.

this situation to a certain extent (see Fig. 1 c). Examples of different types of algorithms can be iterative numerical algorithms, graph algorithms, big data processing algorithms, and so on. Each pair (algorithm type, architecture class) can have its own parallel computation model. Such an approach allows us to reach an acceptable tradeoff between the accuracy of estimations and usability.

G. Bilardi and A. Pietracaprina distinguish the following four typical components of a computational model [19]: an *architectural component*, described as an interconnection of modules of various functionalities; a *specification component*, determining what is a (syntactically) valid algorithm/program; an *execution component*, defining which sequences of states of the architectural modules constitute valid executions of a program/algorithm on a given input; and a *cost*

component, defining one or more cost metrics for each execution. They also formulate the following three conflicting requirements by which a computational model can be evaluated: the ease of algorithm/program design and analysis (us*ability*); the ability of estimating, from the cost metrics provided by the model, the actual performance of a program on a specific real platform (*effectiveness*) and on a class of hardware platforms (hardware portability). In accordance with the approach shown in Fig. 1 c, we add one more component: algorithmic *portability*, determining the type of algorithms to which the model is applicable. In the context of this approach, this article presents a novel model of parallel computations, named bulk synchronous farm (BSF) that is an extension of the BSP model proposed initially by Valiant in [5]. The BSF model is based on the original idea outlined in [20] and intended for evaluating the parallel iterative compute-intensive numerical algorithms on the cluster computing systems. Let us clarify the meaning of the restrictions imposed on the scope of the BSF model's applicability. The "iterative algorithm" means that the time spent on the preparation to perform the iterative computations (data input, memory allocation, variable initialization, and so on) is so markedly less than the time spent on the iterative computations themselves that we can neglect the initialization overhead within the model. In addition, the framework of the "iterative algorithm" includes a sequence of steps implemented as a loop in which the next iteration depends on the results of the previous ones and cannot be executed in parallel with them. The "compute-intensive numerical algorithm" means that the time spent for calculations is greater or comparable to the time spent for input/output and communications between processor nodes. The "cluster com*puting system*" is a set of tightly connected homogeneous processor nodes with private memory that communicate with each other through the MPI library. The model treats the processor node as a black box that can perform scalar and vector operations at a certain speed. The restrictions mentioned above allow us to obtain a simple and reliable equation for estimating the scalability boundary of an iterative numerical algorithm by using the BSF model. No other known model provides such an equation.

The rest of the article is organized as follows. Section 2 briefly reviews parallel computation models for distributed memory multiprocessors. Section 3 introduces the BSF model description determining its architectural, specification and execution components. Section 4 presents the cost metric of the BSF model and contains the derivation of the main equation that estimates the scalability boundary of a parallel algorithm. Section 5 demonstrates how to apply the BSF model to estimate the algorithm scalability boundary using the iterative Jacobi method as an example. Section 6 includes the results of large-scale computational experiments and their comparison with the results obtained analytically using the BSF model. In Section 7, we discuss the strengths and weaknesses of the BSF model. Section 8 concludes this paper.

2. Related works

One of the first parallel computation models for distributed memory multiprocessors was BSP (bulk-synchronous parallel) model proposed by Valiant in [5]. A BSP-computer is a system of K processors that have private memory and are connected by a network allowing data to be transferred from one processor to another. The following cost parameters of the interconnect are defined: g — the time required to transfer a single machine word across the network; L — the time required for a global synchronization. In the BSP-computer, the message transfer is simulated using the notion of h-session. The h-session is an abstraction of an arbitrary communication operation in which each processor transfers no more than h machine words and receives no more than h machine words. In the BSP-computer, the execution time of a single h-session cannot exceed hg. The BSP-program consists of n parallel processes each of that is assigned to a separate processor. The BSP-program is divided into global sequential supersteps that are synchronously executed by all processes. Each superstep includes the following four sequential steps: 1) computations on each processor using only local data; 2) global barrier synchronization; 3) data transfer from any processor to any other processors by performing a single *h*-session; 4) global barrier synchronization. In the BSP model, the *cost metric* is constructed as follows. Let the BSP-program consist of S supersteps. Assume that each processor performs no more than w_i clock cycles during local calculations in the *i*-th superstep. The total time t_i taken by the system to execute the *i*-th superstep is calculated using the equation

$$t_i = w_i + hg + L.$$

The total time T of executing the entire program is determined by the equation

$$T = W + hgS + LS,$$

where $W = \sum_{i=1}^{S} w_i$. The BSP model has a number of disadvantages. The first, the BSP model admits to transfer messages with a length of only one machine word and does not take into account that transferring h machine words as a single message can be more efficient than transferring h messages with the length of one machine word. The second, the amount of data processed by a single processor in each superstep should be approximately equal to the number of words received during the h-session, i.e. h that limits the granularity of parallelism from below. The third, the BSP model assumes hardware support for the global synchronization, but most multiprocessor systems with distributed memory do not have such mechanism.

To overcome these disadvantages, Culler and co-authors proposed a parallel computation model named LogP [12] that extends the BSP model. Like the BSP model, the *LogP-computer* is a system of P processors that have private memory and are connected by a network that allows data to be transferred from one processor to another. The LogP model has the following cost parameters:

- L : *latency* (time of transferring one machine word from one processor to another);
- o : overhead (length of time that a processor is engaged in the transfer or reception of each message);
- g : gap (minimum time interval between consecutive message transfers or

consecutive message receptions at a processor);

P : number of processors.

The LogP model assumes that the network has a finite capacity, such that at most $\lfloor L/g \rfloor$ messages can be in transfer from any processor or to any processor at any time. Furthermore, the LogP model assumes that all messages have a small size (one or a few number of machine words). Large messages need to be fragmented. The time taken by computations using local data in the superstep is calculated in the same way as in the BSP model. The time of transferring one short message from one processor to another is o + L + o. The time of transferring n consecutive short messages is (n - 1)g + o + L + o. An obvious disadvantage of the LogP model is the message size limitation.

The LogP model has been improved in numerous extensions. In [37], an extension of the LogP model called LogGP was proposed. The LogGP model adds the new parameter G (gap per byte) that determines the time of transferring one byte within a long message. The time of transferring a message of length m is o + (m - 1)G + L + o. The LogGPS model [38] extends the LogGP model by introducing an additional parameter S that takes into account the synchronization overhead. Actually, the LogGPS model is an adaptation of the LogPQ model to the features of the communication protocols of the MPICH library [39] that is a portable implementation of MPI. Like the LogGPS model, the $log_n P \mod l$ [40] is dedicated to cluster computing systems that use MPI for message transfers. A distinctive feature of the $log_n P$ model is that it takes into account the implicit overhead of transferring data between different levels of hierarchical memory.

The emergence of heterogeneous cluster systems has generated a new class of parallel computation models. One of these models called *Multi-BSP* was proposed by Valiant in [41]. The *Multi-BSP model* extends the BSP model in the following two directions. First, Multi-BSP is a hierarchical model with an arbitrary number of the levels that represent the actual technical features of the hierarchical memory and various cache levels of modern multiprocessor systems. Second, Multi-BSP includes the amount of memory at each level as an additional parameter. For each level *i* in the hierarchy, the following vector of parameters is introduced: (p_i, g_i, L_i, m_i) . Here, p_i is the number of processors, g_i is a delay in transferring data from level *i* to level i + 1, L_i is the synchronization overhead, and m_i is the amount of memory/cache. For a system that includes *d* levels, the total number of processors is calculated using the equation $P_d = \prod_{i=1}^d p_i$, the total amount of memory is calculated using the equation $M_d = m_d + \sum_{i=1}^{d-1} m_i \prod_{j=i+1}^d p_j$, and the total delay is calculated using the equation $G_d = \sum_{i=1}^d g_i$.

Another model of this class is the $mlog_n P$ [42] extending the $log_n P$ model. The $mlog_n P$ model is dedicated to computing clusters with multi-core processors. In this model, the *n*-level memory hierarchy introduced in the $log_n P$ model is called *vertical*. Additionally, the authors introduce a *horizontal m*-level hierarchy of data channels. The zero-level channel serves data exchange between the cores of a single processor, the first-level channel is used for data exchange between cores of different processors of a single processor node, the second-level channel is dedicated for data exchange between the cores of different processor node, the cores of different processor node, the second-level channel is dedicated for data exchange between the cores of different processor nodes, and so on. The cost of transferring a message at *i*-th horizontal level is calculated using the equation

$$T_i = \sum_{j=0}^{n_i-1} (o_j^i + l_j^i) = \sum_{j=0}^{n_i-1} (f_j^i(s, 1) + f_j^i(s, d)),$$

where the semantics of all parameters within a single horizontal level is inherited from the $log_n P$ model. The model is dedicated to predicting the execution time of a CUDA-program on cluster computing systems with Kepler GPUs.

In [35], an extension of the BSP model for GPUs running CUDA is proposed. This model focuses on predicting the execution time of a CUDA-program on cluster computing systems with Kepler GPUs. The approximated execution time T_k of a *kernel function* with t threads is calculated by the following equation

$$T_k = \frac{t \times (Comp + Comm_{DRAM} + Comm_{GPUM})}{R \times P \times \lambda}.$$

It sums the computational cost (Comp) with the communication cost of global

memory $(Comm_{GM})$ and shared memory $(Comm_{SM})$ accesses, performed by each thread. This cost is multiplied by the number of threads t and divided by the clock rate R times the number of cores P available in the GPU. The parameter λ is used to model the effects of application optimizations, such as divergence, shared bank conflicts and coalesced global memory accesses. The global communication cost is estimated in the same way as in the BSP model.

A large number of other sophisticated parallel computation models for modern cluster computing systems have been proposed in recent years (see surveys [4, 10]). The main disadvantage of these models is the complexity of their practical application when designing and analyzing parallel numerical algorithms for exascale computers. None of these models yields a ready-to-use equation for estimating the scalability boundary of a parallel numerical algorithm.

3. Description of BSF model

In this section, we give a description of the BSF model, namely, determine its architectural, specification and execution components. The *architecture* of the BSF-computer is shown in Fig. 2. A BSF-computer consists of a collection of homogeneous processor nodes with private memory connected by a communication network delivering messages among the nodes. All the nodes have the same capacity. The BSF model treats the processor node of a real computing cluster as a black box. Node interaction is based on the master/slave paradigm [21]: one processor node is the *master node*; all other processor nodes are *worker nodes* (also sometimes referred to as "slaves"). The master node serves as the control and communication hub.

In the BSF model, an algorithm must be *specified* in the form of operations on lists by using the *Map* and *Reduce* functions. The higher-order functions *Map* and *Reduce* defined in the Bird-Meertens formalism [22] are the basis for the parallelization of the BSF-algorithms. Let $[a_1, \ldots, a_l]$ denote a list of length lthat includes elements of a given set \mathbb{A} . Let $F : \mathbb{A} \to \mathbb{B}$ be a function that



Figure 2: Architecture of BSF-computer.

Algorithm 1. Generic BSF-algorithm template.

1: input $A, x^{(0)}$ 2: i := 03: $B := Map(F_{x^{(i)}}, A)$ 4: $s := Reduce(\oplus, B)$ 5: $x^{(i+1)} := Compute(x^{(i)}, s)$ 6: i := i + 17: if $StopCond(x^{(i)}, x^{(i+1)})$ goto 9 8: goto 3 9: output $x^{(i)}$ 10: stop

maps the set \mathbb{A} to a given set \mathbb{B} . The higher-order function Map applies the function F to each element of the list $[a_1, \ldots, a_l]$ and returns a list of results in the same order:

$$Map(F, [a_1, ..., a_l]) = [F(a_1), ..., F(a_l)].$$
 (2)

Let $[b_1, \ldots, b_l]$ denote a list of length l that includes elements of the set \mathbb{B} . Let $\oplus : \mathbb{B} \times \mathbb{B} \to \mathbb{B}$ be a binary associative operation on the set \mathbb{B} . The higherorder function *Reduce* reduces the list $[b_1, \ldots, b_l]$ to a single value by iteratively applying the operation \oplus to its elements:

$$Reduce(\oplus, [b_1, \dots, b_l]) = b_1 \oplus \dots \oplus b_l.$$
(3)

The generic template of an iterative BSF-algorithm is presented as Algorithm 1. The variable *i* denotes the iteration number; $x^{(0)}$ is an initial approximation; $x^{(i)}$ is the *i*th approximation (the approximation can be a number, a vector, or any other data structure); A is the list of elements of a certain set $\mathbb{A},$ which represents the source data of the problem; $F_x:\mathbb{A}\to\mathbb{B}$ is a parameterized function (the parameter x is the current approximation) that maps the set A to a certain set \mathbb{B} ; B is a list of elements of the set \mathbb{B} calculated by applying the function F_x to each element of the list A; \oplus is a binary associative operation on the set \mathbb{B} . Step 1 reads the input data of the problem and the initial approximation. Step 2 assigns the zero value to the iteration counter i. Step 3 calculates the list B by invocating the higher-order function $Map(F_{x^{(i)}}, A)$. Step 4 assigns the result of the higher-order function $Redice(\oplus, B)$ to the intermediate variable s. Step 5 invocates the user function Compute that calculates the next approximation $x^{(i+1)}$ taking two parameters: the current approximation $x^{(i)}$ and the result s of the higher-order function *Reduce*. Step 6 increases the iteration counter i by one. Step 7 checks termination criteria by invocation of the user Boolean function StopCond, which takes two parameters: the new approximation $x^{(i)}$ and the previous approximation $x^{(i-1)}$. If StopCond returns true, the algorithm outputs $x^{(i)}$ as an approximate problem solution and stops working. Otherwise, the control is passed to Step 3 starting the next iteration.

The *parallel execution* of a BSF-algorithm is based on the following theoretical foundation. Let us divide the list $A = [a_1, \ldots, a_l]$ into K sublists of length m:

$$A = A_1 + \dots + A_K \tag{4}$$

(for simplicity, we assume that l is a multiple of K, i.e., l = Km for some $m \in \mathbb{N}$). Here, # denotes the operation of list concatenation. According to the promotion theorem [23], the following equation holds:

$$Reduce (\oplus, Map (F_x, A)) =$$

$$= Reduce (\oplus, Map (F_x, A_1)) \oplus \cdots \oplus Reduce (\oplus, Map (F_x, A_K)).$$
(5)

Equation (5) gives us the generic parallelization scheme shown in Fig. 3. We can run K parallel *worker* threads that independently perform the higher-order functions *Map* and *Reduce* over sublists A_1, \ldots, A_K . Then, the *master* thread joins the produced partial foldings s_1, \ldots, s_K into the single list S and performs



Figure 3: BSF-algorithm parallelization schema.

higher-order function Reduce over it.

The generic parallelization template of an iterative BSF-algorithm is presented as Algorithm 2. It includes K + 1 parallel processes: one master process and K worker processes. The master process runs on the master node. Each worker process runs on a separate worker node. In Step 1, the master process reads the initial approximation $x^{(0)}$ and assigns the zero value to the iteration counter *i*. At the same time, every *j*th worker process reads the sublist A_j that assigned to it for processing and is treated as local data. In Step 2, the master process sends the current approximation $x^{(i)}$ to all worker processes. After that, every jth worker process independently applies higher-order functions Map and Reduce to its sublist (Steps 3 and 4). In Steps 3 and 4, the master process is idle. In Step 5, every jth worker process sends to the master process the partial folding s_i that is a result of the *Reduce* function. In Steps 6-9, the master process performs the following actions: executes the higher-order function *Reduce* over the list of partial foldings $[s_1, \ldots, s_K]$; invocates the user function Compute that calculates the next approximation; and checks the termination criteria by using the user Boolean function StopCond and assigns its result to the Boolean variable *exit*. In Steps 6-9, the worker processes are idle. In Step 10, the master

Master	jth Worker (j=1,,K)
1: input $x^{(0)}; i := 0$	1: input A_j
2: $SendToAllWorkers(x^{(i)})$	2: $RecvFromMaster(x^{(i)})$
3:	3: $B_j := Map(F_{x^{(i)}}, A_j)$
4:	4: $s_j := Reduce(\oplus, B_j)$
5: $RecvFromWorkers(s_1,\ldots,s_K)$	5: $SendToMaster(s_j)$
6: $s := Reduce(\oplus, [s_1, \dots, s_K])$	6:
7: $x^{(i+1)} := Compute(x^{(i)}, s)$	7:
8: $i := i + 1$	8:
9: $exit := StopCond(x^{(i)}, x^{(i+1)})$	9:
10: SendToAllWorkers(exit)	10: RecvFromMaster(exit)
11: if not <i>exit</i> goto 2	11: if not <i>exit</i> goto 2
12: output $x^{(i)}$	12:
13: stop	13: stop

Algorithm 2. Generic BSF-algorithm parallelization template.

process sends the *exit* value to all worker processes. If the *exit* value is false, the master process and worker processes go to the next iteration; otherwise, the master processes output the result and the computation stops. Note that in Steps 2 and 10, all processes perform the implicit global synchronization. In this template, all worker processes execute the same code for different sublists. Since all sublists have the same length, there is no need to balance the workload of the worker nodes.

4. Cost metric of BSF model

The BSF model assumes that the overhead of initializing and terminating a program is negligible compared to the overhead of executing the iterative process. The cost of an iterative process is the sum of the costs of individual iterations. Therefore, to estimate the execution time of an iterative BSF-algorithm, we simply need to obtain an estimation of the time cost of one iteration. The BSF model includes the following cost parameters for a single iteration:

K : number of worker nodes;

- l : length of the list A representing the input data (the same as the length of the list B representing the result of the higher-order function Map);
- L : latency (time of transferring one-byte message node-to-node);
- t_c : time taken by the master node to send the current approximation to and receive a folding from one worker node (including latency);
- t_{Map} : time taken by a single worker node to execute the higher-order function Map over the entire list A;
- t_{Rdc} : time taken by a single worker node to execute the higher-order function *Reduce* over the entire list *B*;
- t_p : time taken by the master node to process the result received from the worker nodes and check the termination criteria (steps 7 and 9 that do not depend on K).

We will also use the parameter t_a that denotes the time taken by a node (master or worker) to execute the operation \oplus being the second parameter of the higherorder function *Reduce*:

$$t_a = \frac{t_{Rdc}}{l-1}.\tag{6}$$

First, let us consider the performance of Algorithm 2 on a BSF-computer consisting of one master node and one worker node (see Fig. 4). Here, the dashed arrows denote the data transfers, and the dotted arrows denote the computation loops. Let T_1 denote the execution time of one iteration of Algorithm 2 by a BSF-computer with one master node and one worker node. Using the cost parameters introduced above, we obtain the following estimation of the time T_1 :

$$T_1 = t_p + t_c + t_{Map} + t_{Rdc}.$$
 (7)



Figure 4: Diagram of Algorithm 2 for the configuration with one master and one worker.

Second, let us consider the performance of Algorithm 2 on a BSF-computer consisting of one master node and K worker nodes (see Fig. 5). We assume from now on that $l \ge K$. Let T_K denote the execution time of one iteration of Algorithm 2 by a BSF-computer with one master node and K worker nodes. It is known that a good MPI implementation would implement a broadcast or allreduce for K processes with $O(\log K)$ [34]. If we use $MPI_Broadcast$ to implement Step 2 and MPI_Reduce to implement Step 5 of Algorithm 2 then we can obtain the following estimation for T_K :

$$T_{K} = (K-1)t_{a} + t_{p} + (\log_{2}(K) + 1)t_{c} + \frac{t_{Map} + (l-K)t_{a}}{K}.$$
 (8)

Note that for K = 1, this equation is converted to equation (7).

In the BSF model, the speedup a as a function of K is calculated as follows:

$$a_{BSF}(K) = \frac{T_1}{T_K} = \frac{t_p + t_c + t_{Map} + t_{Rdc}}{(K-1)t_a + t_p + (\log_2(K) + 1)t_c + \frac{t_{Map} + (l-K)t_a}{K}}, \quad (9)$$

where K is the number of worker nodes. For positive values of all parameters and for $K \ge 1$, the function $a_{BSF}(K)$ defined by equation (9) has the following properties:

$$a_{BSF}(1) = 1;$$
 (10)

$$a_{BSF}(K) > 0; \tag{11}$$

$$\lim_{t_{comp} \to 0} a_{BSF}(K) = \frac{1}{\log_2(K) + 1},$$
(12)



Figure 5: Diagram of Algorithm 2 for the configuration with one master and K workers.

where $t_{comp} = t_{Map} + t_{Rdc} + t_p$. All of these properties follow directly from equation (9) and do not require proofs. From the content point of view, property (10) corresponds to reality: speedup on a single worker node must be equal to 1. Property (11) also confirms the adequacy of equation (9) since the speedup is always a positive quantity. Property (12) says that for very small values of parameters t_{Map} , t_{Rdc} and t_p that determine the total time of computations, equation (9) tends to equation $a_{BSF}(K) = 1/(\log_2(K) + 1)$. The last one, on the interval $[1, +\infty)$, determines a monotonically decreasing function that has a maximum value 1 for K = 1. This means that the BSF model is not applicable for algorithms in which the time spent on data transferring between the processor nodes is incomparably greater than the time spent on computations. In this case, one should use another parallel computation model (see the survey [10]). We state the main property of equation (9) in the form of the following proposition:

Proposition 1. Let $l \in \mathbb{N}$; $L, t_c, t_p \in \mathbb{R}_{>0}$; $t_{Map}, t_a \in \mathbb{R}_{\geq 0}$; $t_{Map} + t_a > 0$. Then, the function $a_{BSF}(K)$ defined by equation (9) has a single extremum on the interval $(1, +\infty)$, which is the maximum. PROOF. To find the extrema of function (9), let us calculate the derivative of the speedup with respect to K:

$$a'(K) = \frac{\left(t_p + t_c + t_{Map} + (l-1)t_a\right) \cdot \left(\frac{t_a K + t_{Map} + (l-K)t_a}{K^2} - t_a - \frac{t_c}{K\ln 2}\right)}{\left(\left(K-1\right)t_a + t_p + \left(\log_2\left(K\right) + 1\right)t_c + \frac{t_{Map} + (l-K)t_a}{K}\right)^2}$$

Multiplying the numerator and denominator by K^2 and combining the like terms in the denominator, we obtain

$$a'(K) = \frac{(t_p + t_c + t_{Map} + (l-1)t_a) \cdot (-t_a K^2 - Kt_c / \ln 2 + t_{Map} + lt_a)}{(K(K-1)t_a + Kt_p + K(\log_2(K) + 1)t_c + t_{Map} + (l-K)t_a)^2}.$$
 (13)

Extrema are reached at points where the derivative is zero. Therefore, we need to solve the following equation:

$$\frac{(t_p + t_c + t_{Map} + (l-1)t_a) \cdot \left(-t_a K^2 - K t_c / \ln 2 + t_{Map} + l t_a\right)}{(K(K-1)t_a + K t_p + K \left(\log_2\left(K\right) + 1\right) t_c + t_{Map} + (l-K)t_a)^2} = 0$$

Under the conditions of the proposition, the first factor of numerator and the denominator in this equation are positive for all K > 0. Hence, this equation is equivalent to the following quadratic equation

$$-t_a K^2 - (t_c / \ln 2 + t_a) K + t_{Map} + lt_a = 0$$

that has only one root on the interval $[1, +\infty)$:

$$K_0 = \frac{1}{2} \sqrt{\left(\frac{t_c}{t_a \ln 2}\right)^2 + \frac{t_{Map}}{t_a} + 4l - \frac{t_c}{t_a \ln 2}}.$$

Since K^2 has a negative coefficient, the derivative a'(K) calculated by equation (13) takes only positive values in the interval $[1, K_0)$ and only negative values in the interval $(K_0, +\infty)$. Therefore, the point K_0 is the maximum of the function $a_{BSF}(K)$ on the interval $[1, +\infty)$. The proposition is proven.

Proposition 1 gives us the following equation to evaluate the scalability boundary of a BSF-algorithm:

$$K_{BSF} = \frac{1}{2} \sqrt{\left(\frac{t_c}{t_a \ln 2}\right)^2 + \frac{t_{Map}}{t_a} + 4l - \frac{t_c}{t_a \ln 2}}.$$
 (14)

It is noteworthy that the scalability boundary of a parallel BSF-algorithm does not depend on the time t_p taken by the master node to process the result received from the worker nodes and check the termination criteria (steps 7 and 9 of Algorithm 2). This is quite natural. Indeed, t_p does not depend on the number of workers K and therefore cannot affect the point of speedup curve maximum that is completely determined by the derivative of the speedup with respect to K.

5. Applying the BSF model to the Jacobi method

In this section, we will show how to apply the BSF model to estimate the scalability boundary of a parallel algorithm without software implementation and computational experiments. As an example, we use the Jacobi iterative method. The *Jacobi method* [24] is a simple iterative method for solving a system of linear equations. This method was originally described by the German mathematician Carl Gustav Jacob Jacobi in [25]. Let us give a brief description of the Jacobi method.

Let a joint square system of linear equations in a matrix form be given in Euclidean space \mathbb{R}^n :

$$Ax = b, (15)$$

where

$$A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix};$$
$$x = (x_1, \dots, x_n);$$
$$b = (b_1, \dots, b_n).$$

It is assumed that $a_{ii} \neq 0$ for all i = 1, ..., n. Let us define the matrix

$$C = \begin{pmatrix} c_{11} & \cdots & c_{1n} \\ \vdots & \ddots & \vdots \\ c_{n1} & \cdots & c_{nn} \end{pmatrix}$$

in the following way:

$$c_{ij} = \begin{cases} -\frac{a_{ij}}{a_{ii}}, \forall j \neq i; \\ 0, \forall j = i. \end{cases}$$

Let us define the vector $d = (d_1, \ldots, d_n)$ as follows: $d_i = b_i/a_{ii}$. The Jacobi method of finding an approximate solution of system (15) consists of the following steps:

Step 1.
$$k := 0; x^{(0)} := d.$$

Step 2. $x^{(k+1)} := Cx^{(k)} + d.$
Step 3. If $||x^{(k+1)} - x^{(k)}||^2 < \varepsilon$, go to Step 5.
Step 4. $k := k + 1$; go to Step 2.
Step 5. Stop.

In the Jacobi method, an arbitrary vector $x^{(0)}$ can be taken as the initial approximation. In Step 1, the initial approximation $x^{(0)}$ is assigned by the vector d. In Step 3, the Euclidean norm $\|\cdot\|$ is used in the termination criteria. The *diagonal dominance* of the matrix A is a sufficient condition for the convergence of the Jacobi method:

$$|a_{ii}| \ge \left(\sum_{j=1}^{n} |a_{ij}|\right) - |a_{ii}|$$

for all i = 1, ..., n, and at least one inequality is strict. In this case, the system (15) has a unique solution for any right-hand side.

Let us represent the Jacobi method in the form of an algorithm on lists. Let c_j denote the *j*-th column of matrix C:

$$c_j = \left(\begin{array}{c} c_{1j} \\ \vdots \\ c_{nj} \end{array}\right).$$

Let G = [1, ..., n] be the list of natural numbers from 1 to n. For any vector $x = (x_1, ..., x_n) \in \mathbb{R}^n$, let us define the function $F_x : \{1, ..., n\} \to \mathbb{R}^n$ as follows:

$$F_x(j) = x_j c_j = \begin{pmatrix} x_j c_{1j} \\ \vdots \\ x_j c_{nj} \end{pmatrix},$$
(16)

i.e., the function $F_x(j)$ multiplies the *j*-th column of the matrix C by the *j*-th coordinate of the vector x. The BSF implementation of the Jacobi method

Algorithm 3. BSF-Jacobi algorithm.

1: input C, G, d2: $k := 0; x^{(0)} := d$ 3: $B := Map(F_{x^{(k)}}, G)$ 4: $s := Reduce(\vec{+}, B)$ 5: $x^{(k+1)} := s\vec{+}d$ 6: k := k + 17: if $||x^{(k+1)}\vec{-}x^{(k)}||^2 < \varepsilon$ goto 9 8: goto 3 9: output $x^{(k)}$ 10: stop

presented as Algorithm 3 can be easily obtained from the generic BSF-algorithm template (Algorithm 1). In Algorithm 3, $\vec{+}$ and $\vec{-}$ denote the operations of vector addition and subtraction, respectively. Note that the matrix C entered in line 1 is implicitly used to calculate the values of the function $F_{x^{(k)}}$ in line 3.

The BSF-Jacobi parallel algorithm (see Algorithm 4) is automatically generated from Algorithm 3 by using the generic BSF-algorithm parallelization template (Algorithm 2). Let us evaluate this parallel algorithm by using the BSF model. We assume that all arithmetic operations (addition and multiplication) as well as the comparison operation of floating-point numbers take the same time, which we denote as τ_{op} . To perform the scalability analysis of the BSF Jacobi algorithm, let us introduce the following notation (all quantities are taken with respect to a single iteration):

- c_c : the quantity of real numbers that the master sends to and receives from a single worker within one iteration;
- c_{Map} : the quantity of arithmetic operations performed in Step 3 of the Algorithm 3;
- c_a : the quantity of arithmetic operations required to calculate the sum of two vectors.

Master	j th Worker $(j=1,\ldots,K)$
1: input d	1: input C_j, G_j
2: $k := 0; x^{(0)} := d$	2:
3: $SendToAllWorkers(x^{(k)})$	3: $RecvFromMaster(x^{(k)})$
4:	4: $B_j := Map(F_{x^{(k)}}, G_j)$
5:	5: $s_j := Reduce(\vec{+}, B_j)$
6: $RecvFromWorkers(s_1,\ldots,s_K)$	6: $SendToMaster(s_j)$
7: $s := Reduce\left(\vec{+}, [s_1, \dots, s_K]\right)$	7:
8: $x^{(k+1)} := s \vec{+} d$	8:
9: $k := k + 1$	9:
10: $exit := \ x^{(k+1)} - x^{(k)}\ ^2 < \varepsilon$	10:
11: SendToAllWorkers(exit)	11: $RecvFromMaster(exit)$
12: if not <i>exit</i> goto 3	12: if not <i>exit</i> goto 3
13: output $x^{(k)}$	13:
14: stop	14: stop

Algorithm 4. BSF-Jacobi parallel algorithm.

Let us calculate these quantities. At the beginning of the iteration, the master sends to each worker the current approximation $x^{(k)}$, which is a vector of length n. At the ending of the iteration, each worker sends the calculated vector s_j of length n to the master. Hence,

$$c_c = 2n. \tag{17}$$

The higher-order function $Map(F_{x^{(k)}}, G)$, in this case, multiplies all columns of the matrix C by the corresponding coordinates of the vector x. Consequently,

$$c_{Map} = n^2. (18)$$

Adding two vectors of length n requires n arithmetic operations. Thus,

$$c_a = n. \tag{19}$$

Let τ_{op} be the average execution time of a single arithmetic or comparison operation by the processor node, and τ_{tr} be the average time for transferring a single floating number across the network excluding latency. Using (17)-(19), we obtain the following values of the cost parameters of the BSF-Jacobi parallel algorithm:

$$t_c = c_c \tau_{tr} + 2L = 2(n\tau_{tr} + L); \tag{20}$$

$$t_{Map} = c_{Map}\tau_{op} = n^2\tau_{op}; \tag{21}$$

$$t_a = c_a \tau_{op} = n \tau_{op}. \tag{22}$$

For the BSF-Jacobi algorithm, the length l of the list is equal to the space dimension n:

$$l = n. (23)$$

Substituting the values of the right-hand sides of equations (20)-(23) into equation (14), we obtain the following equation for estimating the scalability boundary of the BSF-Jacobi parallel algorithm:

$$K_{BSF-Jacobi} = \sqrt{\left(\frac{n\tau_{tr} + L}{n\tau_{op}\ln 2}\right)^2 + \frac{5}{2}n} - \frac{n\tau_{tr} + L}{n\tau_{op}\ln 2}.$$
 (24)

For large values of n, this is equivalent to

$$K_{BSF-Jacobi} \approx O(\sqrt{n}).$$
 (25)

Therefore, we can conclude that the scalability boundary of the BSF-Jacobi parallel algorithm grows in proportion to the square root of the problem dimension n. It should be noted that this result was obtained before a software implementation of the BSF-Jacobi parallel algorithm. In the next section, we verify this analytical estimation by computational experiments on a real cluster system.

Table 1: Specifications of "Tornado SUSU" computing cluster.

Parameter	480
Processor	Intel Xeon X5680 (6 cores, 3.33 GHz)
Processors per node	2
Memory per node	24 GB DDR3
Interconnect	InfiniBand QDR (40 Gbit/s)
Operating system	Linux CentOS

6. Computational experiments

For the rapid development of the parallel BSF-programs, the author implemented an algorithmic skeleton in C++ using the MPI parallel programming library [26]. The source code of this BSF-skeleton [27] is freely available on GitHub, at https://github.com/leonid-sokolinsky/BSF-skeleton. Using the BSF-skeleton, we developed the parallel implementations of several iterative numerical methods and performed the computational experiments on the "Tornado SUSU" computing cluster [28], whose specifications are shown in Table 1. In this section, we present some of the results of these computational experiments and compare them with analytical results obtained by using the BSF cost metric.

The first series of experiments was performed with the BSF-Jacobi parallel algorithm discussed in Section 5. The source code implemented by using the BSF-skeleton is freely available on GitHub, at https://github.com/ leonid-sokolinsky/BSF-Jacobi. To carry out the experiments, we used a scalable system of linear equations (15) having the following coefficient matrix A and the vector of constant terms b:

$$A = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 1 \\ 1 & \cdots & 1 & n \end{pmatrix}; b = \begin{pmatrix} n \\ n+1 \\ \vdots \\ 2n-1 \end{pmatrix}.$$

n	1500	5000	10000	16000
t_c	7.20E-5	1.06E-3	2.17E-3	2.95E-3
t_p	5.01E-6	1.72E-5	3.70E-5	5.61E-5
t_a	1.89E-6	5.27E-6	9.31E-6	2.10E-5
t_{Map}	6.23E-3	9.28E-2	3.73E-1	7.73E-1
$\frac{comp}{comm}$	126	113	215	376

Table 2: Cost parameters for BSF-Jacobi parallel algorithm (seconds).

The specified system has a unique solution x = (1, ..., 1) and has the diagonal dominance property for any $n \ge 2$. We investigated the speedup of the BSF-Jacobi parallel algorithm by varying the number K of working nodes.

The speedup $a_{test}(K)$ was calculated by the equation

$$a_{test}(K) = \frac{T_1}{T_K},$$

where T_1 is the execution time on configuration with one master node and one worker node, and T_K is the execution time on configuration with one master node and K worker nodes. The computations were performed for dimensions n = 1500, n = 5000, n = 10000 and n = 16000. The results are presented in Fig. 6 (the solid curves marked with squares). In the same diagrams, we plotted the curves of speedup calculated analytically by equation (9) (the dotted curves marked with crosses). The diagrams in Fig. 6 also include the values of scalability boundaries obtained using equation (14). This values are flagged with vertical dotted lines. To calculate these analytical estimations, we experimentally determined the values of cost parameters using a configuration with one master and one worker. These values are shown in Table 2. In all cases the latency L (time of transferring one-byte message node-to-node) was 1.5E-5sec. We also added the cost ratio between computations and communications to the table. Here $comp = t_{Map} + (n-1)t_a + t_p$ and $comm = t_c$. The cost of computation significantly exceeds the cost of communications in the BSF-Jacobi algorithm, and this gap tends to increase with the size of the problem. The fact



Figure 6: BSF-Jacobi parallel algorithm speedup obtained empirically and theoretically.

that the value of $\frac{comp}{comm}$ for n = 1500 is greater than the value for n = 5000 is explained as follows. When the cost of a single data exchange becomes comparable to the latency, the latency begins to significantly affect the result.

Let us define the prediction error for the problem of dimension n as follows:

$$Error(n) = \frac{|K_{test}(n) - K_{BSF}(n)|}{\max(K_{test}(n), K_{BSF}(n))},$$
(26)

where $K_{test}(n)$ is the speedup boundary obtained experimentally, and $K_{BSF}(n)$ is the speedup boundary obtained analytically by (14). For the Jacobi-BSF parallel algorithm, we obtain the results shown in Table 3. In all cases, the prediction error does not exceed 15% and tends to decrease with increasing the problem size.

n	1500	5000	10000	16000
K_{BSF}	47	64	112	150
K_{test}	40	60	120	160
Error	0.15	0.06	0.07	0.06

Table 3: Prediction errors for BSF-Jacobi parallel algorithm.

The second series of experiments, the results of which we want to present in this article, relate to a simplified *n*-body problem [29] that describes how a small body will move under the influence of gravitational forces among *n* large motionless bodies. Let us give a brief description of this problem. Let $\mathbb{Y} \subset \mathbb{R}^3$ be a finite set of points representing motionless bodies of large mass. We will denote these points by Y_1, \ldots, Y_n , and their masses by m_1, \ldots, m_n , where $n = |\mathbb{Y}|$. Let the point X represent a body x of low mass m_x moving relative to motionless large mass bodies Y_1, \ldots, Y_n . We assume that no forces other than gravity act on the body x. We know the initial position $X^{(t_0)} \in \mathbb{R}^3$ and the velocity vector $V^{(t_0)} \in \mathbb{R}^3$ of the body x at the instant of time t_0 . The problem is to predict the subsequent motion of the body x using Newton's laws of motion and Newton's law of universal gravitation. To accomplish this, we will sequentially calculate the following positions of the body x using a time slot Δt :

$$X^{(t_0)}, X^{(t_0 + \Delta t)}, X^{(t_0 + 2\Delta t)}, X^{(t_0 + 3\Delta t)}, \dots$$
(27)

According to the law of universal gravitation, the gravitational force F_i of the mass point Y_i acting on the body x can be calculated using the equation

$$F_{i} = G \frac{m_{i}m_{x}}{\left\|Y_{i} - X\right\|^{2}} (Y_{i} - X),$$
(28)

where X represents the current coordinates of the body x. According to Newton's second law of motion, the acceleration α_i of the body x produced by the force F_i can be calculated using the equation

$$\alpha_i = \frac{F_i}{m_x}.\tag{29}$$

Algorithm 5. BSF-Gravity algorithm.

1: input A, X_0, V_0, t_0, T 2: $t := t_0; X := X_0; V := V_0;$ 3: $B := Map(f_X, A)$ 4: $\alpha := Reduce(\vec{+}, B)$ 5: $\Delta t := Delta_t(V, \alpha)$ 6: $V := V \vec{+} \alpha \Delta t$ 7: $X := X \vec{+} V \Delta t$ 8: $t := t + \Delta t$ 9: if t < T goto 3 10: output X 11: stop

The acceleration produced by the all forces F_1, \ldots, F_n is calculated by the equation

$$\alpha = \sum_{i=1}^{n} \alpha_i. \tag{30}$$

Consequently, the velocity vectors required to calculate (approximately) the elements of the sequence (27) can be calculated using the following iterative equation

$$V^{(t+\Delta t)} = V^{(t)} + \alpha^{(t+\Delta t)} \Delta t, \qquad (31)$$

where

$$\alpha^{(t+\Delta t)} = \sum_{i=1}^{n} G \frac{m_i}{\left\|Y_i - X^{(t)}\right\|^2} \left(Y_i - X^{(t)}\right).$$
(32)

Using (31), we obtain

$$X^{(t+\Delta t)} \approx X^{(t)} + V^{(t+\Delta t)} \Delta t.$$
(33)

Using the generic BSF-algorithm template, we obtain Algorithm 5 named the BSF-Gravity that presents an implementation of the simplified *n*-body problem. In the context of this problem, A is the list of pairs (Y_i, m_i) that specify the coordinates and mass of the *i*-th motionless large body:

$$A = [(Y_1, m_1), \dots, (Y_n, m_n)];$$
(34)

Master	jth Worker (j=1,,K)
1: input X_0, V_0, t_0, T	1: input A_j
2: $t := t_0; X := X_0; V := V_0;$	2:
3: SendToAllWorkers(X)	3: $RecvFromMaster(X)$
4:	4: $B_j := Map(f_X, A_j)$
5:	5: $\alpha_j := Reduce(\vec{+}, B_j)$
6: $RecvFromWorkers(\alpha_1,\ldots,\alpha_K)$	6: $SendToMaster(\alpha_j)$
7: $\alpha := Reduce\left(\vec{+}, [\alpha_1, \dots, \alpha_K]\right)$	7:
8: $\Delta t := Delta_t(V, \alpha)$	8:
9: $V := V \vec{+} \alpha \Delta t$	9:
10: $X := X + V \Delta t$	10:
11: $t := t + \Delta t$	11:
12: $exit := t < T$	12:
13: SendToAllWorkers(exit)	13: RecvFromMaster(exit)
14: if not <i>exit</i> goto 3	14: if not <i>exit</i> goto 3
15: output X	15:
16: stop	16: stop

Algorithm 6. BSF-Gravity parallel algorithm.

T is the instant of time that culminates the calculation of the trajectory of the body x; f_X is the parameterized function defined by the equation

$$f_X(Y_i, m_i) = G \frac{m_i}{|Y_i - X|^2} (Y_i - X)$$
(35)

that the higher-order function Map takes as a parameter to apply to items of the list A; and $\vec{+}$ denotes the operation of vector addition in \mathbb{R}^3 . The user function $Delta_t(V, \alpha)$ calculates the time slot Δt depending on the current velocity Vand acceleration α . The parallel implementation of the BSF-Gravity algorithm is obtained automatically by using the generic BSF-algorithm parallelization template (Algorithm 2). This implementation is presented in Algorithm 6.



Figure 7: BSF-Gravity parallel algorithm speedup obtained empirically and theoretically.

Let us evaluate Algorithm 6. Assume that the user function $Delta_t(V, \alpha)$ is defined as follows:

$$Delta_t(V,\alpha) = \frac{\eta}{\left\|V\right\|^2 \cdot \left\|\alpha\right\|^4},$$

where η is a positive constant. Calculating this function takes 13 arithmetic operations. The analysis of Algorithm 6 gives us the following estimations: $t_c = 6\tau_{tr} + 2L$, $t_{Map} = 17n\tau_{op}$, $t_a = 3\tau_{op}$ and l = n. Substituting these quantities in (14), we obtain the following scalability boundary for the BSF-Gravity parallel algorithm:

$$K_{BSF-Gravity} = \frac{1}{2} \sqrt{\left(\frac{6\tau_{tr} + 2L}{3\tau_{op}\ln 2}\right)^2 + \frac{29}{3}n} - \frac{6\tau_{tr} + 2L}{3\tau_{op}\ln 2}.$$
 (36)

Table 4: Prediction errors for BSF-Gravity parallel algorithm.

n	300	600	900	12000
K_{BSF}	69	141	210	279.1
K_{test}	60	140	200	280
Error	0.13	0.01	0.05	3.6E-4

For $n \to \infty$, equation (36) asymptotically tends to the following estimation:

$$K_{BSF-Gravity} \approx O(\sqrt{n}),$$
 (37)

where n is the number of motionless large bodies.

We implemented the BSF-Gravity parallel algorithm using the BSF-skeleton. The source code of this implementation is freely available on GitHub, at https://github.com/leonid-sokolinsky/BSF-gravity. Using this implementation, we conducted the computational experiments on the "Tornado SUSU" computing cluster. The computations were performed for the following numbers of large bodies: n = 300, n = 600, n = 900 and n = 1200. A comparison of the results obtained empirically and theoretically (by using equation (9)) is shown in Fig. 7. To plot the curve tagged as "BSF", we experimentally determined the following cost parameters that are independent of n: $t_c = 5 \cdot 10^{-5}$, $t_p = 9.5 \cdot 10^{-7}, t_a = 4.7 \cdot 10^{-9}$ and $L = 1.5 \cdot 10^{-5}$ (in seconds). For the parameter t_{Map} that depends on n, the following values in seconds were obtained: $3.6 \cdot 10^{-3}$, $7.46 \cdot 10^{-3}$, $1.12 \cdot 10^{-2}$, $1.5 \cdot 10^{-2}$ for *n* equal to 300, 600, 900 and 1200, respectively. The cost ratio between computations and communications varied from 100 at n = 300 to 411 at n = 1200. The error values calculated using equation (26) are shown in table 4. In all cases, the error does not exceed 13%and tends to decrease with increasing the problem size.

Another example of using the BSF model can be found in [30], where we investigate an iterative numerical method for solving nonstationary systems of linear inequalities. In this article, we also compare the speedup curves constructed analytically using the BSF model and the speedup curves obtained by conducting experiments on the real cluster computing system. All conducted experiments confirm the adequacy of the BSF model.

7. Discussion

In this section, we discuss the strengths and weaknesses of the BSF model and answer the following questions.

- 1. Is it possible to use the BSF model for the algorithms that process sets?
- 2. Can we apply the BSF model to an algorithm that uses only the *Map* function without the *Reduce* function?
- 3. Can we apply the BSF model to a numerical algorithm that is not iterative?
- 4. What is the difference between MapReduce and the BSF model?
- 5. Does the BSF model admit a configuration of the BSF-computer with two or more master nodes?
- 6. Does the BSF model take into account the multicore structure of the processor node?
- 7. Is the BSF model the best predictor of the execution time of a parallel algorithm on a target multiprocessor system?

Let us start by discussing the advantages of the BSF model. The main contribution of the BSF model and this article is equation (14) that allows us to estimate the scalability boundary of a parallel program at an early stage of its design. No other known model yields such an equation. In addition, the BSF model is easy to use when designing and analyzing parallel algorithms and programs. Based on the BSF model, we constructed a compilable algorithmic skeleton using the MPI library that allows quick creation of a syntactically valid BSF program. However, to acquire this result, we introduced a number of constraints.

First, the BSF model requires the representation of a numerical method as an algorithm over lists using the higher-order functions *Map* and *Reduce*. How strong is this restriction? We will start by answering the question whether the BSF model can be applied to multisets (sets that allows duplicates) [31]. Let $\mathcal{A} = \{a_1, \ldots, a_n\}$ be a finite multiset. By ordering, we can represent \mathcal{A} as the list $A = [a_1, \ldots, a_n]$. Let $B = Map(f, A) = [f(a_1), \ldots, f(a_n)]$. Ignoring ordering, we can transform B to the multiset $\mathcal{B} = \{f(a_1), \ldots, f(a_n)\}$. Thus, we can apply the BSF model to multisets without any modifications. Applying the BSF model to sets is also possible if the function f does not generate duplicates. Taking into account the generic nature of the *Map* function, we can conclude that the representation of the numerical method as an algorithm over lists is not a very strong restriction.

The second question is: Can we apply the BSF model to an algorithm that uses only the *Map* function without the *Reduce* function? The answer is "yes". In this case, the parameter t_a denoting the time of execution of the operation \oplus is assumed to be zero. An example of applying the BSF model to an algorithm using only the *Map* function is presented in [32].

Next, we limited the scope of the BSF model application to the iterative compute-intensive numerical algorithms. The iterative nature of the algorithm assumes that the iterative process takes much longer than initialization (reading or generating the problem data, allocating the memory, etc.), and we can neglect the cost of the last (there is no corresponding parameter in the BSF cost metric). Consequently, we can use the BSF model if our noniterative algorithm is presented in the form of operations over lists, and the initialization cost is negligible compared to the execution cost of the *Map* and *Reduce* functions. This is the answer to the third question stated above.

The compute-intensive nature of the algorithm assumes that the cost of computations is greater than or comparable to the cost of interprocessor communications and input/output operations. If we have the opposite situation, when the communications and input/output operations significantly exceed the calculations, then, due to property (12), we have $K_{BSF} = 1$, and the BSF model becomes inapplicable. This is the main difference between the BSF model and the programming model MapReduce [33] that is intended for processing and generating big data sets. This is the answer to the 4th question.

The 5th question is: Does the BSF model admit a configuration of the BSF-computer with two or more master nodes? The answer is "no". All attempts to derive an equation such as (14) failed in the case of configurations with two or more master nodes.

The 6th question is: Does the BSF model take into account the multicore structure of the processor node? The answer is also "no": the model treats the processor node as a black box that can perform scalar and vector operations at a certain speed. Most of the modern cluster computing systems have the processor nodes that include multicore processors and GPUs. The efficient usage of such systems is impossible without the use of intranode parallelism and vectorization. However, in this situation, how can we determine the adequate values of the cost parameters t_{Map} (the time taken by a single processor node to process the whole list by using the higher-order function Map and t_a (the time taken by a processor node to execute the operation \oplus)? Let us explain the possible solution of this problem by the following example. Let \oplus be the addition operation of two vectors of dimension 1 000 000. We write a program that computes the sum of $1\,000\,000$ such vectors using all resources of the intranode parallelism and vectorization. Then, we run this program on a single processor node of the target system and measure the execution time (excluding initialization). After that, we divide the resulting time by 1000000 and obtain the value of the parameter t_a . In the same way, we can obtain the value of the parameter t_{Map} . This is a quite rough method for obtaining the values of these parameters. However, the main goal of the BSF model is to predict the scalability boundary being the maximum number of processor nodes to which the speedup increases. If a more accurate prediction of the execution time of a parallel program on a computing cluster with multicore nodes is needed, then other parallel computation models should be used (see, for example, [35, 36]). Therefore, the BSF model is not the best predictor of the execution time of a parallel algorithm on a target multiprocessor system. This is the answer to the last question of our discussion.

8. Conclusions

In this paper, we presented a novel parallel computation model named bulk synchronous farm (BSF) that is an extension of the BSP (bulk-synchronous parallel) model. The main advantage of the proposed model is that it allows to estimate the scalability of a parallel algorithm before its implementation. Another important feature of the BSF model is the representation of problem data in the form of lists that greatly simplifies the logic of building applications. To develop the BSF model, we used a new approach that restricts not only the class of multiprocessor architectures but also the type of algorithms admitted by the model. The application scope of the BSF model is algorithms of the computeintensive iterative type performed on cluster computing systems. In the BSF model, the processor nodes are organized by using the master/slave paradigm. The BSF model uses the Map/Reduce scheme to parallelize applications. We constructed a cost metric of the BSF model. This metric requires the representation of a numerical method as an algorithm over lists using the higher-order functions Map and Reduce. Using this metric, we derived equation (14) that allowed us to estimate the scalability boundary of a parallel algorithm before its software implementation. No other known model yields such an equation. Based on the BSF model, we constructed a compilable algorithmic skeleton using the MPI library that allows the quick creation of a syntactically valid BSF-program. Using this skeleton, we implemented several iterative algorithms and conducted large-scale computational experiments on a real cluster computing system. In all cases, experiments have shown that the scalability boundary of the algorithm obtained analytically using the BSF cost metric is very close to the scalability boundary obtained experimentally. Our experience with the BSF model has shown that this model is accurate and easy to use when designing and analyzing parallel algorithms and programs.

9. Declaration of competing interests

The author declare that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

10. Acknowledgments

This work was supported by the Russian Foundation for Basic Research [project No. 20-07-00092-a] and the Ministry of Science and Higher Education of the Russian Federation [gov. order FENU-2020-0022].

References

- J. Dongarra, S. Gottlieb, W.T.C. Kramer, Race to Exascale, Comput. Sci. Eng. 21 (2019) 4-5. https://doi.org/10.1109/MCSE.2018.2882574.
- [2] TOP500 Supercomputer Sites. https://www.top500.org (accessed December 13, 2020).
- [3] X. Wu, W. Li, Performance models for scalable cluster computing, J. Syst. Archit. 44 (1998) 189-205. https://doi.org/10.1016/S1383-7621(97) 00036-2.
- Y. Zhang, G. Chen, G. Sun, Q. Miao, Models of parallel computation: a survey and classification, Front. Comput. Sci. China. 1 (2007) 156-165. https://doi.org/10.1007/s11704-007-0016-1.
- [5] L.G. Valiant, A bridging model for parallel computation, Commun. ACM.
 33 (1990) 103-111. https://doi.org/10.1145/79173.79181.
- [6] C.C. Elgot, A. Robinson, Random-Access Stored-Program Machines, an Approach to Programming Languages, J. ACM. 11 (1964) 365-399. https: //doi.org/10.1145/321239.321240.

- J. Hartmanis, Computational complexity of random access stored program machines, Math. Syst. Theory. 5 (1971) 232-245. https://doi.org/10. 1007/BF01694180.
- [8] S.A. Cook, R.A. Reckhow, Time bounded random access machines, J. Comput. Syst. Sci. 7 (1973) 354-375. https://doi.org/10.1016/ S0022-0000(73)80029-7.
- [9] D.B. Skillicorn, D. Talia, Models and languages for parallel computation, ACM Comput. Surv. 30 (1998) 123-169. https://doi.org/10.1145/ 280277.280278.
- [10] J.A. Rico-Gallego, J.C. Diaz-Martin, R.R. Manumachu, A.L. Lastovetsky, A Survey of Communication Performance Models for High-Performance Computing, ACM Comput. Surv. 51 (2019) 1-36. https://doi.org/10. 1145/3284358.
- [11] S. Fortune, J. Wyllie, Parallelism in random access machines, in: Proc. Tenth Annu. ACM Symp. Theory Comput. — STOC'78, ACM Press, New York, New York, USA, 1978: pp. 114-118. https://doi.org/10.1145/ 800133.804339.
- [12] D. Culler, R. Karp, D. Patterson, A. Sahay, K.E. Schauser, E. Santos, R. Subramonian, T. Von Eicken, LogP: towards a realistic model of parallel computation, in: Proc. Fourth ACM SIGPLAN Symp. Princ. Pract. Parallel Program. PPOPP'93, ACM Press, New York, New York, USA, 1993: pp. 1-12. https://doi.org/10.1145/155332.155333.
- [13] L. Yuan, Y. Zhang, Y. Tang, L. Rao, X. Sun, LogGPH: A Parallel Computational Model with Hierarchical Communication Awareness, in: Proc. 2010 13th IEEE Int. Conf. Comput. Sci. Eng. CSE'10, IEEE Computer Society, Washington, DC, US, 2010: pp. 268-274. https://doi.org/10. 1109/CSE.2010.40.

- [14] F. Lu, J. Song, Y. Pang, HLognGP: A parallel computation model for GPU clusters, Concurr. Comput. Pract. Exp. 27 (2015) 4880-4896. https: //doi.org/10.1002/cpe.3475.
- [15] X. Qiao, S. Chen, L.T. Yang, HPM: a hierarchical model for parallel computations, Int. J. High Perform. Comput. Netw. 1 (2004) 117-127. https://doi.org/10.1504/IJHPCN.2004.007571.
- [16] J.-A. Rico-Gallego, J.-C. Diaz-Martin, τ-Lop: Modeling performance of shared memory MPI, Parallel Comput. 46 (2015) 14-31. https://doi. org/10.1016/J.PARCO.2015.02.006.
- [17] J.-A. Rico-Gallego, A.L. Lastovetsky, J.-C. Diaz-Martin, Model-Based Estimation of the Communication Cost of Hybrid Data-Parallel Applications on Heterogeneous Clusters, IEEE Trans. Parallel Distrib. Syst. 28 (2017) 3215-3228. https://doi.org/10.1109/TPDS.2017.2715809.
- [18] G. Bilardi, C. Fantozzi, A. Pietracaprina, G. Pucci, On the Effectiveness of D-BSP as a Bridging Model of Parallel Computation, in: Proc. Int. Conf. Comput. Sci. — ICCS '01. Part II. Lect. Notes Comput. Sci. Vol. 2074, Springer, Berlin, Heidelberg, 2001: pp. 579-588. https://doi.org/ 10.1007/3-540-45718-6_63.
- [19] G. Bilardi, A. Pietracaprina, Models of Computation, Theoretical, in: Encycl. Parallel Comput., Springer US, Boston, MA, 2011: pp. 1150-1158. https://doi.org/10.1007/978-0-387-09766-4_218.
- [20] L.B. Sokolinsky, Analytical Estimation of the Scalability of Iterative Numerical Algorithms on Distributed Memory Multiprocessors, Lobachevskii J. Math. 39 (2018) 571-575. https://doi.org/10.1134/ S1995080218040121.
- [21] S. Sahni, G. Vairaktarakis, The master-slave paradigm in parallel computer and industrial settings, J. Glob. Optim. 9 (1996) 357-377. https://doi. org/10.1007/BF00121679.

- [22] R.S. Bird, Lectures on Constructive Functional Programming, in: M. Broy (Ed.), Constr. Methods Comput. Sci. NATO ASI Ser. F Comput. Syst. Sci. Vol. 55, Spring-er, Berlin, Heidlberg, 1988: pp. 151-216.
- [23] G. Malcolm, Homomorphisms and promotability, in: van de S. J.L.A. (Ed.), Math. Progr. Constr. MPC 1989. Lect. Notes Comput. Sci. Vol. 375, Springer, Berlin, Heidelberg, 1989: pp. 335-347. https://doi.org/10.1007/3-540-51305-1_20.
- H. Rutishauser, The Jacobi Method for Real Symmetric Matrices, in: Bauer F.L. (Ed.), Handb. Autom. Comput. Vol 2. Linear Algebr., Springer Berlin Heidelberg, Ber-lin, Heidelberg, 1971: pp. 202-211. https://doi.org/10.1007/978-3-662-39778-7_12.
- [25] C.G.J. Jacobi, Ueber eine neue Auflösungsart der bei der Methode der kleinsten Quadrate vorkommenden lineären Gleichungen, Astron. Nachrichten. 22 (1845) 297-306. https://doi.org/10.1002/asna. 18450222002.
- [26] W. Gropp, MPI 3 and Beyond: Why MPI Is Successful and What Challenges It Faces, in: J.L. Träff, S. Benkner, J.J. Dongarra (Eds.), Recent Adv. Messag. Passing Interface. EuroMPI 2012. Lect. Notes Comput. Sci. Vol. 7490, Springer, Berlin, Heidelberg, 2012: pp. 1-9. https://doi.org/10.1007/978-3-642-33518-1_1.
- [27] L.B. Sokolinsky, BSF-skeleton: user manual, ArXiv:2008.12256 [Cs.DC].
 (2020) 22. http://arxiv.org/abs/2008.12256.
- [28] P. Kostenetskiy, P. Semenikhina, SUSU Supercomputer Resources for Industry and fundamental Science, in: Proc. — 2018 Glob. Smart Ind. Conf. GloSIC 2018, Art. No. 8570068, IEEE, 2018: p. 7. https://doi.org/10. 1109/GloSIC.2018.8570068.
- [29] F. Diacu, The solution of the n-body problem, Math. Intell. 18 (1996) 66-70. https://doi.org/10.1007/BF03024313.

- [30] L.B. Sokolinsky, I.M. Sokolinskaya, Scalable parallel algorithm for solving non-stationary systems of linear inequalities, Lobachevskii J. Math. 41 (2020) 1571-1580. https://doi.org/10.1134/S1995080220080181.
- [31] A. Syropoulos, Mathematics of multisets, in: C.S. Calude, G. Paun, G. Rozenberg, A. Salomaa (Eds.), Multiset Process. WMC 2000. Lect. Notes Comput. Sci. Vol. 2235, Springer, Berlin, Heidelberg, 2001: pp. 347-358. https://doi.org/10.1007/3-540-45523-x_17.
- [32] N.A. Ezhova, L.B. Sokolinsky, Scalability Evaluation of Iterative Algorithms Used for Supercomputer Simulation of Physical processes, in: Proc. 2018 Glob. Smart Ind. Conf. GloSIC 2018, Art. No. 8570131, IEEE, 2018: p. 10. https://doi.org/10.1109/GloSIC.2018.8570131.
- [33] J. Dean, S. Ghemawat, MapReduce: Simplified Data Processing on Large Clusters, Commun. ACM. 51 (2008) 107-113. https://doi.org/10.1145/ 1327452.1327492.
- [34] T. Hoefler, W. Gropp, R. Thakur, J.L. Traff, Toward Performance Models of MPI Implementations for Understanding Application Scaling Issues, in: R. Keller, E. Gabriel, M. Resch, J. Dongarra (Eds.), Recent Adv. Messag. Passing Interface. EuroMPI 2010. Lect. Notes Comput. Sci. Vol 6305., Springer, Berlin, Heidelberg, 2010: pp. 21–30. https://doi.org/10.1007/ 978-3-642-15646-5_3.
- [35] M. Amaris, D. Cordeiro, A. Goldman, R.Y. de Camargo, A Simple BSPbased Model to Predict Execution Time in GPU Applications, in: 2015 IEEE 22nd Int. Conf. High Perform. Comput., IEEE, 2015: pp. 285-294. https://doi.org/10.1109/HiPC.2015.34.
- [36] A. V. Gerbessiotis, Extending the BSP model for multicore and out-of-core computing: MBSP, Parallel Comput. 41 (2015) 90-102. https://doi.org/ 10.1016/j.parco.2014.12.002.

- [37] A. Alexandrov, M.F. Ionescu, K.E. Schauser, C. Scheiman, LogGP: Incorporating Long Messages into the LogP Model for Parallel Computation, J. Parallel Distrib. Comput. 44 (1997) 71-79. https://doi.org/10.1006/ jpdc.1997.1346.
- [38] F. Ino, N. Fujimoto, K. Hagihara, LogGPS: A parallel computational model for synchronization analysis, ACM SIGPLAN Not. 36 (2001) 133-142. https://doi.org/10.1145/568014.379592.
- [39] W. Gropp, E. Lusk, N. Doss, A. Skjellum, A high-performance, portable implementation of the MPI message passing interface standard, Parallel Comput. 22 (1996) 789-828. https://doi.org/10.1016/0167-8191(96) 00024-5.
- [40] K.W. Cameron, R. Ge, X.-H. Sun, log_nP and log₃P: Accurate Analytical Models of Point-to-point Communication in Distributed Systems, IEEE Trans. Comput. 56 (2007) 314–327. https://doi.org/10.1109/TC.2007. 38.
- [41] L.G. Valiant, A bridging model for multi-core computing, J. Comput. Syst. Sci. 77 (2011) 154–166. https://doi.org/10.1016/j.jcss.2010.06.012.
- [42] B. Tu, J. Fan, J. Zhan, X. Zhao, Performance analysis and optimization of MPI collective operations on multi-core clusters, J. Supercomput. 60 (2012) 141–162. https://doi.org/10.1007/s11227-009-0296-3.