



UNIVERSITÀ DI PISA



Scuola Superiore  
Sant'Anna

di Studi Universitari e di Perfezionamento

UNIVERSITÀ DI PISA AND SCUOLA SUPERIORE SANT'ANNA

DEPARTMENT OF COMPUTER SCIENCE

Master's Degree in Computer Science and Networking

---

# A Bulk Synchronous Parallel pattern in FastFlow

---

*Candidate:*  
Orlando Leombruni

*Supervisor:*  
Prof. Marco Danelutto

---

Academic Year 2018/2019

# Ringraziamenti

Si chiude qui un altro importante capitolo della mia vita. Il percorso di questa Laurea Magistrale è stato lungo e a volte difficile, ma non sono pentito delle mie scelte: ho imparato tanto, sia a livello di conoscenze che a livello umano. Ho avuto al mio fianco persone con cui ho condiviso gioie, soddisfazioni, piccoli e grandi traguardi, e perché no anche un po' di tristezza<sup>1</sup>. Purtroppo la situazione attuale non mi permette di festeggiare con queste persone — che con ogni probabilità sono le stesse che stanno leggendo ora questa tesi — il raggiungimento di un risultato così importante: spero che queste poche righe di ringraziamento possano comunque esprimere un minimo la mia stima ed affetto verso di loro.

Un sentito ringraziamento al Prof. Marco Danelutto, per la pazienza ed il supporto dimostrato durante il lavoro di tesi.

Un grazie gigantesco a mamma Mara, di cui mi basta sentire la voce per riacquistare tranquillità. Grazie di cuore *madreh*, mi sei sempre stata vicina anche a 500 chilometri di distanza e non mi hai mai fatto mancare nulla. (E mi regali un sacco di biglietti di concerti, cosa molto importante.)

Subito dopo — ma non per importanza, ovviamente — vorrei ringraziare Alessandra, che mi accompagna tenendomi la mano praticamente dall'inizio di questa avventura universitaria. Si dice che al fianco di ogni grande uomo vi è una grande donna: io non sono di certo un grande uomo (beh, stazza a parte), ma sono sicuro che al mio fianco c'è una donna superlativa.

Un grazie enorme a nonna Maria, che non ha mai mancato di farmi sentire il suo affetto e il suo supporto. Abbiamo superato tante difficoltà assieme, abbiamo chiacchierato approssimativamente cinquemila ore al telefono da quando sono a Pisa, e spero di poter continuare a farlo ancora per molto.

Un grazie a *fratello* Maurizio, a cui ogni tanto tocca sorbirsi le mie lamentele videoludiche e non solo.

Un grazie di cuore a Valerio e alla sua famiglia per esserci sempre, anche se ci sentiamo poco. (Un po' per colpa mia, un po' perché non so mai qual è il tuo numero di telefono, caro Zione.)

Un ringraziamento particolare va a Francesco, Maria e Marco, che mi hanno sempre riservato un'accoglienza e un calore senza pari. Sono contento che le nostre strade si siano incrociate.

Un grazie speciale ad Antonio e Mattia, amici di lunghissima data su cui so di poter contare nonostante la lontananza. (Ehi, ora posso prendervi in giro in quanto extracomunitari! Bella 'sta Brexit.)

Un grazie ad Alfredo, ~~Lamberto~~ Luigi e (l'altro) Mattia, per le serate spassose e per le chat deliranti di cui non posso scrivere il nome, perché Cricca Panzoni non è un nome adatto per essere stampato su una tesi. (Ops.)

Un grazie a Marco (quello *tetesko*), perché anche se ci sentiamo poco la nostra amicizia resta sempre salda come ai vecchi tempi.

Un grazie ai "Mammut di Proust" Francesco, Lorenzo e Matteo per la compagnia e i meme sull'informatica (ma non solo).

Un ringraziamento speciale ad Alberto e Paola per la professionalità, la cortesia e la pazienza dimostratami.

Un grazie a tutti quelli che hanno condiviso con me un pezzetto del loro percorso di vita, sia pure solo virtualmente. Vorrei ringraziarvi ad uno ad uno, ma lo spazio purtroppo è poco. Confido nella vostra comprensione.

Vorrei infine dedicare questo traguardo a nonno Vincenzo, che sono sicuro sarà fiero di me ovunque egli sia. Ciao nonno, mi manchi.

Orlando Leombruni

---

<sup>1</sup>Che comunque non va vissuta come un valore negativo, come insegnano gli Elit.

# Contents

<b>1</b>	<b>Introduction</b>	<b>7</b>
<b>2</b>	<b>Background</b>	<b>11</b>
2.1	The Bulk Synchronous Parallel model . . . . .	11
2.1.1	The BSP Computer . . . . .	12
2.1.2	Towards an object-oriented shared memory BSP . . . . .	15
2.2	FastFlow . . . . .	20
2.2.1	Slab Allocator . . . . .	22
<b>3</b>	<b>Architectural design</b>	<b>24</b>
3.1	Overview . . . . .	24
3.2	Nodes . . . . .	25
3.3	Memory manager . . . . .	25
3.4	Synchronization facility . . . . .	26
3.5	Shared memory . . . . .	27
3.6	Summary . . . . .	29
<b>4</b>	<b>Implementation</b>	<b>30</b>
4.1	Library structure . . . . .	30
4.2	Architecture . . . . .	31
4.3	bsp_program.hpp . . . . .	31
4.4	bsp_node.hpp . . . . .	34
4.5	bsp_internals.hpp . . . . .	35
4.5.1	Requests . . . . .	35
4.5.2	Shared memory containers . . . . .	36
4.6	bsp_communicator.hpp . . . . .	38
4.6.1	Implementation of the Shared Memory entity . . . . .	38
4.6.2	Implementation of the Memory Manager entity . . . . .	40
4.7	bsp_variable.hpp and bsp_array.hpp . . . . .	43
4.8	bsp_barrier.hpp . . . . .	45
4.9	stl_allocator.hpp . . . . .	47
<b>5</b>	<b>Experimental validation</b>	<b>48</b>
5.1	Performance metrics . . . . .	48
5.2	Machine architecture . . . . .	50
5.3	Benchmarks . . . . .	50
5.3.1	Vector inner product . . . . .	52
5.3.2	Parallel sorting by regular sampling . . . . .	54

5.3.3	Fast Fourier Transform . . . . .	60
5.3.4	Matrix LU decomposition . . . . .	63
5.3.5	Other programs . . . . .	69
5.4	Summary . . . . .	69
<b>6</b>	<b>Conclusion</b> . . . . .	<b>70</b>
6.1	Future work . . . . .	72
<b>A</b>	<b>Library quick start</b> . . . . .	<b>74</b>
A.1	API documentation . . . . .	74
A.1.1	bsp_program Class Reference . . . . .	74
A.1.2	bsp_node Class Reference . . . . .	75
A.1.3	bsp_variable<T> Template Class Reference . . . . .	79
A.1.4	bsp_array<T> Template Class Reference . . . . .	81
A.2	Variable/array request mechanism . . . . .	88
A.3	A note on BSPunsafe methods . . . . .	90
<b>B</b>	<b>Source code</b> . . . . .	<b>91</b>
B.1	BSP Library for FastFlow . . . . .	91
bsp_array.hpp	. . . . .	91
bsp_barrier.hpp	. . . . .	96
bsp_communicator.hpp	. . . . .	97
bsp_internals.hpp	. . . . .	104
bsp_node.hpp	. . . . .	109
bsp_program.hpp	. . . . .	112
bsp_variable.hpp	. . . . .	115
stl_allocator.hpp	. . . . .	118
B.2	Parallel test programs . . . . .	119
BSPinprod.cpp	. . . . .	119
BSPpsrs.cpp	. . . . .	121
BSPfft.cpp	. . . . .	125
BSPlu.cpp	. . . . .	131
B.3	Sequential test programs . . . . .	136
sequential_inprod.cpp	. . . . .	136
sequential_sorting.cpp	. . . . .	137
sequential_fft.cpp	. . . . .	138
sequential_lu.cpp	. . . . .	142
B.4	Miscellaneous . . . . .	142
commstress.cpp	. . . . .	142
BSPbench.cpp	. . . . .	145
ff_example.cpp	. . . . .	148
mwe.cpp	. . . . .	150
unit_tests.cpp	. . . . .	150
B.5	Java programs . . . . .	162
BSPpsrs.java	. . . . .	162
IntArrayList.java	. . . . .	166
SequentialInprod.java	. . . . .	166
SequentialSort.java	. . . . .	167
SequentialFFT.java	. . . . .	168
SequentialLU_commons.java	. . . . .	172

# List of Figures

1.1	Workflow for the design and implementation of the library. . . .	9
2.1	The architecture of a generic BSPC. . . . .	12
2.2	The superstep flow of a BSP computation. . . . .	13
2.3	The BSPRAM architecture. . . . .	17
2.4	The superstep flow of a BSPRAM computation. . . . .	17
2.5	The FastFlow framework layered design (from [2]). . . . .	20
2.6	An example of arbitrary compositions of FastFlow computational graphs (from [21]). . . . .	21
2.7	The slab allocation memory organization. . . . .	22
3.1	The BSPC architecture presented in this work. . . . .	25
3.2	The sub-components of a BSP node. . . . .	26
3.3	The Memory Manager component. . . . .	27
3.4	The Shared Memory entity. . . . .	28
4.1	The FastFlow graph of the implementation. . . . .	32
4.2	Logical tasks performed by a runnable <code>bsp_node</code> object. . . . .	34
5.1	An example of cyclic distribution of an array over 3 processors. . . . .	52
5.2	Average completion time for the vector inner product benchmark on 2B elements. . . . .	53
5.3	Speedup for the vector inner product benchmark. . . . .	53
5.4	Scalability for the vector inner product benchmark. . . . .	54
5.5	Efficiency for the vector inner product benchmark. . . . .	55
5.6	An example of block distribution of an array over 3 processors. . . . .	55
5.7	Average completion time for the PSRS benchmark on $2^{24}$ elements. . . . .	56
5.8	Speedup for the PSRS benchmark. . . . .	57
5.9	Scalability for the PSRS benchmark. . . . .	58
5.10	Efficiency for the PSRS benchmark. . . . .	58
5.11	The group-cyclic distribution with $n = 16$ , $p = 4$ and $c = 2$ . . . . .	60
5.12	Average completion time for the FFT benchmark on $2^{25}$ elements. . . . .	61
5.13	Speedup for the FFT benchmark. . . . .	62
5.14	Scalability for the FFT benchmark. . . . .	62
5.15	Efficiency for the FFT benchmark. . . . .	63
5.16	Cyclic distribution of a $5 \times 5$ matrix over $2 \times 3$ processors. . . . .	64
5.17	A two-phase broadcast of a vector of twelve elements over four processors (taken from [9]). . . . .	65

5.18	Average completion time for the LU decomposition of a $5000 \times 5000$ matrix. . . . .	66
5.19	Speedup for the matrix LU decomposition benchmark. . . . .	67
5.20	Scalability for the matrix LU decomposition benchmark. . . . .	68
5.21	Efficiency for the matrix LU decomposition benchmark. . . . .	68
5.22	BSP parameters for the Phi machine as reported by bspbench. . . . .	69
6.1	FastFlow graph of a possible implementation of the MultiBSP model. The solid line nodes compose the MultiBSP tree. . . . .	72
A.1	An example of the difference between bsp_variables and their underlying containers. . . . .	90

# List of Listings

4.1	The <code>bsp_program</code> constructor. . . . .	32
4.2	The emitter's <code>svc</code> method. . . . .	33
4.3	The collector's <code>svc</code> method. . . . .	33
4.4	The <code>start</code> method of the <code>bsp_program</code> class. . . . .	34
4.5	The <code>bsp_sync</code> method. . . . .	35
4.6	Possible request types. . . . .	35
4.7	Fields of a request data structure. . . . .	36
4.8	The <code>inner_var</code> class. . . . .	37
4.9	The <code>inner_array</code> class. . . . .	37
4.10	Fields of the <code>bsp_communicator</code> class used to implement the Shared Memory logical entity. . . . .	39
4.11	The hashing function for variables and arrays. . . . .	39
4.12	The <code>get_variable</code> method. . . . .	41
4.13	Fields of the <code>bsp_communicator</code> class used to implement the Mem- ory Manager logical entity. . . . .	41
4.14	The variable put request handler of the <code>bsp_communicator</code> . . . .	42
4.15	Request processing. . . . .	42
4.16	The <code>bsp_communicator</code> final cleanup function. . . . .	43
4.17	The <code>bsp_variable</code> class. . . . .	44
4.18	The <code>bsp_barrier</code> class. . . . .	46
4.19	The FastFlow allocator "STL interface". . . . .	47
A.1	The code for the example. . . . .	89

# Chapter 1

## Introduction

Writing parallel software, even in its simplest form, is harder than writing sequential software. The amount of details that need to be taken into account is much higher, and traditional programming languages are for the most part designed for sequential programs. In these languages, support for parallel computations is not a central focus and is often limited to a set of basic mechanisms for spawning and synchronizing parallel executors.

More recently, however, parallel computing has started to become a central element of the technological landscape. Commodity processors feature lower operating frequencies and an increasing number of cores with respect to the ones available ten years ago [52]. GPUs, which feature a high number of computing units that can only perform basic arithmetic operations, are increasingly used as coprocessors to accelerate heavy numerical algorithms [15, 44]. Smart devices run on low-power multicore SoCs (often equipped with small GPUs) that greatly benefit from a balanced computational load [58]. The ability to write good parallel code is therefore becoming a basic tool in a programmer's arsenal.

We believe that there are two fundamental ingredients in efficient parallel programs: a **clear design** of the computation to be performed and a **suitable and efficient framework** for its implementation. The role of the second ingredient is often played by libraries and APIs paired with existing sequential languages. One such example is *FastFlow*, a C++ library developed at University of Pisa and University of Torino [3], which provides the programmer with easy-to-use mechanisms for the design and implementation of parallel programs.

Designing the parallel algorithm is not an easy task either. One can start assuming virtually no limitations on what a parallel computer can do, i.e. to have infinite processing elements and infinite memory without concerns to regulations of concurrent accesses. The algorithms designed under these assumptions, however, cannot be directly implemented “as-is”, since the actual hardware *does have* limitations.

The Bulk Synchronous Parallel paradigm was introduced by Leslie G. Valiant in the late 1980s as a model with realistic limitations on parallel hardware features [61]. BSP algorithms are divided in *supersteps*: every processing element must wait that its peers complete the current superstep before advancing the computation to the next one. Moreover, while communications with other pro-



cessing elements can be scheduled at any time, they actually only take place after the end of the superstep. Together, these two constraints impose a “superstructure” to parallel computations, which makes it easier to reason about the effects of operations.

The BSP model has not been widely adopted, at least in comparison to other parallel models like PRAM and dataflow, until recent times. For a long while its use was limited to resolution of numerical problems and niche interests in theoretical computer science. Since the early 2010s, however, the BSP model got a “second life” having been adopted in Google Pregel [36], a framework for distributed large-scale graph computing, due to its clear and effective programming style and performance model. Nowadays, it is implemented in a number of frameworks for both distributed (Apache Hama [7], Apache Giraph [6]) and parallel (MulticoreBSP [67, 66], Bulk [16]) computing.

We believe that BSP is a good model for designing efficient parallel algorithms. Its role as a “bridging model” also allows for direct execution of those algorithms on a suitable runtime support, therefore freeing the programmer from having to adapt them to specific frameworks’ programming models. Unfortunately, there are few BSP implementations that offer modern features like automatic memory management on multicore, shared memory architectures. Therefore, we considered implementing a BSP library that offers the above-mentioned modern features, with an API that is simple to use and a run-time support that frees users from menial tasks like having to register shared variables. We chose FastFlow as the underlying framework due to its easy-to-use algorithmic skeletons and its compositional capabilities. Since the BSP pattern is not currently provided by FastFlow, we also decided to make the library a fully-featured FastFlow component, able to be used inside a broader computation.

The objective of this work therefore consists in the design and implementation of a modern C++ parallel programming library providing the BSP model *on top of* and *within* the FastFlow framework. The library will target shared-memory multicore machines. We want to provide programmers with an easy-to-use interface that allows them to design and code efficient BSP programs. Our implementation will fully support the object-oriented paradigm and will relieve the programmer from the burden of having to manually manage memory and shared variables. The library has been designed to be

- fully **BSP compliant** — existing BSP algorithms must be able to be easily ported to our library, and
- fully **FastFlow compliant** — the whole BSP computation will act as a single FastFlow parallel component, thus being able to become part of a broader FastFlow structured computation; it will receive input data from other FastFlow nodes and send output data to other FastFlow nodes.

The library will be itself implemented as a FastFlow graph, therefore exploiting the framework’s compositional capabilities. This fact has helped in the implementation design phase, as solutions based on different algorithmic skeletons for the library could be considered and prototyped quickly without having to change the internal structure of the nodes.

Our work started with an in-depth study of the state of the art, in the form of both theoretical advancements and practical implementations, to better un-

derstand the BSP model and its scope. After coming up with an initial draft, we carried on both the architectural and implementation design at the same time, one aspect influencing the other. Of course, this meant going through multiple revisions of the library, from a rough prototype – which aggressively enforced type correctness between a superstep’s output and its successor’s input – up to the final version presented in this thesis (Figure 1.1), sporting good scalability as well as highly expressive API and a run-time support that relieves the user from burdens like memory management and variable registration. We took inspiration for some of the features from “competitors”, mainly MulticoreBSP for Java, but ultimately we designed both the implementation structure and the API from scratch. Lastly, we performed extensive tests for both correctness and performance of our implementation, achieving good results for the latter.

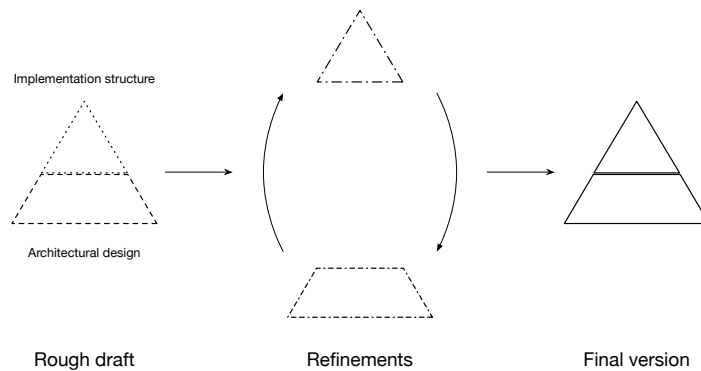


Figure 1.1: Workflow for the design and implementation of the library.

The rest of the thesis is organized as follows.

- **Chapter 2** provides the reader with the necessary concepts used throughout the work. The theory behind the Bulk Synchronous Parallel model is laid out, together with the relevant definitions and theorems that help frame the model into the bigger picture of parallel programming. The FastFlow framework is also introduced, together with its main features, to provide familiarity with some of the terminology and mechanisms presented in the rest of the thesis. A brief digression on a memory allocation strategy provided by FastFlow and used in the library closes the Chapter.
- **Chapter 3** introduces the architectural model of our implementation. This model is composed by various entities: the design of each entity is discussed, along with a description of how it fits and behaves inside the architecture.
- **Chapter 4** presents an overview of the library’s implementation. Each class is described in depth, focusing on its peculiarities and relation with respect to the entities of Chapter 3. Extensive code snippets detailing the main portions of the implementation are also provided and discussed.
- **Chapter 5** contains the experimental results from the various test programs we implemented. The testing methodology and machine architecture are detailed at the start of the Chapter. For each test program we define the communication scheme and input distribution, then we provide

and discuss performance plots confronting the results obtained from our library with the corresponding MulticoreBSP for Java implementation.

- **Chapter 6** draws the conclusions relative to the whole work. The experimental outcomes are once again commented, and an outline for possible future works is provided.
- **Appendix A** contains the documentation for the library API, together with some clarifications about advanced mechanisms such as direct access methods.
- **Appendix B** contains the source code for the library and for the test programs. The source code is also available at [41].

## Chapter 2

# Background

This Chapter provides the reader with basic concepts relative the Bulk Synchronous Parallel (BSP) paradigm. The BSP model has been designed in the late 1980s to overcome the limitations of the PRAM model. BSP divides a program in a series of supersteps that contain both computation performed by a number of processing elements and communication between them. Each superstep is separated from the others by means of a global synchronization (e.g. barrier). Recent developments on the paradigm (e.g. MultiBSP, Multi-coreBSP) are touched upon.

An overview of the FastFlow framework closes the Chapter.

### 2.1 The Bulk Synchronous Parallel model

The Bulk Synchronous Parallel model was introduced by Leslie G. Valiant in 1990 [61] as a “bridging model” for distributed computing. Valiant considered the main reason for the “success” and ubiquity of the sequential computation the availability of a *central unifying model*, the Von Neumann computer. This unified view serves as a connecting bridge between software and hardware: software developers can focus on writing programs that run efficiently on it – abstracting from the complexities of the hardware. Conversely, hardware designers only have to realize efficient Von Neumann machines, without having to mind about the software that will be executed on them. Valiant thought that – in order to obtain a widespread adoption of distributed computing – a model with a similar purposes as Von Neumann’s one for sequential computation was needed; he then introduced the BSP model as a viable candidate for this role.

NOTE. Valiant’s original article actually references *parallel computing*, hence the name *Bulk Synchronous Parallel*; nevertheless, it was written at a time when the predominant source of parallelism derived from the exploitation of an interconnection of processing elements, each with its own private memory and lacking a global clock or synchronization, in which data is exchanged by messages sent and received via the network — i.e. what is called today a “distributed architecture” [50]. (In fact, Valiant even mentions implementations of a BSP computer based on optical crossbars or packet-switched networks, reinforcing this concept.) Throughout the remainder of this text we will refer to

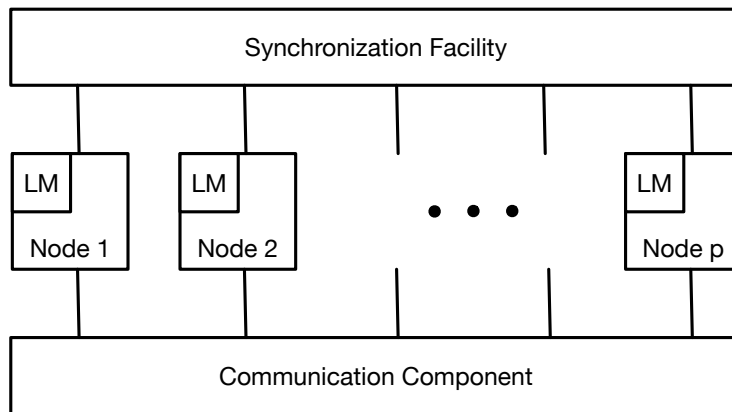


Figure 2.1: The architecture of a generic BSPC.

the modern-day definition of *parallel computing* (i.e. a single system with multiple processors and cores, with a shared memory and global clock), so Valiant's original BSP article will be treated as referencing a distributed architecture.

### 2.1.1 The BSP Computer

The BSP model defines – along the lines of the Von Neumann model – an abstract **Bulk Synchronous Parallel Computer**, or BSPC for short. A BSPC is the combination of three attributes:

1. a number of *nodes* for performing processing and memory operations;
2. a *communication component* that delivers point-to-point messages between nodes;
3. a *synchronization facility* that coordinates the behavior of the nodes.

A BSP computation consists in a sequence of **supersteps**. During a superstep, each node performs a task consisting of local computation and message sending primitives. The superstep ends when all nodes finish their task. The synchronization facility is the entity responsible for checking this condition. An important feature of this mechanism is that the effect of any communication primitive performed during a superstep will take place only during the successive superstep, i.e. the communication between nodes actually takes place at the end of the superstep.

The way the synchronization facility works has an impact on the performance of a BSPC. First, we define the following.

**Definition 2.1.** A *time unit* is the time spent by a node performing a single operation over data available in its local memory.

**Definition 2.2** (Periodicity of a BSPC). The *periodicity*  $L$  of a BSPC is the number of time units in the interval between two consecutive checks of the “superstep ended” condition.

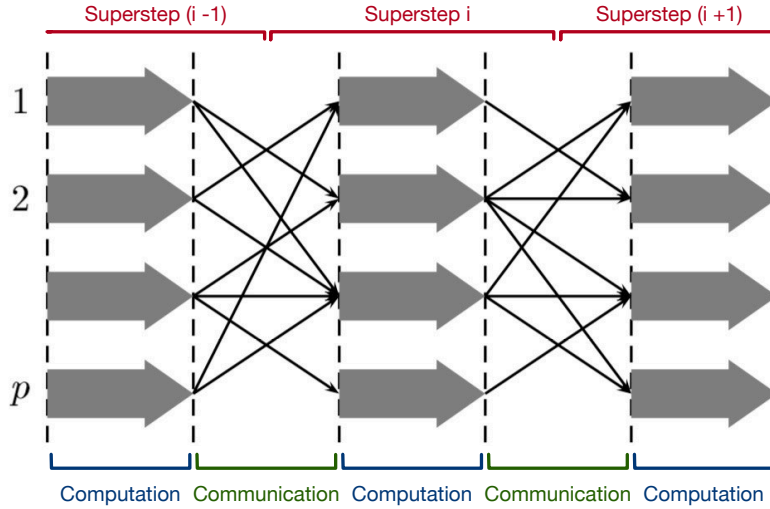


Figure 2.2: The superstep flow of a BSP computation.

In other words, the synchronization facility checks that the current superstep has ended every  $L$  time units. This definition is purposefully vague, since no assumption is made on how the synchronization facility actually works. For example, if the system can continuously check whether a superstep is completed, then  $L$  can be defined as the number of time units needed to perform the check itself. The periodicity concept also brings two interesting consequences:

- $L$  is the minimum effective duration of a superstep (i.e. a superstep that theoretically completes in less than  $L$  time units will actually behave like a superstep that completes in  $L$  time units);
- the maximum efficiency for a BSPC of periodicity  $L$  can be achieved with a program that perfectly balances tasks in such a way that, in every superstep, every node completes its task in exactly  $L$  time units.

Depending on the synchronization facility, some portion of the periodicity  $L$  can be overlapped with local computations. In this case, with  $l$  we refer to the non-overlapping time interval of the periodicity.

The task executed by each node in a given superstep can be further divided into three distinct portions: the *input phase*, the *local computation phase* and the *output phase*. In the input phase, a node receives data in its local memory; then, some local computation is performed over this data; in the output phase, the node sends some data to other nodes. This functional partition of a task allows us to define more useful parameters for a BSPC.

**Definition 2.3** (Local computation cost for a superstep). Let  $i$  be a superstep. The *local computation cost*  $w_i$  is the maximum number of local operations performed by any node during superstep  $i$ . In other words, let  $w_{ij}$  be the number of local operations performed by node  $j$  during superstep  $i$ . Then we define

$$w_i = \max_j w_{ij}. \quad (2.1)$$

NOTE. Since we defined the time unit to be the time to perform a single operation it also holds that, for superstep  $i$ ,  $w_i$  is the time spent in the local computation phase.

**Definition 2.4** (Number of data units sent and received). Let  $h'_{ij}$  and  $h''_{ij}$  be respectively the number of data units received and sent by node  $j$  during superstep  $i$ . We define the *maximum number of received data units* during  $i$  as

$$h'_i = \max_j h'_{ij} \quad (2.2)$$

and the *maximum number of sent data units* during  $i$  as

$$h''_i = \max_j h''_{ij}. \quad (2.3)$$

**Definition 2.5** ( $h$ -relation). Let  $i$  be a superstep. The *total data units* exchanged during this superstep by a generic node  $j$  is defined as

$$h_{ij} = h'_{ij} + h''_{ij}. \quad (2.4)$$

The maximum data units exchanged during superstep  $i$  is called the  *$h$ -relation* of  $i$  and is defined as

$$h_i = \max_j h_{ij}. \quad (2.5)$$

Sometimes it is also useful to define the *maximum  $h$ -relation* for a BSP program as

$$h = \max_i h_i. \quad (2.6)$$

We said before that the communication component is the entity responsible for delivering point-to-point messages (carrying data) between nodes. Without loss of generality, we can assume that this component will take a fixed amount of time to delivery a single data unit after reaching steady state.

**Definition 2.6.** Let  $s$  be the startup time of the communication component, i.e. the time needed to reach steady state, and  $\bar{g}$  be its basic throughput when in continuous use, i.e. the time spent in delivering a single unit of data at steady state. Then, for superstep  $i$ , the communication component will take  $\bar{g}h_i + s$  units of time.

From this definition it follows that in order to achieve optimality we need  $\bar{g}h_i$  to be at least of the same order of magnitude as  $s$ , i.e.

$$\bar{g}h_i \geq s. \quad (2.7)$$

This condition is required, otherwise a lot of time is “wasted” in the startup phase of the communication component. This suggests a way to simplify the cost model.

**Definition 2.7** (Gap of a BSPC). Let  $g = 2\bar{g}$ . Then, it holds

$$gh_i = 2\bar{g}h_i \geq \bar{g}h_i + s \quad (2.8)$$

where the second inequality comes from equation 2.7. The parameter  $g$  is called *gap* or *communication throughput ratio* of a BSPC.

We now have all the necessary ingredients for a basilar cost model for a BSPC.

**Definition 2.8** (Cost model for a superstep). Let  $i$  be a superstep. Then, the cost for  $i$  (i.e. time spent in this superstep) can be modeled by the following equation:

$$c_i = w_i + gh_i + l. \quad (2.9)$$

**Definition 2.9** (Cost model for a BSP program). Let  $S$  be the number of supersteps of a BSP program. Then, the total time spent executing the program can be modeled by the following equation:

$$c = \sum_{i=1}^S c_i = \sum_{i=1}^S (w_i + gh_i + l) = \sum_{i=1}^S w_i + g \sum_{i=1}^S h_i + S \cdot l. \quad (2.10)$$

If we define

$$W = \sum_{i=1}^S w_i \quad H = \sum_{i=1}^S h_i$$

then equation 2.10 can be simplified as

$$c = W + g \cdot H + S \cdot l \quad (2.11)$$

where  $W$  is called the *total local computation cost*,  $H$  the *total communication cost* and  $S$  the *total synchronization cost*.

Lastly, we define when a BSP program is *balanced*.

**Definition 2.10** (Balanced BSP computation). Let  $p$  be the number of nodes of a BSPC. We define the *local computation volume*  $\mathcal{W}$  as the total number of local operations executed during the whole program. Similarly, we define the *communication volume*  $\mathcal{H}$  as the total number of data units transferred between all nodes during the whole program. If both following conditions hold:

$$W = \mathcal{O}(\mathcal{W}/p) \quad H = \mathcal{O}(\mathcal{H}/p) \quad (2.12)$$

then the BSP program is said to be *balanced*.

### 2.1.2 Towards an object-oriented shared memory BSP

The aim of this work is to develop a BSP implementation which focuses on *objects* as the atomic unit of data, designed for multicore shared-memory single machines. This scenario introduces its own set of complexities and challenges that have been tackled throughout the years by different people. This section presents the main concepts behind the most relevant adaptations of the BSP model for a shared-memory architecture.

#### PRAM simulations and the BSPRAM model

In Valiant's original article [61], particular emphasis was placed in providing an efficient way to simulate PRAM algorithms on a BSPC. The PRAM model was introduced in [64] as a way to design parallel algorithms and perform quantitative analyses on their performance, in a similar manner to Von Neumann's RAM. The PRAM model relied on a set of simplifying assumptions, namely



- any processor can access any location of the shared memory uniformly (i.e. with the same cost of access)
- there is no resource contention between processors.

A PRAM is categorized according to the strategy it uses to resolve read/write conflicts in memory:

- if a memory location can be accessed only by a processor at a time, regardless if it is a read or a write, then the PRAM is said to be *EREW* (Exclusive Read, Exclusive Write);
- if a single processor can write in a memory location at a given time, but multiple processors can read from it simultaneously, then the PRAM is said to be *CREW* (Concurrent Read, Exclusive Write);
- if multiple processors can read and write in a memory location at a time, then the PRAM is said to be *CRCW* (Concurrent Read, Concurrent Write).<sup>1</sup>

In [59] it is shown that PRAM algorithms can be efficiently simulated on a BSPC<sup>2</sup>; the single shared memory of a PRAM is implemented in a BSPC by mapping its memory positions to local memories of BSP nodes according to a hashing function. The simulation is *optimal* provided that the PRAM has more processors than the BSPC.

**Definition 2.11.** A model  $M$  can *optimally simulate* a model  $N$  if there exist a transformation that maps any problem with cost  $T(n)$  on  $N$  to a problem with cost  $O(T(n))$  on  $M$ .

**Definition 2.12.** Let  $p$  be the number of nodes of a BSPC. A PRAM algorithm is said to have *slackness*  $\sigma$  if at least  $\sigma p$  PRAM processors perform a memory operation (read or write) at every step.

A PRAM algorithm with slackness  $\sigma$  has *at least*  $\sigma p$  processors that communicate at each step. A simulation of this PRAM algorithm on a BSPC with  $p$  nodes means that at least  $\sigma$  processors are mapped to each BSP node. In order to achieve an optimal simulation, some constraints must be put on  $\sigma$ , as showed by the following theorem.

**Theorem 2.1.** *Suppose to have a BSPC with  $p$  nodes, with  $g$  constant (not depending on  $p$  or the problem size, i.e.  $g = O(1)$ ) and  $l = O(\sigma)$ . This BSPC can optimally simulate*

- any *EREW* PRAM algorithm with slackness  $\sigma \geq \log p$
- any *CRCW* PRAM algorithm with slackness  $\sigma \geq p^\epsilon$ , with  $\epsilon > 0$ .

The interested reader can find the proof of this theorem in [59].

In [55] Tiskin proposed a BSP model for shared memory systems, which replaced the communication network with a shared memory component; this allowed for the exploiting of data locality, something that was not possible with the original BSP model. Tiskin called his model *BSPRAM*; the architecture for a BSPRAM is shown in Figure 2.3.

<sup>1</sup>The Exclusive Read, Concurrent Write (ERCW) strategy is never considered.

<sup>2</sup>The paper shows a PRAM simulation on a XPRAM, which is a simple BSPC.

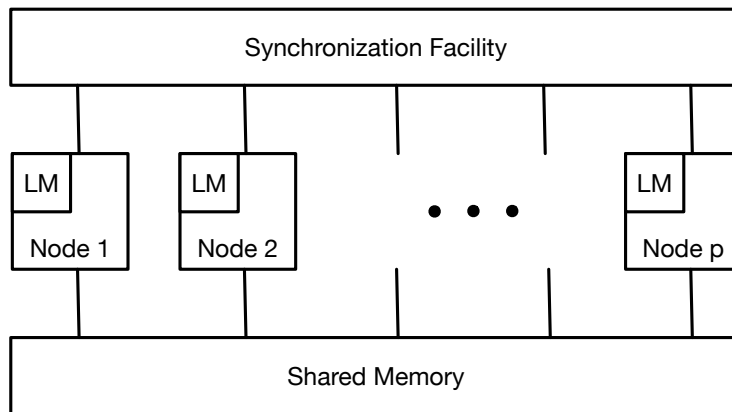


Figure 2.3: The BSPRAM architecture.

As with a standard BSPC, a BSPRAM has  $p$  processors (nodes) that perform operations on data in their local memory, a synchronization facility and a communication component, which in this case is a shared memory unit. By adapting slightly the definition of  $h$ -relation as the maximum data units *read or written into the shared memory*, the same basic cost model of equation 2.9 can be used. The BSPRAM supersteps differ a little from standard BSP supersteps: in the latter model, the input phase was mostly a “passive” one, since the communication component wrote data directly into the nodes’ local memories; communication could be essentially considered concluded when the output phase ended. In the BSPRAM, instead, nodes must actively read from the shared memory in the input phase, so the three phases of a superstep are more distinct (Figure 2.4).

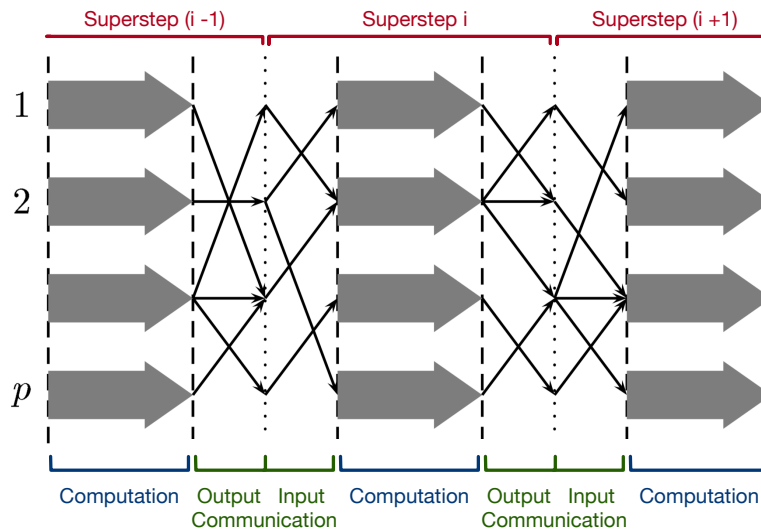


Figure 2.4: The superstep flow of a BSPRAM computation.

As Valiant did for PRAM algorithms, Tiskin proved that BSPRAM programs could be optimally simulated onto a “classical” BSPC. First, he defined the *slackness* parameter for BSPCs.

**Definition 2.13.** A BSP or BSPRAM algorithm is said to have *slackness*  $\sigma$  if for any superstep  $i$  it holds

$$gh_i \geq \sigma \quad (2.13)$$

i.e. the communication cost for any superstep is at least  $\sigma$ .

The following theorem directly mirrors Theorem 2.1.

**Theorem 2.2.** *An optimal simulation on a BSPC with parameters  $p, g, l$  can be achieved for*

- *any algorithm designed for a BSPRAM with EREW strategy and parameters  $p, g, l$ , that has slackness  $\sigma \geq \log p$*
- *any algorithm designed for a BSPRAM with CRCW strategy and parameters  $p, g, l$ , that has slackness  $\sigma \geq p^\epsilon$ , with  $\epsilon > 0$ .*

Another important property of a BSPRAM is that it can optimally simulate generic BSP algorithms that have at least a certain level of slackness.

**Theorem 2.3.** *An optimal simulation on an EREW BSPRAM with parameters  $p, g, l$  can be achieved for any algorithm designed for a BSPC with parameters  $p, g, l$  that has slackness  $\sigma \geq p$ .*

Proof for Theorem 2.2 follows from the one for Theorem 2.1, while the interested reader can find proof for Theorem 2.3 in [55].

Tiskin’s results are interesting, since they show that classical BSP algorithms can be efficiently run on shared-memory systems if the amount of data exchanged at every superstep (the  $h$ -relation) is at least equal to the number of BSP nodes, which is an easily achievable condition in most cases.

### Object-oriented BSP and further developments

Most BSP models presented so far do not place any condition on how the memory is organized, nor on the structure and contents of messages exchanged by nodes. (An example of BSP model that explicitly considers memory hierarchies can be found at [22].) Most BSP implementations — especially the ones based on BSPlib [14, 12], but also including other independent implementations [42, 65] — require that the programmer directly manages the memory, allocating and freeing blocks of data “by hand” [27]. Nowadays, the most popular programming languages tend instead to reduce the burden on the programmer by automatically taking care of memory management. We think that our BSP implementation should adapt this philosophy by providing the user with a simple memory management scheme. We choose to embrace the *object-oriented* paradigm [51] and place *objects* as the atomic unit of both BSP messages (i.e. each message carries a single object) and memory management.

The idea of an object-oriented extension to the BSP model is not a new one. In [31] Lecomber proposed a distributed memory object-oriented BSPC and BSP++, its C++ implementation. Some of the ideas proposed in this work derive from Lecomber’s solution, such an emphasis on easy-to-use mechanisms

and primitives, while others – for example, tightly coupling an object with the set of processes (nodes) allowed to use it, as a mean to express whether the object is in a node’s local memory – only make sense in a distributed environment.

There exist, of course, many implementations of the BSP model for modern languages, some of which support the object-oriented paradigm [25, 26, 33]. The library which bears the closest similarity to our work is, however, *Multi-coreBSP for Java* by Yzelman and Bisseling [39, 66]. This library targets shared memory, multicore machines and place *objects* as the atomic unit of data; it provides the user with few, clear-cut memory operations (mainly variants of *put* and *get*) and an interface similar to the language’s Collections framework. The definition of a BSP program in *Multi-coreBSP for Java* also follows the object-oriented paradigm: not only the whole computation is encapsulated into a class that inherits from `BSP_PROGRAM`, but also variables are accessed and manipulated via methods of the object that represents them. We will emphasize our solution’s similarity to *Multi-coreBSP for Java* in more detail in Chapter 4; here we highlight the fact that most programs written using the *Multi-coreBSP for Java* library are very easily ported to this work’s C++ library, once taken care of the differences between the two languages.

Over the course of the years, new refinements of the BSP model have been proposed. In [60], BSP’s original author Valiant proposed an extension of its model to account for multicore machines and memory hierarchies called *Multi-BSP*.

A *Multi-BSP* is organized in a tree structure where the leaves (level 0 nodes) are processors with no local memory; every node at tree level  $i$  contains a portion of level  $i - 1$  nodes and acts as a sub-BSP with some memory space  $m_i$  and internal synchronization, where internal level  $i - 1$  nodes act as the processing elements of the sub-BSP. For example

- a multicore processor is a level 1 node with cache memory  $m_1$  where each core is a level 0 node;
- a multiprocessor machine that uses the above processor is a level 2 node with shared random access memory  $m_2$ ;
- a LAN with a small number of such multiprocessor machines is a level 3 node, and so on.

One key element of the model is the ability to have *nested BSP runs*: *Multi-BSP* programs can have subroutines which are themselves *Multi-BSP* programs, and if the tree structure has enough depth the execution of those subroutines can be performed as a BSP computation over lower-level nodes. The *Multi-coreBSP for C* library [67], by the same authors of *Multi-coreBSP for Java*, is an implementation of the *Multi-BSP* model for shared-memory machines using the C language. Although a C++ interface is provided, the library does not focus on the same object-oriented principles as its Java counterpart and the two projects are mostly independent [38].

Lastly, we mention some of the most relevant distributed frameworks that are inspired or directly based on BSP. This model gained popularity in the recent years due to its inclusion in a highly influential proposal for a large-scale graph processing system, Google Pregel [36], in which graph computation follows the bulk-synchronous pattern and is organized in supersteps. Pregel is a

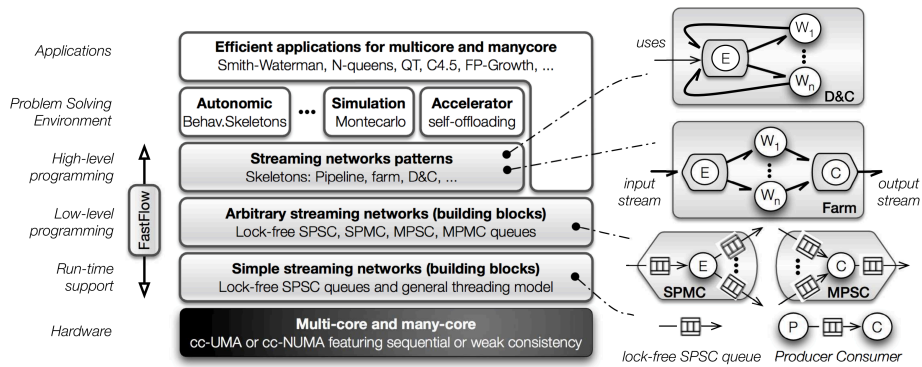


Figure 2.5: The FastFlow framework layered design (from [2]).

proprietary Google technology, but an open-source counterpart called Giraph has been developed at Apache [6, 45, 17]. A more general example of distributed BSP framework, also by Apache, is the Hadoop-based Hama [7, 47].

## 2.2 FastFlow

FastFlow is a parallel programming framework for shared-memory multicore machines, developed by the Parallel Processing research groups at University of Pisa and University of Torino [3]. It aims to be an easy-to-use yet highly efficient platform for the development of parallel applications targeting shared memory multi-/many-cores. It has been designed according to four fundamental principles to achieve those goals:

- a **layered design** that allows progressive abstraction from lower-level hardware concepts up to high-level programming constructs (Figure 2.5). The core of the framework is the *run-time support* layer, built on top of a threading library like POSIX. The run-time support layer features an extremely efficient implementation of one-to-one, Single-Producer Single-Consumer (SPSC) FIFO queues [57]. These queues are then used as building blocks for the above *low-level programming* layer, which provides support for other types of queues (one-to-many, many-to-one, many-to-many) used in creating lock-free and wait-free data-flow graphs. Lastly, the *high-level programming* level provides programmers with an interface for building parallel programs based on the *parallel patterns* concept. Users are not required to build their applications only on top of the high-level programming layer, but can instead use constructs from the full stack of layers;
- a **focus on base mechanism efficiency**; the SPSC queue structure that forms the base mechanism of FastFlow is wait-free and lock-free, with different optimizations for different underlying architectures. Those queues are used as synchronization mechanisms for pointers, in a producer/consumer pattern, and can be used to build networks of entities (threads) that communicate according to the data-flow pattern;

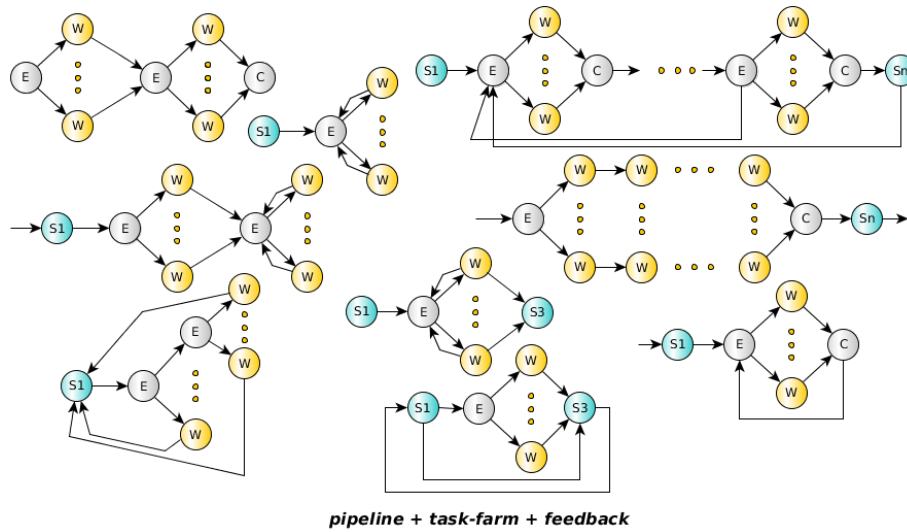


Figure 2.6: An example of arbitrary compositions of FastFlow computational graphs (from [21]).

- the adoption of **stream parallelism** as the main supported form of parallelism. This reduces the complexity of the library and is not restrictive, as other forms of parallelism can be implemented via stream parallel patterns. A stream program is a graph of independent stages that communicate via data channels; the streaming computation is a sequence of transformations on data streams done by nodes in the graph. The data streams can either be generated and consumed entirely inside the streaming application (*endo-streams*) or entirely outside (*exo-streams*);
- the **algorithmic skeletons** approach to parallel programming. The skeletons encompass common parallel programming concepts (e.g. Divide and Conquer, map-reduce, ...) and make them available to the user as high-level constructs with proper semantics.

FastFlow is designed to be portable over different architectures and to support different kinds of accelerators like GPUs and FPGAs. The authors have also started to work on a distributed-memory implementation of the library [4].

FastFlow is provided as a C++11 header library, and its mechanisms are available as template classes. The main component of the framework is the FastFlow *node*, an object of the `ff_node` class; in its simplest form, a `ff_node` contains the portion of sequential computation to be executed by a single runnable entity (thread). The parallel computation can be expressed as a graph of `ff_nodes` and, in true compositional fashion, can be considered a node itself, thus becoming a portion of a broader computation. According to this principle, the library provides algorithmic skeletons such as *pipeline* (`ff_pipe`) and *farm* (`ff_farm`) that are also `ff_node` objects. This supports the construction of graphs of any complexity to be built (Figure 2.6). In Chapter 4 we will see how to implement a BSPC using FastFlow components.

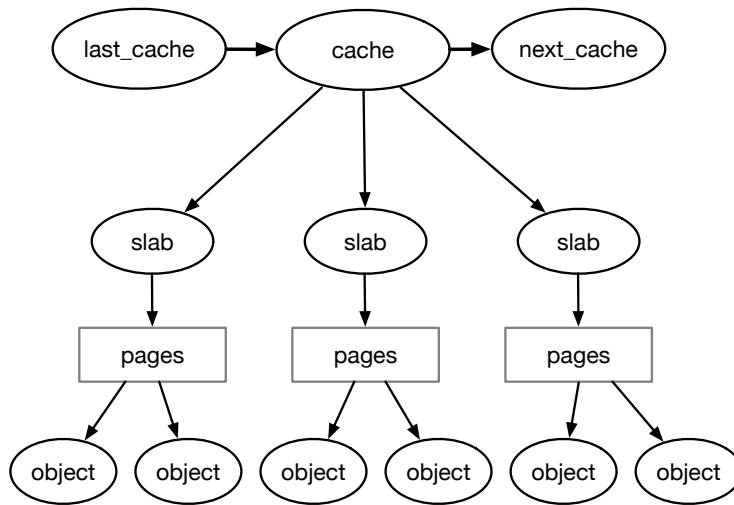


Figure 2.7: The slab allocation memory organization.

### 2.2.1 Slab Allocator

The FastFlow framework contains a custom allocator which may achieve higher performance when the allocation structure is composed of small memory areas [56]; this allocator is based on the *slab allocator* introduced by Bonwick in [13] and that will be briefly discussed here.

When dealing with frequent allocations and deallocations of small data objects, the CPU time for initialization and destruction of those objects may outweigh the proper allocation and deallocation cost, thus degrading the overall performance. A solution to this problem is *object caching*: when an object is released, the system keeps it in memory instead of releasing it; when another object with the same type is needed, the previous one is re-used and the cost of initialization is spared.

A memory managed with slab allocation is organized into multiple *caches* with no shared state and thus able to be locked independently. Each cache is composed by several *slabs* (Figure 2.7); a slab consists in one or more portions of contiguous memory divided into equal-sized chunks, with a reference count (of already-requested chunks) and a free list. When a chunk is requested, it is removed from the free list and the reference count is updated; the opposite happens when releasing it. A slab is the minimum memory size that a cache can grow or shrink; when a cache is full and a new object is requested, an entire slab is allocated and the object is built from a chunk of this new slab, and conversely, when the last chunk of a slab is released, the whole slab gets deallocated and removed from the cache.

The FastFlow implementation of the slab allocator is optimized for the following scenarios:

- a single thread allocates memory, and a single thread (not necessarily the same one) deallocates it;
- a single thread allocates memory, and multiple threads deallocate it.

FastFlow provides two interfaces for this allocator: the low-level one, `ff_allocator`, exclusively supports the allocation/deallocation patterns above, while the “high-level” one, `FFAllocator`, can be used by any thread regardless of the memory operation to be performed. A C++11 “wrapper” for the `FFAllocator` class is provided in the BSP library (see Section 4.9).



## Chapter 3

# Architectural design

After having introduced the necessary background topics in the previous Chapter, here we provide an overview on how the BSP library for FastFlow is logically organized and how the various tasks are assigned to different entities. One important thing to note is that the actual implementation of the library will not necessarily map these logical entities to “physical” ones, as discussed later in Chapter 4.

### 3.1 Overview

In the previous Chapter we discussed the three main components to the BSPC, the abstract machine that will execute BSP programs in an analogous way to how the abstract Random Access Machine executes sequential programs. These components are:

- the *nodes* that will provide computing capabilities and local memory;
- the *communication component* that will deliver point-to-point messages between any two nodes;
- the *synchronization facility* that will periodically check if all nodes exhausted their computation and communication tasks for the current superstep, allowing the global computation to advance to the next one.

These components form the basis of the structure of many BSP implementations, and the BSP library for FastFlow makes no exception. Since the library is designed for shared-memory systems, another logical component is needed:

- the *shared memory*, which will host the data needed by multiple BSP nodes.

Note that this BSPC is slightly different from the BSPRAM: in a BSPRAM, the shared memory effectively replaces the communication component, while in the BSPC presented here both entities coexist. Due to its specialized usage, the communication component for this BSPC is also called *memory manager* (MM).




Figure 3.1: The BSPC architecture presented in this work.

## 3.2 Nodes

BSP nodes are the entities responsible for requesting memory operations and performing computation over data that is already available in local memory. A BSP node can be thought of being composed by four sub-components:

- a *processing unit* (PU) that performs arithmetic and logical operations over local data;
- a *local memory unit* (LMU) that holds data available to be used exclusively by the PU of the same node. This memory cannot be accessed directly from other BSP nodes, but can be written to by the memory manager (see Sections 3.3, 3.5);
- a bidirectional *communication channel* (CC) with the memory manager. The node can perform read/write *requests* for shared memory locations over this channel, and conversely the memory manager will write data on the node's local memory through it;
- a bidirectional *synchronization channel* (SC) with the synchronization facility. The node will signal the end of the local computation phase of the current superstep through this channel, while the facility will signal the start of the next one.

Note that in this BSPC the shared memory component is not accessible directly from the nodes. This avoids memory contention in case of simultaneous accesses, which are instead mediated by the memory manager.

## 3.3 Memory manager

The memory manager (MM), as its name implies, is the entity that provides BSP nodes with access to the shared memory. In the BSPC we're considering, it

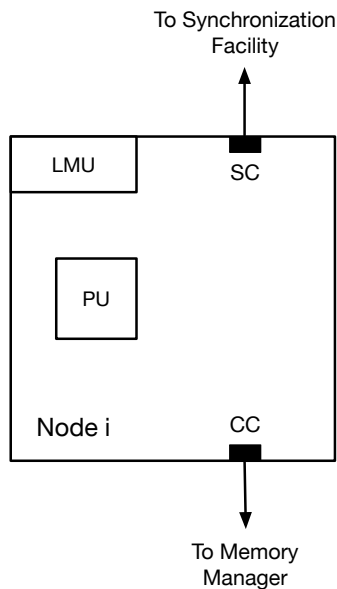


Figure 3.2: The sub-components of a BSP node.

is also the logically centralized entity that effectively realizes the interconnection network between nodes via communication channels (Section 3.2). For this reason, it corresponds to the original BSP’s definition of *communication component*.

The MM is connected to *all* BSP nodes via communication channels. On those channels, during the computation phase of the superstep, the MM receives read/write requests from nodes. Since these memory operations are required to be asynchronous in the BSP model (i.e., they immediately appear as completed but must actually take place at the end of a superstep), these requests are *stored* inside the MM. When the computation phase ends, the MM executes each request’s operation on the shared memory. The memory contention resolution strategy is not fixed; for our purposes, the MM uses the *CRCW* strategy, but it’s conceptually simple to enact the *EREW* or *CREW* strategies. After performing a memory operation, the MM sends the result of that operation to the relevant node (see Section 3.5).

Lastly, the memory manager has a bidirectional channel that connects it to the synchronization facility. The facility will notify the end of the local computation phase of all nodes to the MM, while the latter will notify to the former the end of all memory operations and – thus – of the whole superstep.

### 3.4 Synchronization facility

The behavior of the synchronization facility is heavily implementation-dependent, so its architectural specification is loosely defined.

The synchronization facility is the entity that enforces the computation structure of the BSP model, i.e. the local computation and communication phases of each superstep. It communicates with the BSP nodes via the *synchronization*

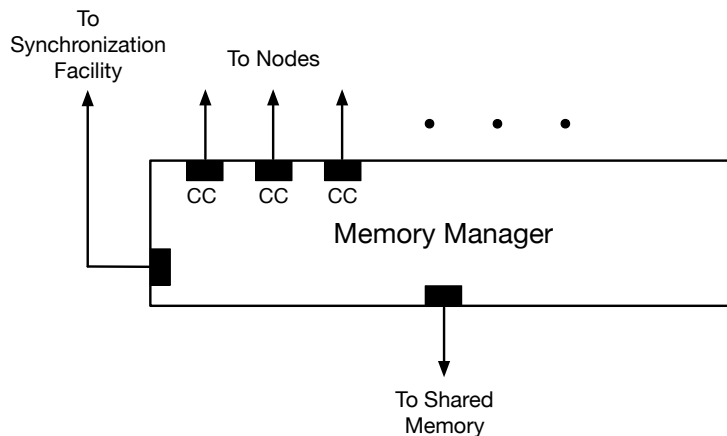


Figure 3.3: The Memory Manager component.

*channels* (SCs, see Section 3.2), and with the memory manager via a single bidirectional channel.

The synchronization facility periodically checks (via the SCs) if every node has finished its local computation for the current superstep. As soon as this condition is true, it asks the memory manager to perform the end-of-superstep memory operations and waits for it to signal the end of those operations. When the synchronization facility receives this signal, it broadcasts the superstep advancement message to all BSP nodes.

### 3.5 Shared memory

The defining characteristic of the BSPC that we are implementing is the shared memory component. Unlike its role in the BSPRAM model, where it was effectively the entity responsible for communication between nodes, here the shared memory takes the more passive role of “merely” being the memory space where the data to be used by multiple nodes is stored.

As previously discussed (Section 3.3), nodes can only access the shared memory by performing asynchronous requests via the MM. This means that data to be used in local computation inside a node must be copied in that node’s local memory, in order to avoid having to wait for the shared memory operation to conclude. The MM must perform this copy as efficiently as possible and maintain coherence between both types of memory (local node and shared). Thus, an adequate organization of the data in both the shared and local memories is needed.

First of all, we remind that the whole BSP computation is in fact a combination of smaller local computations which – as the name implies – are performed over *local data*. The result of a local computation is also stored in the local memory of the node that performs it. In order for a local computation to use data that belongs to other nodes (i.e. *global data*), that data must have been requested in advance and already copied inside the node; conversely, the result of a local computation can be sent to other nodes only after it has been written

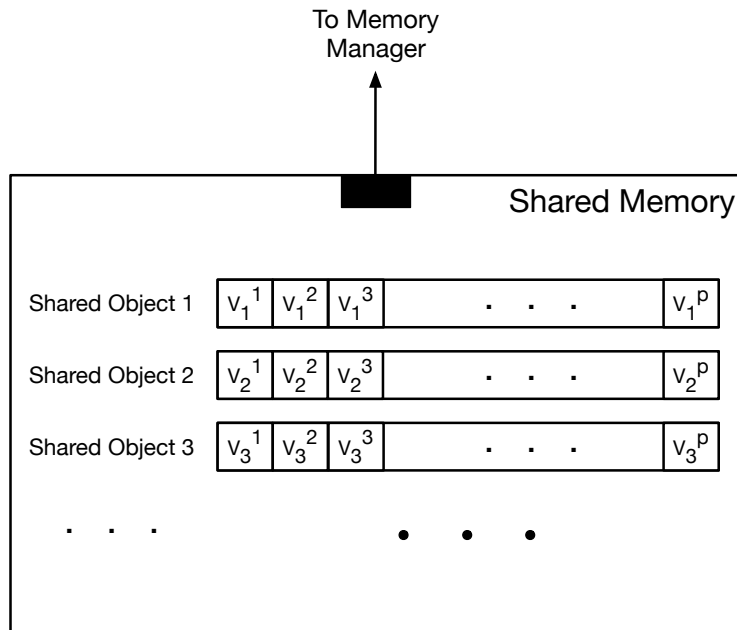


Figure 3.4: The Shared Memory entity.

in local memory. There is, therefore, a certain level of *data redundancy* that must be present in order for the BSP computation to take effect; this applies to any BSPC, not only to the one subject of this work.

The main ideas for data management in the architecture of our BSPC are the following:

1. each global data element is replicated on the shared memory  $p$  times (where  $p$  is the number of nodes). Basically, each node has its own private copy of the data element on the shared memory;
2. a write request  $write(element, payload, recipient)$  is fulfilled by the MM by writing the request payload in both the recipient's local memory and private copy of the element in the shared memory. The MM can perform both writes in parallel (one on the CC with the recipient node, one on the shared memory channel);
3. a read request  $read(element, source, destination)$  is fulfilled by the MM in the same way as a write request, i.e. the payload is the source's private copy of the element in the shared memory, which is copied into the destination's private copy and local memory in parallel.

The first idea trades off some shared memory space that may go unused (a private copy for every node is created for each data element, even if that element is not used at all in some nodes) for an increase in performance for read operations, since the element to be read is already present in the (fast) shared memory and there is no need to copy it from the source's local memory. The second and third ideas guarantee coherence between the data in the shared memory and their copies in the nodes' local memories.

## 3.6 Summary

In this Chapter we introduced an architectural design for an object-oriented BSP model on shared memory machines. We defined *nodes* as the computing entities, with fast local memory and a series of channels for communication and synchronization. We described a shared memory organization that allows for efficient retrieval of information and the network component in charge of communication and memory management. Finally, we detailed the tasks of a synchronization facility that regulates the BSP supersteps.

In the next Chapter we will describe a C++ implementation of a BSPC that follows this design. Each of the components described above will be represented by (one or more) C++ objects.

# Chapter 4

## Implementation

This chapter describes the implementation design and choices for the BSP library for FastFlow. The architecture of the BSPC described in Chapter 3 is realized over common shared-memory, multicore machines using the FastFlow framework discussed in Section 2.2. The proposed implementation fully embraces the object-oriented paradigm by having the *object* be the atomic data element of the BSP computation. The programmer is therefore relieved of having to manage memory space manually.

### 4.1 Library structure

The library is provided as a header-only library, for consistency with the FastFlow framework. Programmers who wish to use it can simply include the `bsp_program.hpp` file in their code.

Here is a list of the library files, together with a short description of their contents.

- `bsp_program.hpp` (Section 4.3)  
The FastFlow core of the BSP library. Defines the underlying FastFlow structure, i.e. a farm with custom emitter and collector, and provides the user with means to build the BSP computation by providing `bsp_node` objects and (optionally) functions to be called before and after the computation.
- `bsp_node.hpp` (Section 4.4)  
Specialization of `ff_node` to work as the computing entities of a BSPC. Users must write their own classes that inherit from `bsp_node` and specialize the `parallel_function` method.
- `bsp_internals.hpp` (Section 4.5)  
Forward declarations and definitions for the communication component, `bsp_container` and its specializations (`bsp_variable` and `bsp_array`).

- `bsp_communicator.hpp` (Section 4.6)  
Implementation of the communication component, including the requests mechanism and the methods responsible for the creation of new shared data elements.
- `bsp_variable.hpp` (Section 4.7)  
Implementation of the data structure that represents a shared element.
- `bsp_array.hpp` (Section 4.7)  
Specialization of the `bsp_variable` structure to handle multiple elements of the same type.
- `bsp_barrier.hpp` (Section 4.8)  
Implementation of a simple, reusable barrier that partially fulfills the role of the synchronization facility in a BSPC.
- `stl_allocator.hpp` (Section 4.9)  
C++11 wrapper for the `FFAllocator` class (FastFlow's version of the Slab Allocator, see Section 2.2.1).

## 4.2 Architecture

The BSPC architecture discussed in Chapter 3 forms the basis for this implementation, which targets single shared memory multi-core machines and relies on the FastFlow framework. It is easy to see that the simplest way to realize the BSPC is to map every entity on a FastFlow node, except for the shared memory entity which is mapped on the program's virtual memory. This mapping is unfortunately inefficient: the memory manager node remains idle while "processing" (BSP) nodes perform their local computation, and vice-versa all BSP nodes must wait that the MM performs its end-of-superstep routine without doing nothing.

A more reasonable mapping, which is used in this work, *distributes* the MM and synchronization facility tasks to FastFlow nodes, effectively mapping the two entities onto the latter.

## 4.3 `bsp_program.hpp`

The BSP implementation presented in this work uses FastFlow as the target parallel programming framework, namely it is implemented *on top* of it. The `bsp_program` class encapsulates the whole BSP computation. It is a specialization of the `ff_node` class and, as such, it can be used as the component of a broader pattern. It should be noted, however, that in the current state of the library the support for FastFlow composition is still basic, although fully functional: the BSP nodes can access the object received in the `bsp_program` FastFlow node as input, and can forward objects to the next stage. The actual BSP computation can be performed either via the `.start()` method, which is completely unrelated to the broader FastFlow computation, or by sending a token to the `bsp_program` node.






Figure 4.1: The FastFlow graph of the implementation.

The `bsp_program` is internally organized as a `ff_Farm` with custom Emitter and Collector nodes, organized in a `ff_Pipe`. The Worker nodes of the farm properly represent the BSPC, i.e. implement all non-memory entities of the BSPC discussed in Chapter 3. The specialized Worker nodes, objects of the `bsp_node` class, will be discussed in Section 4.4.

```

1  bsp_program(std::vector<std::unique_ptr<bsp_node>>&& _processors,
2              std::function<void(void)> _pre = nullptr,
3              std::function<void(void)> _post = nullptr):
4      nprocs{static_cast<int>(_processors.size())},
5      comm{nprocs},
6      barr{nprocs},
7      E{std::move(_pre), nprocs},
8      processors{std::move(_processors)},
9      C{std::move(_post), nprocs} {
10     for (size_t i{0}; i < nprocs; ++i) {
11         processors[i]->nprocs = nprocs;
12         processors[i]->id = i;
13         processors[i]->barrier = &barr;
14         processors[i]->comm = &comm;
15         C.master = this;
16     }
17 };

```

Listing 4.1: The `bsp_program` constructor.

The user is responsible for the creation of the worker nodes, which must then be provided to the `bsp_program` object by passing an `std::vector` of `std::unique_ptr`s to the nodes. The constructor for the `bsp_program` class will inject the necessary information into the `bsp_node` objects: in Listing 4.1, lines

13–14, it is shown how the logical Communication Channels and Synchronization Channels (see Section 3.2) are realized, i.e. as pointers to respective objects. Optionally, the user can also provide two functions `pre` and `post` that will be executed respectively before and after the BSP computation. These functions cannot take arguments and cannot return values. Internally, they are executed respectively by the `ff_Farm`'s Emitter and Collector nodes (Listing 4.2, line 6 and Listing 4.3, line 7).

```

1  struct emitter: ff::ff_node {
2
3      [...]
4
5      void* svc(void* in) override {
6          if (preprocessing != nullptr) preprocessing();
7          for (int i = 0; i < count; i++) {
8              // ENDCOMP is a special value to stop the
9              // FastFlow execution after the BSP
10             // computation ended
11             ff_send_out(ENDCOMP);
12         }
13         return EOS;
14     }
15 };

```

Listing 4.2: The emitter's `svc` method.

```

1  struct collector: ff::ff_node {
2
3      [...]
4
5      void* svc(void* in) override {
6          f (in == ENDCOMP) {
7              if (++count == threshold) {
8                  if (postprocessing != nullptr) postprocessing();
9              }
10             return GO_ON;
11         } else { // Request to forward onto later stages
12             master->forward(in);
13         }
14         return in;
15     }
16 };

```

Listing 4.3: The collector's `svc` method.

Listing 4.4 shows the method that actually builds and runs the whole BSP computation. First, the FastFlow graph is built by moving the `bsp_nodes` into the `Workers` entities of the `ff_Farm` (lines 2–6, 11). Then, the eventual FastFlow token is forwarded to the communicator so that it can be available to all BSP

```

1  void start(void* in = nullptr) {
2      std::vector<std::unique_ptr<ff::ff_node>> workers;
3      for (size_t i{0}; i < nprocs; ++i) {
4          auto d = static_cast<ff_node*>(processors[i].release());
5          workers.emplace_back(std::unique_ptr<ff_node>(d));
6      }
7
8      // Let all BSP nodes see the FastFlow input
9      comm.set_fastflow_input(in);
10
11     ff::ff_Farm<> farm(std::move(workers), E,C);
12
13     if (farm.run_and_wait_end() < 0)
14         std::cout << "error in running farm" << std::endl;
15
16     comm.end();
17 }

```

Listing 4.4: The start method of the `bsp_program` class.




Figure 4.2: Logical tasks performed by a runnable `bsp_node` object.

nodes (line 9). The `ff_Farm` is then started (line 13) and at the end (line 16) the cleanup function of the `bsp_communicator` class is called, in order to properly deallocate all relevant data structures (see Section 4.6).

## 4.4 `bsp_node.hpp`

The central unit of a BSP computation is the BSP node, an entity that is capable of processing data stored inside its own local memory and that can send and receive data to and from other nodes. The BSP node abstraction is realized into the `bsp_node` class, which is itself a specialization of the FastFlow `ff_node` and as such is the entity that is mapped onto the actual machine's processing elements, in this case POSIX threads.

As it will be discussed later (Sections 4.6 and 4.8), the communicator and synchronization entities are not implemented into runnable, standalone entities. Their tasks will, instead, be executed by the processing elements that run the `bsp_node` objects. In particular, communication requests are done by call-

ing the *put* and *get* methods of the `bsp_variable` and `bsp_array` classes (see Section 4.7). The end-of-superstep operations, be it signaling the end of local computation or processing the memory requests, are executed when the user calls the `bsp_sync` method (see Listing 4.5).

```
1 void bsp_sync() {
2     barrier->wait();
3     comm->process_requests(id);
4     barrier->wait();
5 }
```

Listing 4.5: The `bsp_sync` method.

The `bsp_node` class is quite simple. It provides protected methods to access information such as the number of BSP nodes or request new BSP variables. The intended usage for this class is for the programmer to create a class that inherits from `bsp_node` and implements the virtual `parallel_function` method. An example of inherited class can be found in the program `mwe.cpp` that has been provided with the library.

## 4.5 `bsp_internals.hpp`

This file mainly provides forward declarations and interfaces for other classes, namely `bsp_communicator` (Section 4.6), `bsp_variable.hpp` and `bsp_array.hpp` (Section 4.7). Nevertheless, two important data structures are defined here: communication *requests* and containers that implement the first idea for data management explained in Section 3.5, i.e. for every global data element the shared memory holds a private copy for each node.

### 4.5.1 Requests

Requests are represented by a record-like `struct` that holds the relevant fields as shown in Listing 4.7. The `request_type` enum (Listing 4.6) is used for discerning how the request itself must be processed at the end of the superstep. Note that not all actual request categories are represented — in particular, the “variable get” and “array element get” requests are stored with the same type as their “put” counterparts, but with the source and destination fields reversed.

```
1 enum request_type {
2     var_put,
3     arr_put_el,
4     arr_put,
5     arr_get
6 };
```

Listing 4.6: Possible request types.

In the “get” request types, the data object is gathered from the relevant memory element by the communicator. In the “put” ones, the data is instead provided by the node that performs the request. In both cases, the data must be *copied* to allow the original element to be modified by its owner without repercussions on the memory operation. The request object contains a pointer to the copy, wrapped in a `std::shared_ptr` object to allow for simpler dynamic memory management.

```
1  struct request {
2      request_type t;
3      int reference;
4      int source;
5      int destination;
6      int src_offset;
7      int dest_offset;
8      int length;
9      std::shared_ptr<void> element;
10
11     // Constructor...
12 };
```

Listing 4.7: Fields of a request data structure.

## 4.5.2 Shared memory containers

As outlined in Section 3.5, the ideal memory organization for shared memory elements in our BSPC consists in duplicating each element  $p$  times, where  $p$  is the number of BSP nodes. This allows for efficient memory operations at the cost of increased space usage<sup>1</sup>.

As it will be discussed further in Section 4.7, the BSP library for FastFlow explicitly place C++ *objects* as the unit of data exchanged in communication between nodes. `bsp_variables` and `bsp_arrays` are therefore template classes whose type arguments are essentially without restraints. This allow maximum flexibility from the user point of view, but provides the additional challenge to having to store heterogeneous data in the single data container that implements the shared memory entity. This would not be a problem in a language that supports type erasure and reflection like Java, where a container like `ArrayList<Object>` would have sufficed and type information could be saved in a Class object, but unfortunately C++ does not provide such mechanisms. We decided to implement our own type erasure variant by using `void` pointers and closures to obtain respectively storage of heterogeneous data and a way to save relevant type information.

The `inner_var` (Listing 4.8) and `inner_array` (Listing 4.9) are private inner classes of the `bsp_communicator` class. They both contain a vector of  $p$  `void*`

<sup>1</sup>Actually, in this implementation we avoided the redundancy by making nodes hold *pointers* to their private data, instead of having to copy entire elements to the nodes' local memories (represented as class fields in the implementation). These pointers are furthermore invisible to the user, so “BSP-unsafe” modifications of the shared object are impossible unless explicitly requested (Section 4.7).

```

1  struct inner_var {
2      // vector of nprocs elements
3      std::vector<void*, ff::FF_Allocator<void*>> element;
4      // bookkeeping functions needed to work with void*
5      // function to replace a variable with another value
6      void (*swap)(void* el, void* other);
7      // function to safely free memory occupied by an element
8      void (*free_el)(void* el);
9
10     // Constructor...
11 };

```

Listing 4.8: The inner\_var class.

objects and functions that will enclose some type information. Objects of any type can be stored in an inner\_var or inner\_array object and, since the latter classes are not templated, their objects can be stored in a regular container like std::vector regardless of what they contain.

```

1  struct inner_array {
2      // vector of nprocs arrays
3      std::vector<void*, ff::FF_Allocator<void*>> element;
4      // bookkeeping functions needed to work with void*
5      // function to replace an element of the array with another value
6      void (*put)(void* el, void* toput, int pos);
7      // function to replace a portion of the array
8      void (*replace)(void* el, int srcof, int dstof,
9                      int len, void* toput);
10     // function to safely free memory occupied by an element
11     void (*free_el)(void* el);
12
13     // Constructor...
14 };

```

Listing 4.9: The inner\_array class.

A inner\_var object that stores an element of type T must be provided with a swap(element, other) function that copies the value pointed by other into the variable pointed by element. Both element and other are passed as void\* and must be converted to T\* inside the function, which effectively saves the type information for T inside the function. The put(element, other, position) and replace(element, offset1, offset2, length, other) functions of the inner\_array class have the same purpose. Both inner\_var and inner\_array, moreover, have a free\_el(element) function to deallocate an element. The task of creating these functions is not left to the user and are instead generated by the bsp\_communicator class when a new BSP variable or array is requested.

## 4.6 bsp\_communicator.hpp

As said in the introduction to this chapter, the implementation for the BSPC discussed in Chapter 3 does not assign every logical entity to a physical one. This is especially true for the `bsp_communicator` class, which not only represents two logical entities — the Memory Manager and the Shared Memory — but is also distributed over the FastFlow nodes instead of being an autonomous centralized node.

### 4.6.1 Implementation of the Shared Memory entity

The Shared Memory entity is implemented simply as a vector of `inner_vars` or `inner_arrays` that represent shared elements (see Section 4.5.2). To obtain a higher performance when accessing those containers, lookup maps – which associate *hashes* of the variables (or arrays) to their position in the vector – are used. The hashing function used for this purpose is the one commonly known as djb2 ([53, 62], shown in Listing 4.11), which is a simple yet effective [29] hash that takes a string and an integer and returns another integer. In our case, the string argument is an unique representation of the variable’s type obtained with the `typeid` builtin, while the integer argument is obtained from the current superstep number and the total number of variables of type `T` requested by the current node, for reasons that will become clear shortly. The hash result is then used as the key for a lookup map, whose values are positions in the vector of `inner_vars`. C++17 `std::shared_mutexes` regulate concurrent access to the shared storages: those mutexes allow for a single writer-multiple readers pattern, so any number of nodes can access (and modify!) already-existing shared elements, but only a single node at a time can create a new variable or array (lines 10 and 12 of Listing 4.10).

One of the problems that arised when devising the implementation was how to allow users to request access to the *same* shared variable<sup>2</sup>. For example, take this pseudo-code:

```
bsp_variable v = get_variable<double>();
if (node_id == 0)
    v.put(payload=5.0, destination=1);
```

Suppose we execute this code on a BSPC with two nodes. The communicator entity receives a request for a `bsp_variable` of type `double` from both nodes, but does not know if they are asking for access to the same variable or for the creation of two separate variables. There are two ways to solve this problem:

- let the user specify unique IDs when requesting variables (e.g. as strings or integers), and treat variables of the same type with the same ID as the same shared object;
- let the library manage variable requests automatically, providing the user with a clear rule for defining when a fresh shared object is created.

We chose the second option, according to the library’s philosophy of relieving the user from the burden of explicitly managing memory aspects. We use the following proposition.

---

<sup>2</sup>Wherever we refer to variables in this section, the same concept applies for arrays.

```

1 // number of current superstep
2 int generation = 1;
3
4 // counts the variables requested by each worker in this superstep
5 std::map<std::string, int>* var_count;
6 // counts the arrays requested by each worker in this superstep
7 std::map<std::string, int>* arr_count;
8
9 // mutex for multiple readers-single writer access to the variables
10 mutable std::shared_mutex var_mutex;
11 // mutex for multiple readers-single writer access to the arrays
12 mutable std::shared_mutex arr_mutex;
13
14 // dictionary for quick retrieving of variables based on their hash
15 std::map<int, int> variable_dict;
16 // dictionary for quick retrieving of arrays based on their hash
17 std::map<int, int> array_dict;
18
19 // actual storage of variables
20 std::vector<inner_var> variables_storage;
21 // actual storage of arrays
22 std::vector<inner_array> arrays_storage;

```

Listing 4.10: Fields of the `bsp_communicator` class used to implement the Shared Memory logical entity.

```

1 static int get_hash(const char* s, int seed) {
2     int hash = seed + 5381;
3     while (*s) {
4         hash = hash * 33 ^ (*s++);
5     }
6     return hash;
7 }

```

Listing 4.11: The hashing function for variables and arrays.



**Proposition 1.** During the  $i$ -th superstep, for any node the  $n$ -th request for a variable of type  $T$  will return a handle to the  $n$ -th shared element of type  $T$ .

In practice, this means that when a node requests a variable of type  $T$  for the first time in a superstep, and it is the first node in chronological order to do so, a new shared element of type  $T$  is created. When any other node request their first variable of type  $T$ , they receive a handle to (their private data of) this shared element, and so on. The method to request a variable is shown in Listing 4.12, and the one to request an array is similar. Lines 2–12 show how the hash value used in the lookup map is obtained: the method then checks if a variable with the same hash is present (lines 13–17) and, if not, a new `inner_var` shared element is created. Lines 22–28 show how type information is captured inside the functions passed to the `inner_var` constructor. The newly-created shared element is then pushed into the shared memory (line 30) and the lookup map is updated (line 31). The shared element’s private data – relative to the node that requested it – is updated with the provided value (line 37), and a `bsp_variable` containing that pointer is returned to the caller (line 39).

## 4.6.2 Implementation of the Memory Manager entity

The `bsp_communicator` class also fulfills the role of the Memory Manager, i.e. the entity that collects communication requests during the local computation phase of a superstep and executes them at the end of it. As we said before, there is no autonomous entity that performs the Memory Manager tasks in a centralized way. FastFlow nodes – the same that run BSP nodes – will handle both request management and execution. The Memory Manager implementation shown here (Listing 4.13) has  $p$  request queues, one for each BSP node. Concurrent access to these queues is regulated by a series of mutexes, one mutex per queue. The  $i$ -th request queue will hold all requests that have the  $i$ -th node as *destination*. In this way, at the end of the superstep, the FastFlow node that is running BSP node  $i$  can process all requests directed to it with no concerns over concurrent accesses.

Listing 4.14 shows an example of how requests are inserted and managed. First of all (line 4), exclusive access to the destination node’s request queue is obtained by locking the relevant mutex. Then, the request proper is created and enqueued (lines 5–6), and finally the mutex is released. `m_shared_ptr` is a macro to simplify the creation of a `std::shared_ptr` using the custom `stl_allocator` described in Section 4.9.

Listing 4.15 shows instead how requests are processed by the various FastFlow nodes at the end of a superstep. Each node processes the requests which feature it as destination (i.e. the ones who have effect on the node’s private data of shared elements). In this way, concurrent access to the same portion of shared memory is avoided and no locking mechanism is needed. Each request is then processed according to their type; lines 5–11 of Listing 4.15 show, for example, how a variable put request is managed. The “type information-holding” functions of the `inner_var` and `inner_array` classes are used here: since the data is stored as `void*` pointers, these functions are needed to cast it back to its proper type before it can be modified (Section 4.5).

```

1  bsp_variable<T> get_variable(int holder, T* initial_val) {
2      auto tname = typeid(T).name();
3      // no. of variables of type T requested by the current worker
4      int get_count = 0;
5      try {
6          get_count = var_count[holder].at(std::string(tname));
7      } catch (const std::out_of_range&) {
8          var_count[holder].insert({std::string(tname), 0});
9      }
10     // hash on type, superstep and number of vars
11     int hash = get_hash(typeid(T).name(),
12                      (generation * 5000000) + get_count);
13     int ref;
14     // try to find a variable w/ the same hash
15     try {
16         std::shared_lock lock(var_mutex);
17         ref = variable_dict.at(hash);
18     } catch (const std::out_of_range&) {
19         std::unique_lock lock(var_mutex);
20         auto iter = variable_dict.find(hash);
21         if (iter == variable_dict.end()) {
22             inner_var var{nprocs, [] (void* el, void* other){
23                 auto tptr = static_cast<T*>(el);
24                 *tptr = *(static_cast<T*>(other));
25             }, [] (void* el){
26                 auto tptr = static_cast<T*>(el);
27                 delete tptr;
28             }};
29             ref = variables_storage.size();
30             variables_storage.push_back(var);
31             variable_dict.insert({hash, ref});
32         } else {
33             ref = iter->second;
34         }
35     }
36     // initialize with the value requested by the worker
37     variables_storage[ref].element.at(holder) = initial_val;
38     var_count[holder].at(std::string(tname))++;
39     return bsp_variable<T>{ref, holder, this, initial_val};
40 }

```

Listing 4.12: The get\_variable method.

```

1  std::mutex* mutexes;
2  std::vector<request, ff::FF_Allocator<request>>* requests;
3
4  std::atomic_int process_count{0};

```

Listing 4.13: Fields of the bsp\_communicator class used to implement the Memory Manager logical entity.

```

1  template <typename T>
2  void variable_put(int what, int source,
3      int destination, const T& element) {
4      mutexes[destination].lock();
5      requests[destination].emplace_back(request_type::var_put,
6          what, source, destination, 0, 0, 0, m_shared_ptr(T, element));
7      mutexes[destination].unlock();
8  }

```

Listing 4.14: The variable put request handler of the bsp\_communicator.

```

1  void process_requests(int id) {
2      // Request filtering by worker id
3      for (const auto& req: requests[id]) {
4          switch (req.t) {
5              case request_type::var_put: {
6                  auto ptr = variables_storage.at(req.reference)
7                      .element[req.destination];
8                  variables_storage.at(req.reference)
9                      .swap(ptr, req.element.get());
10                 break;
11             }
12             case request_type::arr_put_el: { ... }
13             case request_type::arr_put: { ... }
14             case request_type::arr_get: { ... }
15         }
16     }
17     if (++process_count == nprocs) {
18         generation++;
19         delete[] arr_count;
20         arr_count = new std::map<std::string, int>[nprocs]();
21         delete[] var_count;
22         var_count = new std::map<std::string, int>[nprocs]();
23         delete[] requests;
24         requests = new std::vector<request,
25             ff::FF_Allocator<request>>[nprocs]();
26         process_count = 0;
27     }
28 }

```

Listing 4.15: Request processing.

Lines 17–27 are only executed by the last node that finishes processing its requests. They “clean the state” relative to the current superstep, i.e. reset the variable and array counters to 0 and empty all request queues.

```
1  void bsp_communicator::end() {
2
3      for(auto& var: variables_storage) {
4          auto del = var.free_el;
5          for (auto& el: var.element) {
6              del(el);
7          }
8      }
9
10     for(auto& arr: arrays_storage) {
11         auto del = arr.free_el;
12         for (auto& el: arr.element) {
13             del(el);
14         }
15     }
16
17     delete[] arr_count;
18     arr_count = nullptr;
19     delete[] var_count;
20     var_count = nullptr;
21     delete[] requests;
22     requests = nullptr;
23     delete[] mutexes;
24     mutexes = nullptr;
25 }
```

Listing 4.16: The `bsp_communicator` final cleanup function.

Finally, Listing 4.16 shows the `end()` method of the `bsp_communicator` class. It is a cleanup function that gracefully deallocates all shared objects and data structures used by the Shared Memory and Memory Manager entities. It will be called by the encompassing `bsp_program` object at the very end of the BSP computation, after all nodes have exhausted their local computations (see Section 4.3).

## 4.7 `bsp_variable.hpp` and `bsp_array.hpp`

As outlined in Section 3.5, even in our shared-memory BSFC each node must have all data needed for computation in its local memory, and this includes shared objects. There is therefore a need for a representation of those shared elements inside the BSP node. The `bsp_variable` and `bsp_array` fulfill this task. A `bsp_variable<T>` is a template class that holds the local copy of a shared object of type `T`, and it also represents the handle of the shared object that the user can employ to perform memory requests on it. Due to limitations in the STL functions and data structures used in the implementation (in particular the `std::allocate_shared` function, see [49]), the type `T` *cannot* be a C-style

array type, e.g. E[]. This restriction is lifted in the C++20 version of the language. The `bsp_array<T>` class is similar to the previous one, except for the fact that the shared object is a resizable vector of elements of type T. Therefore, the class provides methods for vector manipulation like iterators and access to elements. Unlike `bsp_variable`, in the `bsp_array<T>` class T can be a C-style array type.

```

1  template <typename T>
2  class bsp_variable: public bsp_container {
3
4      // The template type must be copy-constructible and copy-assignable
5      static_assert(std::is_copy_constructible<T>::value &&
6                    std::is_copy_assignable<T>::value,
7                    "Type of bsp_variable must be "
8                    "copy-constructible and copy-assignable");
9  private:
10     ...
11     T* element; // Pointer to the actual element
12     ...
13  public:
14     // Replaces the element on another worker
15     void bsp_put(const T& elem, int destination) {
16         comm->variable_put<T>(reference, holder, destination, elem);
17     }
18
19     ...
20
21     // Returns the element in another worker's memory
22     // directly (without waiting for a superstep sync)
23     T bsp_direct_get(int source) {
24         return comm->variable_direct_get<T>(reference, source);
25     }
26
27     // Allows direct access to the element
28     // in this worker's memory
29     T& direct_access() {
30         return *element;
31     }
32 };

```

Listing 4.17: The `bsp_variable` class.

Listing 4.17 shows a portion of the `bsp_variable<T>` class. First of all, in order to perform BSP memory operation, elements of type T must be copy-constructible and copy-assignable (lines 4–7); this is needed, since objects have to be copied inside requests and successively into the shared memory. The local copy of the private data of the shared object is implemented simply as a pointer to the data stored inside the `bsp_communicator` object, to save some copies. The user can call the methods `bsp_put` and `bsp_get` of the class (with different signatures according to the desired operation to be executed) to perform memory requests on the current shared object. The class provides also

two “BSP-unsafe” functions: `bsp_direct_get` and `BSPunsafe_access`. The first one (lines 23–25) returns a *copy* of a node’s private data of the shared element; This violates the BSP condition that a node cannot access another node’s local or private memory, thus the function is deemed “BSP-unsafe”. Still, the fact that a *copy* of the desired data is returned means that this function can reasonably be used inside a BSP computation without having to consider unwanted side effects.

The second “BSP-unsafe” method is more problematic, since it gives access to the *actual element* in the shared memory (lines 29–31). This is a big violation of our BSPC’s premise that the Memory Manager entity is the only one allowed to operate on shared memory, and in general of the whole BSP concept of only allowing communication effects to take place at the end of the superstep. Both `direct_get` and `BSPunsafe_access` functions are provided for compatibility with similar methods used in the MulticoreBSP for Java library (see [40]).

The `bsp_array<T>` class is very similar to the `bsp_variable<T>` one, so it won’t be discussed as much. `T` still needs to be a copy-constructible and copy-assignable type, and the class provides specialized `bsp_put` and `bsp_get` methods for single elements of the vector. Arrays also have another BSP unsafe method, `BSPunsafe_put`, that places an element into the node’s private copy without waiting for the end of the superstep.

## 4.8 `bsp_barrier.hpp`

The synchronization facility of our BSPC (Chapter 3) is the entity that properly manages the flow of supersteps of the whole computation: its tasks are to

1. check wherever all BSP nodes have finished their local computation;
2. wait that the Memory Manager finishes processing requests;
3. begin the next superstep, updating the BSPC’s relevant structures (see Listing 4.15) and signaling the nodes to begin local computations again.

The request processing and the update portion of task 3 are implemented in our library in the `bsp_communicator` class (see Section 4.6). This leaves the implementation of the actual check and the signal to start the new superstep. A synchronization mechanism that implements them in a very simple way is the *barrier*. Nodes that finish their local computation will wait at the barrier for their peers to do the same. When all nodes reach the barrier, the latter is lifted and they can begin the local operations of the next superstep. As seen in Listing 4.5, the barrier is called twice — once when the nodes finish their local computation, and once when the Memory Manager finishes processing requests (remember that both operations are effectively executed by the physical entity that runs the underlying FastFlow node).

The barrier structure used in the implementation is shown in Listing 4.18. It’s a simple structure that uses a condition variable (line 4) and its wait/notify mechanism (lines 35–38), together with a generation counter to allow for multiple uses [63].

```

1  class bsp_barrier {
2      private:
3          std::mutex mutex;
4          std::condition_variable cond_var;
5          // Number of workers needed to release the barrier
6          int threshold = 0;
7          // Count of workers currently at the barrier
8          int count = 0;
9          // Number of times the barrier has been used
10         int generation = 0;
11
12     public:
13         bsp_barrier() = delete;
14
15         explicit bsp_barrier(int size):
16             threshold{size},
17             count{size},
18             generation{0} {
19         }
20
21         // Copy constructor
22         bsp_barrier(const bsp_barrier& other):
23             threshold{other.threshold},
24             count{other.count},
25             generation{other.generation}{
26         }
27
28         // Barrier function
29         void wait() {
30             std::unique_lock<std::mutex> lock{mutex};
31             auto lastgen = generation;
32             if (!(--count)){
33                 generation++;
34                 count = threshold;
35                 cond_var.notify_all();
36             } else {
37                 cond_var.wait(lock,
38                     [this, lastgen]() { return lastgen != generation; });
39             }
40         }
41     };

```

Listing 4.18: The bsp\_barrier class.

## 4.9 `stl_allocator.hpp`

The FastFlow allocator discussed in Section 2.2.1 performs better than the standard C++ allocator when the program has to do multiple allocations of small memory areas ([56], also discussed in Chapter 5). Oftentimes, the request mechanism of the Memory Manager exhibits this behavior, so it can be helpful to use the FastFlow allocator in place of the standard one. Unfortunately, this allocator does not provide an “STL-conforming” interface [5], so it has to be used manually (e.g. calling the `malloc` and `free` functions of an `FFAllocator` instance). The `stl_allocator.hpp` file contains a very simple interface (Listing 4.19) that allows the use of the FastFlow allocator wherever an STL allocator is required, e.g. in the `std::allocate_shared` function.

```
1  template <typename T>
2  class FF_Allocator {
3  public:
4      using value_type = T;
5      using propagate_on_container_move_assignment = std::true_type;
6      using is_always_equal = std::true_type;
7
8      FF_Allocator() noexcept = default;
9      template <class U>
10     explicit FF_Allocator(const FF_Allocator<U>& other) noexcept {};
11
12     value_type* allocate(std::size_t n) {
13         return static_cast<value_type*>(FFAllocator::instance()
14             ->malloc(n* sizeof(value_type)));
15     }
16
17     void deallocate(value_type* ptr, std::size_t) noexcept {
18         FFAllocator::instance()->free(ptr);
19     }
20 };
```

Listing 4.19: The FastFlow allocator “STL interface”.



# Chapter 5

## Experimental validation

### 5.1 Performance metrics

Before we delve into the analysis of the performance of the library presented in this work, it is important to lay the foundations for the metrics used throughout the Chapter. First of all, we formally define the most straightforward parameter of the set: the time needed for a computation to take place from start to finish.

**Definition 5.1.** The *total time* it takes to compute a single result is called **latency**, usually referred to with the symbol  $L$ .

The latency is a first indicator of how well the program behaves. Clearly, the lower  $L$  is, the better the program is considered. Nevertheless, if to obtain a slightly lower latency we must massively increase the resources in play, the tradeoff is not always worthwhile. We need a metric to define how “well-utilized” is the machine that performs the computation. Since we are dealing with parallel programs, the number of workers (nodes, cores, processors) used in the computation has an effect on the latency.

**Definition 5.2.** A parallel computation over  $p$  workers is said to have latency  $T_{par}(p)$  (or  $T_p$ ).

As a particular case,  $T_{par}(1)$  is the latency of a computation over a single worker. Oftentimes, executing a parallel algorithm on a single worker will still build a set of data structures and mechanisms for parallel synchronization and communication which is not actually needed inside the computation itself, therefore inflating the latency. For this reason, analyses on the performance of a parallel algorithm or library are conducted with regards to the latency of *the best sequential algorithm* that solves the same problem without overheads from parallel-specific mechanisms.

**Definition 5.3.** A purely-sequential algorithm with no parallel overhead is said to have latency  $T_{seq}$ .

We can now introduce relevant metrics for performance evaluation.

**Definition 5.4** (Speedup). The ratio between the latency of a sequential computation and the latency of the corresponding parallel computation over  $p$  workers is called **speedup**:

$$sp(p) = \frac{T_{seq}}{T_{par}(p)} \quad (5.1)$$

The speedup roughly represents “how much” the latency changes when switching from a purely sequential computation to a parallel one over  $p$  workers. The best result one can hope to achieve is to perfectly cut the latency of the sequential computation down by a factor  $p$ :

$$T_{par}(p) = \frac{1}{p}T_{seq} \implies sp(p) = \frac{pT_{seq}}{T_{seq}} = p$$

This behavior is called *linear speedup* and it is rarely encountered, due to parallel computations needing coordination and communication between them; in fact, the speedup for a single worker

$$sp(1) = \frac{T_{seq}}{T_{par}(1)}$$

is almost never equal to 1 due to the abovementioned need to build all structures needed for the parallel computation. This suggests a new metric that closely encompasses “how better” a parallel computation behaves when increasing the number of workers.

**Definition 5.5** (Scalability). The ratio between the latency of a parallel computation over a single worker and the latency of the same parallel computation over  $p$  workers is called **scalability**:

$$sc(p) = \frac{T_{par}(1)}{T_{par}(p)} \quad (5.2)$$

This metric is sometimes called *relative speedup*, while the quantity of equation 5.1 is called *absolute speedup* [37].

Finally, we define a measure that represents “how much” of the hardware is used during a parallel computation.

**Definition 5.6** (Efficiency). The **efficiency** is the ratio between the speedup of a parallel computation and the number of workers:

$$\epsilon(p) = \frac{sp(p)}{p} = \frac{T_{seq}}{pT_{par}(p)} \quad (5.3)$$

A parallel computation that achieves linear speedup  $sp(p) = p$  will have efficiency

$$\epsilon(p) = \frac{p}{p} = 1$$

i.e. it fully utilizes the available hardware resources. Using the above definition of scalability, one can also define a *relative efficiency* as

$$\epsilon_r(p) = \frac{sc(p)}{p}.$$

In some cases, one can attain an efficiency greater than 1, i.e.  $sp(p) > p$ . This phenomenon is called *superlinear speedup* and may be due to various factors such as

- a better cache exploitation in the parallel program w.r.t. its sequential counterpart;
- a different organization of data objects for parallel programming, improving data locality even when the parallel computation uses a single thread;
- the parallel algorithm is simply more efficient than the sequential one, in the sense that it may skip unnecessary work (e.g. earlier branch pruning in search tree problems).

These situations are rarely encountered, therefore sublinear speedups are more common.

## 5.2 Machine architecture

Tests were conducted on a machine of the Department of Computer Science, University of Pisa. It features an Intel processor of the Xeon Phi x200 (Knights Landing) family, sporting 64 cores at 1.3 GHz [28]. Each core can have up to four threads, for a total of 256 threads. The machine has 48 GB of RAM and runs the GNU/Linux CentOS 7.2.1511 operating system, kernel version 3.10. The C++ programs were compiled using gcc version 7.3.0 (libstdc++ version 6.0.24), while the Java programs (packaged into JAR files) were executed using OpenJDK 1.8.0\_201. The JAR files were instead compiled with the official Oracle Java SDK version 13.0.1 (at language level 8) on a different machine with a 2.3 GHz dual-core Intel Core i5 processor, 8 GB RAM, running macOS 10.14.6. The FastFlow version used is 2.2.0, checked out from the official repository at commit 8b9105d [20, 19].

## 5.3 Benchmarks

The BSP model has traditionally been used for solving numerical problems, especially in the field of linear algebra. In [9], Bisseling introduced a series of programs written in C using the BSP model, the *BSP educational package*. It features three main BSP algorithms for solving the following problems:

- *LU* decomposition of a square matrix;
- the Fast Fourier Transform of a vector of complex numbers;
- the sparse matrix-vector multiplication.

The BSPedupack also includes a program for testing the parameters of the BSPC in use and a “toy” program (used as a tutorial of sorts for the BSP model) that performs the inner product of two vectors.

The BSPedupack suite is one of the most widespread examples of BSP programs: most BSP implementations use it for testing performance and correctness. Bisseling’s original version targets the BSPlib interface for C [8], but a version also exists for the MulticoreBSP for Java library.

When choosing the algorithms to use for testing our library, the BSPedupack ones seemed the obvious choice. We chose, though, to substitute one of the programs – the sparse matrix-vector multiplication – with a non-numerical algorithm, to show an example of how the BSP model is used for solving problems in other fields. The “replacement” program implements a parallel sorting algorithm.

In general, when possible, the programs written for the C++ library derive directly from the BSPedupack for Java versions. This is to show that is easy to port existing programs to our library, requiring only slight adjustments besides having to change language-specific constructs and data structures. The sorting program has been instead written originally for our library and later ported to Java.

All the programs are self-contained and require no input other than the problem size and the number of desired parallel workers. Each program will (often randomly) populate the data structures on which the computation will be performed.

For the purpose of performance metrics calculations, for each parallel program a corresponding sequential version has been implemented. When possible, the sequential algorithm used is the closest version to the parallel one, using the same data structures.

All the tests have been executed using the same methodology. A suitable problem size is chosen in advance, large enough for the computation time to be much higher than the time needed for maintaining parallel structures. The program is then executed six times with this problem size, starting with parallelism degree 1 and doubling it each time up to a ceiling of  $p = 64$ . The time spent in the parallel portion of every execution is logged. The corresponding sequential program is also executed with the same problem size, logging the actual computation time (i.e. not counting data generation).

We call the set of these seven executions a *run*. Ten runs for each programs are executed; for each run and each parallelism degree we discard the maximum and minimum times and calculate the average of the remaining ones. This average time spent in the computation forms the basis for the calculations of the abovementioned performance metrics. The tests aim to draw comparisons between three kinds of BSP libraries:

- the one presented in this work that uses the FastFlow slab allocator (Section 2.2.1)
- a variant of the same one that uses the standard library (STL) allocator;
- the MulticoreBSP for Java library.




Figure 5.1: An example of cyclic distribution of an array over 3 processors.

### 5.3.1 Vector inner product

The first program is meant to be a simple introduction to BSP programming. Given two vectors of  $n$  integer numbers

$$\vec{x} = [x_0, x_1, \dots, x_{n-1}] \quad \vec{y} = [y_0, y_1, \dots, y_{n-1}]$$

the **inner product** (or *dot product*) of the two vectors is defined as follows:

$$\vec{x} \cdot \vec{y} = \vec{x} \vec{y}^T = \sum_{i=0}^{n-1} x_i y_i \tag{5.4}$$

Given the input size  $n$ , the program creates the vector  $\vec{x} = [1, 2, \dots, n]$  (i.e.  $x_i = i + 1$ ) and calculates the inner product of  $\vec{x}$  with itself:

$$\vec{x} \cdot \vec{x} = \sum_{i=0}^{n-1} x_i x_i = \sum_{i=1}^n i^2.$$

The sequential algorithm simply scans the vector once, accumulating the squares of the  $x_i$  elements in a variable. To check for correctness, the known formula

$$\sum_{i=1}^n i^2 = \frac{1}{6} n(n+1)(2n+1) \tag{5.5}$$

is used for rapidly calculating the correct value.

The BSP algorithm distributes the input array over  $p$  processors according to a *cyclic distribution*: processor  $k$  is assigned all the elements  $x_i$  such that  $i \bmod p = k$ . In Figure 5.1 the array cells with the same color are assigned to the same processor. Each processor then computes (sequentially) the inner product of its portion of vector and broadcasts the value to all other processors. Finally, all the processors (redundantly, but in parallel) sum the received data to obtain the final result.

The performance metrics for this program were evaluated by running it with a problem size of  $n = 2 \cdot 10^9$ . Figure 5.2 shows a comparison of average latencies (completion time) for the three libraries. (The series shown as “0” on the plot refers to the sequential program execution.) The MulticoreBSP for Java implementation performs better (by less than 10%) than either C++ ones when using a single thread. As the parallelism degree increases, the two C++ implementations prove to be faster, albeit by a narrow margin. As this is the simplest BSP program, with very few communications and a sequential portion that can be easily vectorized by compiler optimizations, these kinds of results are to be expected. Note that, between the two C++ implementations, the one that uses the standard allocator achieves lower latencies at any parallelism degree.

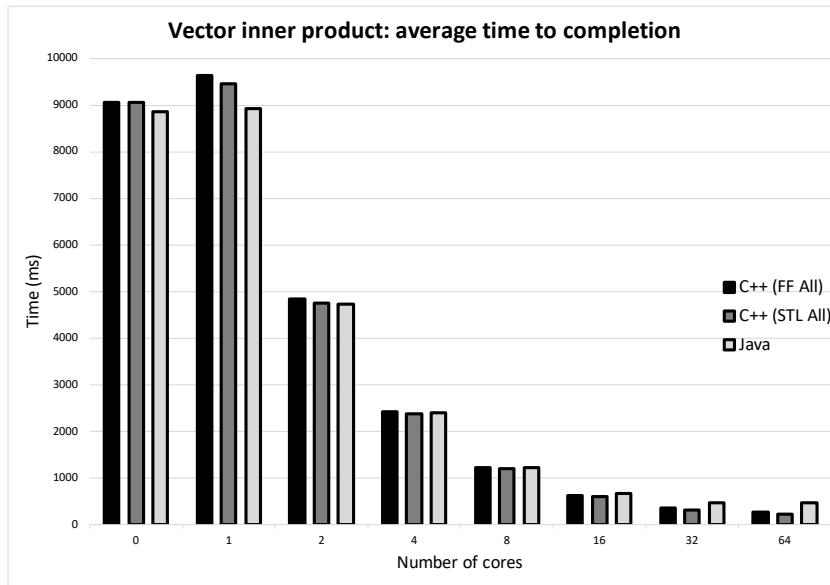


Figure 5.2: Average completion time for the vector inner product benchmark on 2B elements.




Figure 5.3: Speedup for the vector inner product benchmark.

Figure 5.3 shows the speedups achieved by the three libraries. The Java implementation cannot maintain a near-linear speedup for parallelism degrees higher than 8, and after  $p = 16$  the gains are minimal (quadrupling the cores only increases the speedup of a factor  $\sim 0.6$ ). Both C++ implementations ex-

hibit a decent speedup up to 32 workers, and even after that the speedup increase is still acceptable, especially for the implementation with the STL allocator. The slab allocator version reaches a sort of middle ground between the two others.

The trends shown in Figure 5.4 about scalability are similar to the ones of Figure 5.3. In particular, the Java implementation scalability exhibits almost the same behavior as its speedup, meaning that it holds  $T_{seq} \approx T_{par}(1)$  (i.e., the creation and maintenance of structures for supporting parallelism has a negligible effect on the latency). Conversely, both C++ implementations exhibit consistently higher scalability than speedup.




Figure 5.4: Scalability for the vector inner product benchmark.

Lastly, Figure 5.5 sums up the behavior of the three implementations with regards to this program by showing their efficiency parameter. As said before, the Java library offers good efficiency at low parallelism degrees, but it is out-classed by the other implementations after  $p \geq 4$  (and starts to decline rapidly for  $p \geq 8$ ). The C++ library on the other hand exhibits a good level of efficiency, between 0.9 and 0.95 for both variants up to  $p = 16$ . At  $p = 32$  the implementation that uses the slab allocator begins to fall behind, while the one with the STL allocator maintains an efficiency of more than 0.9. Finally, at  $p = 64$  both implementations show a definite decline in efficiency.

### 5.3.2 Parallel sorting by regular sampling

Sorting is one of the most important and well-studied operations in computing, as it is used as a subroutine of a large number of programs. Many parallel algorithms have been designed for this purpose [30, 1, 10], and with the advent of both new technologies (e.g. GPGPUs) and new use cases (e.g. big data analysis) sorting continues to be a relevant field of research [48, 43].




Figure 5.5: Efficiency for the vector inner product benchmark.




Figure 5.6: An example of block distribution of an array over 3 processors.

There exists a number of algorithms for parallel sorting on the BSP model [23, 24]. Between them, we chose Tiskin’s version of the sorting by regular sampling algorithm [55, 46, 32], which is easy to understand and allows us to illustrate another form of distribution of the input, the *block distribution* (Figure 5.6). In this kind of distribution, each processor is assigned a contiguous block of the input: with a problem of size  $n$  and  $p$  processors, each one receives  $n/p$  elements. The block distribution can lead to poor load balancing if  $n/p$  is not an integer number.

The algorithm, as its name implies, is guided by *regularly-spaced samples* of the data elements. First of all, each processor sorts its own subarray using a sequential algorithm. After that, it selects  $p + 1$  regularly-spaced elements from the sorted subarray (the first, the  $(p + 1)$ -th, the  $(2p + 1)$ -th, ..., the last). We call these elements the *primary samples* for that subarray. The  $p$  sets of  $p + 1$  primary samples are then collected and sorted sequentially. The sampling process is applied again, this time to the sequence of sorted primary samples, to obtain a set of  $p + 1$  *secondary samples* which are broadcast to all the processors. The secondary samples partition the original array into  $p$  blocks. Each processor sends all elements of the  $i$ -th block to the  $i$ -th processor. After all the blocks are distributed, the processors sort (sequentially) the received elements. The input array is now sorted in the sense that the first processor holds the first (ordered)



block, the second processor holds the second block, and so on. Optionally, all blocks can be collected in a single processor afterwards.

The test creates an array  $\vec{x}$  containing the first  $n = 2^{24}$  natural numbers and shuffles it randomly. In this way, the algorithm correctness check is simply  $\forall i. 0 \leq i < n : x_i = i$ .

The sequential algorithms simply use the respective language library's *sort* function to sort the input. The input array is represented in the same way in both parallel and sequential applications: in Java it is stored as an `ArrayList<Integer>`, in C++ as a `std::vector<int>`. This is significant in the Java case, since sorting an `ArrayList<Integer>` (using the `Collections.sort()` function) can be up to *ten times* slower with respect to sorting an `int[]` using `Arrays.sort()`<sup>1</sup>. The performance metrics in the Java case are therefore to be interpreted having in mind the fact that the sequential program can be vastly improved by changing the input representation. Conversely, in C++ the difference between sorting a `std::vector<int>` and a `int[]` is negligible.

We could not obtain data for  $p = 64$  for the Java parallel program. The execution with this parallelism degree on the machine detailed in Section 5.2 constantly failed with a `NullPointerException` error encountered when performing communication at the end of a superstep. We did not manage to reproduce the error on the machine used in the development, which used a more recent version of the Java Runtime Environment. As the `MulticoreBSP` for Java library has not been updated in years, it is unknown what causes this problem. As a consequence, in the following performance plots, data for  $p = 64$  is only available for the C++ library.

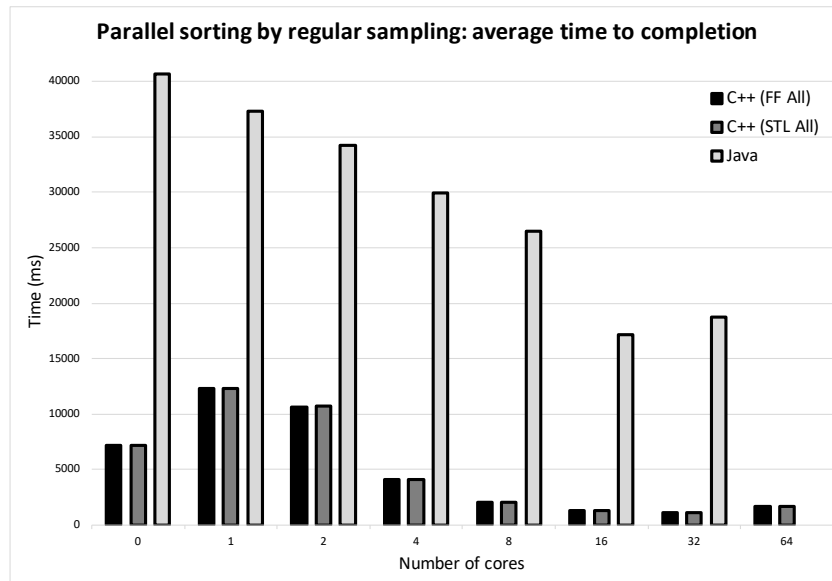


Figure 5.7: Average completion time for the PSRS benchmark on  $2^{24}$  elements.

<sup>1</sup>The bottleneck is not in the sorting function, as `Collections.sort()` actually uses `Arrays.sort()` as the sorting subroutine. To the best of our knowledge, the problem lies in the additional book-keeping needed for maintaining Integer wrappers and `ArrayList` objects.

Figure 5.7 shows the average completion time for the sorting benchmark. The Java implementation actually achieves superlinear speedup for  $p = 1$ : in line with what we said before, this is due to the fact that in the first phase of the algorithm, the input array is distributed in a structure which relies on an `int[]` as the underlying data organization (the `BSP_INT_ARRAY` class of the `MulticoreBSP` for Java library). The other phases of the algorithm use `ArrayList<Integer>`, but at that point the array is at least partially ordered. This may explain the slight advantage in using the parallel program with a single thread over the sequential one. The C++ implementations do not benefit from different data organization, so at  $p = 1$  they show the increased latency due to multiple sorting passes. As the parallelism degree increases, the Java library shows a modest decrement of latency up to  $p = 16$ . The C++ implementations follow a similar trend, instead managing to decrease latency up to  $p = 32$ . Overall, the C++ implementations are remarkably faster than the Java one, with the slab allocator version gaining a very slight edge over the STL allocator one.

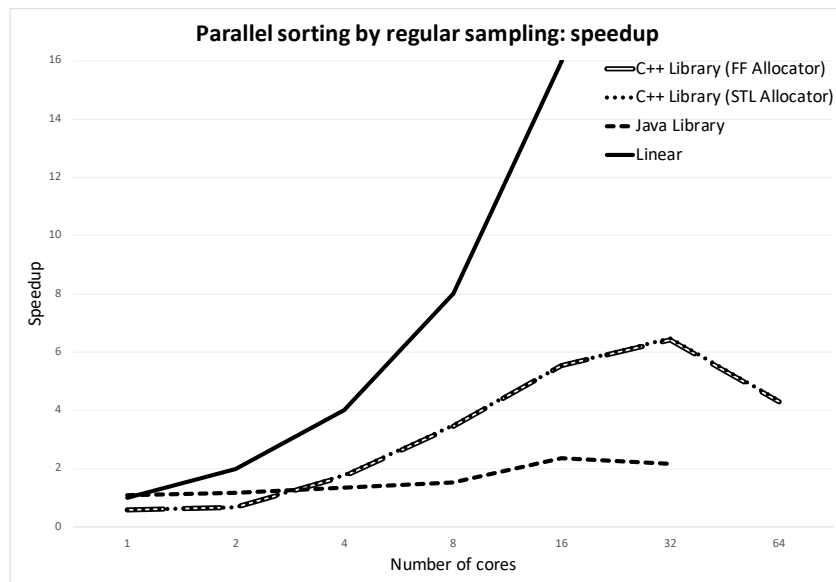


Figure 5.8: Speedup for the PSRS benchmark.

The speedup plot of Figure 5.8 offers an overview of the implementations' behavior for this benchmark. None of them manage to reach linear speedup, except for the Java one at  $p = 1$  that we have discussed before. The two C++ implementations do not show noticeable differences between them; they both achieve better results with regards to the Java implementation for  $p \geq 4$  and perform worse for  $p = 64$  than for  $p = 32$ . The Java implementation shows almost constant speedup throughout the test.

Figure 5.9 refers to the scalability metrics of the benchmark. It is similar to the speedup plot of Figure 5.8 in that no implementation manages to get closer to linear scalability. We notice, although, that unlike in the other metrics




Figure 5.9: Scalability for the PSRS benchmark.

the Java implementation does not perform better than the C++ ones for lower values of  $p$ .




Figure 5.10: Efficiency for the PSRS benchmark.

Lastly, Figure 5.10 show the efficiency plot for this benchmark. As we anticipated when talking about speedup, the Java implementation with one worker

performs slightly better than the sequential program, thus achieving an efficiency greater than 1 (about 1.08). The efficiency rapidly declines as the parallelism degree increases: at  $p = 8$  it reaches 20% of the maximum, and the trend continues only to worsen as  $p$  increases. The C++ implementations immediately sink under 0.5 efficiency for  $p \geq 2$ , but manage to outperform the Java implementation for  $p \geq 4$ . Overall, none of the three libraries provide satisfying performances for this program.




Figure 5.11: The group-cyclic distribution with  $n = 16$ ,  $p = 4$  and  $c = 2$ .

### 5.3.3 Fast Fourier Transform

Periodic functions are encountered in basically every field, from music to medical imaging to telecommunications and more. On computers, these functions can only be represented by the value they take at sample points: a song on an audio CD is sampled 44100 times per second, an MRI image is typically composed by  $512 \times 512$  pixels, and so on. The **Discrete Fourier Transform** (DFT) of a sequence of points sampled from a continuous function converts it into a sequence of points in the frequency domain, which are themselves sample points of the *discrete-time Fourier transform* (DTFT). If the DTFT of a continuous function is known, then – under certain conditions – the original continuous function can be perfectly recovered from it. The DFT is therefore one of the fundamental component in digital signal processing and as such has been extensively studied.

Let  $\vec{x}$  be a vector of  $n$  complex numbers  $\vec{x} = (x_0, \dots, x_{n-1}) \in \mathbb{C}^n$ . The DFT of  $\vec{x}$  is the vector  $\vec{y} = (y_0, \dots, y_{n-1}) \in \mathbb{C}^n$  such that

$$y_k = \sum_{j=0}^{n-1} x_j e^{-2\pi i j k / n} \quad 0 \leq k < n \quad (5.6)$$

The straightforward application of the mathematical definition of DFT over a sequence of  $n$  samples has complexity  $O(n^2)$ . The **Fast Fourier Transform** is a class of algorithms that compute the DFT of a sequence of length  $n$  in time  $O(n \log n)$  instead. The BSPedupack (see Section 5.3) contains an adaptation of the Cooley-Tukey algorithm [18] for calculating the FFT, which has therefore been ported to our C++ library. A detailed explanation of the BSP algorithm is outside the scope of this work, and the interested reader can find it in [9]. Here we limit ourselves to note that the `bspfft` program employs yet another kind of distribution for the input, which in a sense is an intermediate step between the cyclic distribution of Figure 5.1 and the block distribution of Figure 5.6. The *group-cyclic* distribution of  $n$  elements over  $p$  processors with cycle  $c$  is defined as follows: the  $j$ -th element of the input is assigned to the  $k$ -th processor, where

$$k = \left( \left( j \operatorname{div} \left\lceil \frac{cn}{p} \right\rceil \right) c + \left( \left( j \bmod \left\lceil \frac{cn}{p} \right\rceil \right) \bmod c \right) \right) \quad 0 \leq j < n. \quad (5.7)$$

The distribution is defined for every  $c$  such that  $1 \leq c \leq p$  and  $p \bmod c = 0$ . Figure 5.11 shows an example with  $n = 16$ ,  $p = 4$ ,  $c = 2$ . This type of distribution partitions the input into blocks of size  $\lceil cn/p \rceil$ ; each block is then assigned to a group of  $c$  processors according to the cyclic distribution. For  $c = 1$ , the group-cyclic distribution becomes the “normal” block distribution, while for  $c = p$  this reduces to the normal cyclic distribution. The `bspfft` program computes the FFT of a vector by redistributing the input at each step, according to group-cyclic distributions with different parameter  $c$ , and then

computing a sequential FFT over the elements in the same processor. In fact, the redistribution subroutine is the only one that performs communication in this program.

The sequential implementation of the FFT follows the same general behavior of the parallel version, for consistency. Both sequential and parallel version actually execute the algorithm twice (once for the direct FFT, once for the inverse), so the latency for a single pass can be roughly deducted by halving the time spent executing the whole program.

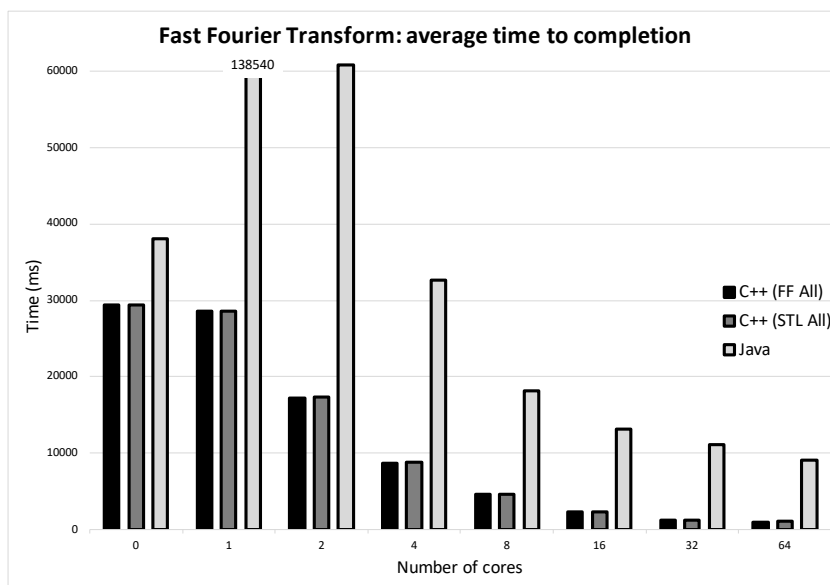


Figure 5.12: Average completion time for the FFT benchmark on  $2^{25}$  elements.

Figure 5.12 shows the completion time chart relative to this benchmark. For  $p = 1$  we notice that the C++ implementations perform faster than the sequential algorithm, albeit only slightly so. As there's no communication between processors even in the parallel version with one worker (since no redistribution ever happens), similar results with regards to the sequential version were expected. Both C++ implementations show a good decrease in latency when the parallelism degree increases. The version that uses the slab allocator generally performs a little better than its STL allocator counterpart. The Java implementation, on the other hand, performs poorly for low values of  $p$ , requiring at least 4 workers to marginally reduce the latency with respect to the sequential program. For  $p \geq 8$  the decrease in latency is acceptable, though the performance of the Java library is constantly inferior to both C++ implementations.

The speedup plot of Figure 5.13 confirms the analysis of the previous paragraph. The Java library never reaches even radical ( $\sqrt{p}$ ) speedup, while the C++ implementations achieve near-linear speedup for  $p \leq 8$  and good results up to the ceiling of  $p = 64$ , with the slab allocator version performing slightly better.

Figure 5.14 offers a surprisingly different scenario for low parallelism degrees: the Java library manages to attain a slightly superlinear scalability up to

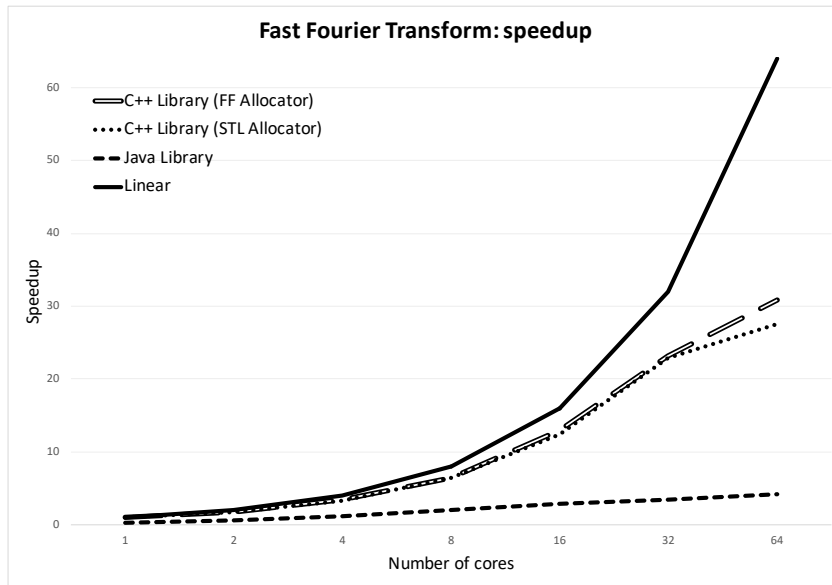


Figure 5.13: Speedup for the FFT benchmark.




Figure 5.14: Scalability for the FFT benchmark.

$p \leq 8$ . This suggests that the poor absolute performance of this library is due to setup costs that get amortized at higher parallelism degrees. Once again, the C++ implementations show good performances for any value of  $p$ , although the scalability increases slowly for  $p \geq 32$ .

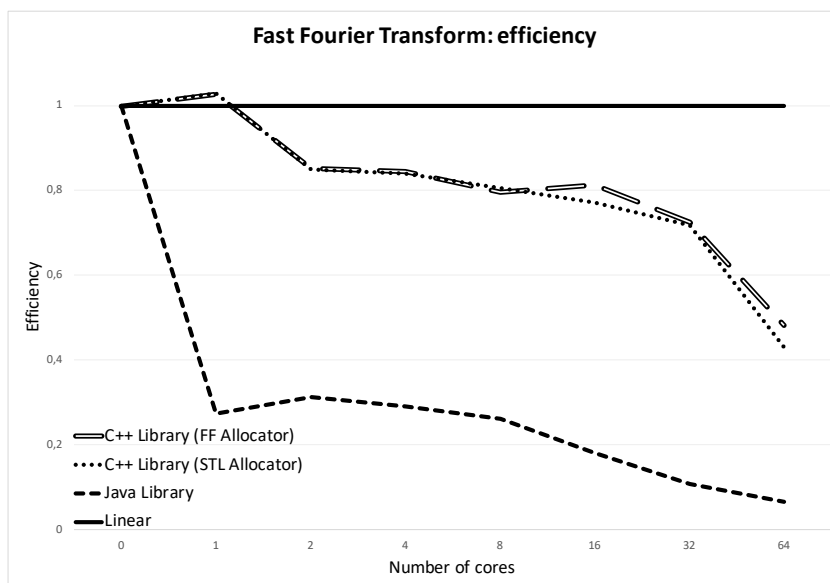


Figure 5.15: Efficiency for the FFT benchmark.

The efficiency plot (Figure 5.15) sums up the observations made in the last few paragraphs. Both C++ versions of the library manage to never fall lower than 0.8 efficiency until  $p = 8$ , with the slab allocator implementation maintaining the advantage over its counterpart throughout all parallelism degrees. The Java library, on the other hand, falls early and – after a modest increase for  $p = 2$  – manages only to maintain a slowly descending rate for the efficiency metrics.

### 5.3.4 Matrix LU decomposition

The last algorithm we consider is another numerical one that is widely used in a broad series of contexts. Let  $A$  be a  $n \times n$  nonsingular matrix and  $\mathbf{b}$ ,  $\mathbf{x}$  be two vectors of length  $n$ , of which the second is not known.  $\mathbf{x}$  is the solution of the linear system

$$A\mathbf{x} = \mathbf{b}. \quad (5.8)$$

One of the possible methods for solving this system is by employing the **LU decomposition** of the matrix  $A$ , i.e. finding two  $n \times n$  matrices  $L$  and  $U$  such that  $L$  is unit lower triangular and  $U$  is upper triangular and

$$A = LU.$$

A  $n \times n$  matrix  $M$  is *unit lower triangular* if  $m_{ii} = 1$  and  $m_{ij} = 0$  for  $i < j$ , and is *upper triangular* instead if  $m_{ij} = 0$  for  $i > j$ .

The matrices  $L$  and  $U$  can be used to find solutions of the system  $A\mathbf{x} = \mathbf{b}$  by solving the triangular systems

$$L\mathbf{y} = \mathbf{b} \quad \text{and} \quad U\mathbf{x} = \mathbf{y} \quad (5.9)$$



since solving triangular systems is easy (in a mathematical sense). The advantage of LU decomposition over other approaches (e.g. Gaussian elimination) is that the  $L$  and  $U$  matrices can be reused when solving a different system such as

$$Ax = \mathbf{b}'.$$

LU decomposition is also used as a step of the solution of various other problems such as inverting a matrix or calculating its determinant.

	0	1	2	0	1
0	00	01	02	00	01
1	10	11	12	10	11
0	00	01	02	00	01
1	10	11	12	10	11
0	00	01	02	00	01

Figure 5.16: Cyclic distribution of a  $5 \times 5$  matrix over  $2 \times 3$  processors.

As with the FFT algorithm, the BSPedupack version of the algorithm was ported to the C++ library implemented in this work. We again point the reader to refer to [9] for a detailed explanation of the algorithm, and focus only on the input distribution and communication patterns. The algorithm actually solves the related problem

$$PA = LU$$

where  $P$  is a permutation matrix that reorders the rows of  $A$ . This family of LU decomposition algorithms are called *with partial pivoting*, since at each step they find *row pivots* that guide the computation by indicating which rows to swap.

Like the vector inner product program, this algorithm makes use of a cyclic distribution. In order to evenly distribute a matrix, though, processors need to be (at least logically) organized in a 2D grid. Each processor is assigned a row index  $s$  and a column index  $t$  and will be therefore referred as  $p_{st}$ . The BSP LU decomposition algorithm distributes the input matrix according to a cyclic distribution over *both* rows and columns. This means that processors  $p_{s*}$  will each hold an element of row  $s$  and processors  $p_{*t}$  will each hold an element of column  $t$ . Figure 5.16 shows how a  $5 \times 5$  matrix is distributed over 6 processors in a  $2 \times 3$  grid.

Both the matrix distribution and the processor organization help in reducing the number of communications at each step. The algorithm requires that every update of an element in the matrix is broadcast to all the other processors, so that they can update their own elements. Since this operation is costly, a *two-phase broadcast* strategy is employed: in the first phase, the node sends each element of the vector it must broadcast to a different intermediate peer. In the second phase, each intermediate node send copies of the received ele-

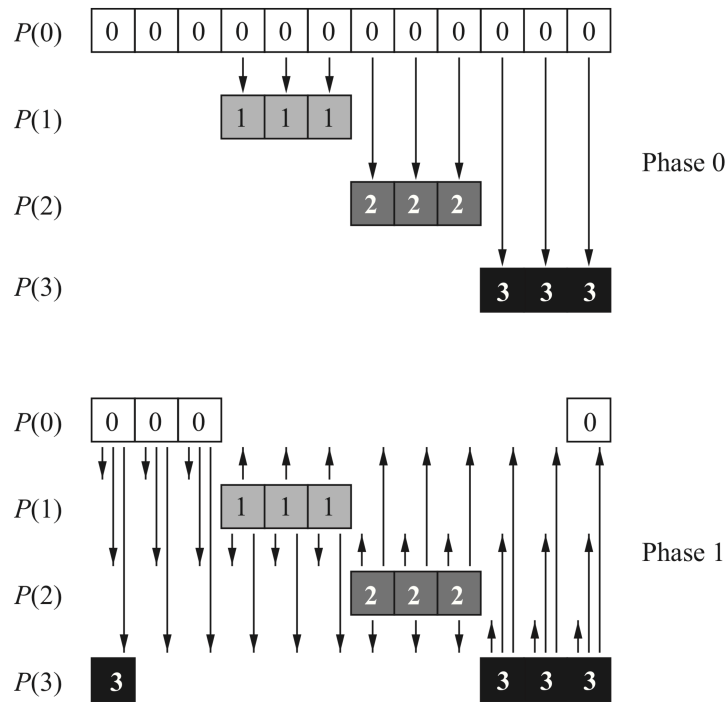


Figure 5.17: A two-phase broadcast of a vector of twelve elements over four processors (taken from [9]).

ment to the final destination. An example of two-phase broadcast is shown in Figure 5.17.

Initially, the sequential version of both the C++ and Java programs employed the sequential Doolittle LU decomposition algorithm [54, 11]. Their performance, however, was much lower than the parallel program using one thread. We therefore decided to use external libraries for this purpose: ALGLIB for the C++ version [34] and Apache Commons Math for the Java version [35].

All tests are run on a  $5000 \times 5000$  input matrix. The parallel versions use the following processor grid configurations:

- $1 \times 1$  ( $p = 1$ );
- $1 \times 2$  ( $p = 2$ );
- $2 \times 2$  ( $p = 4$ );
- $2 \times 4$  ( $p = 8$ );
- $4 \times 4$  ( $p = 16$ );
- $4 \times 8$  ( $p = 32$ );
- $8 \times 8$  ( $p = 64$ ).

The average completion time shown in the chart of Figure 5.18 immediately reveals that this benchmark is by far the heaviest for all the implementations,

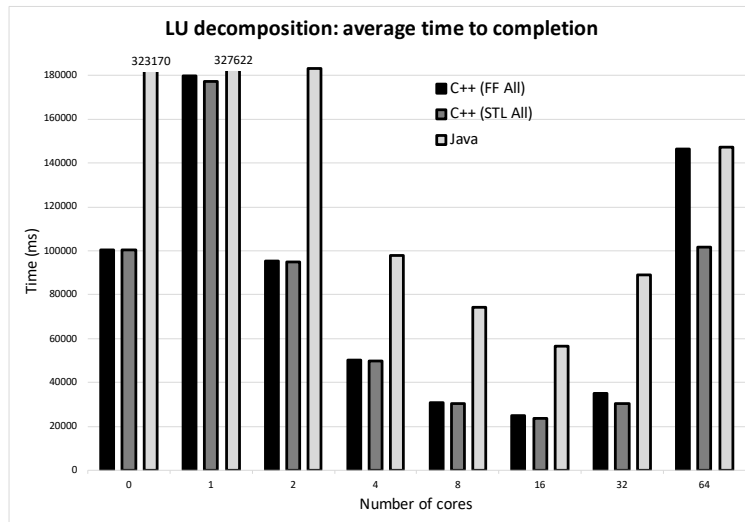


Figure 5.18: Average completion time for the LU decomposition of a  $5000 \times 5000$  matrix.

with latencies that in some cases reached over 350 seconds. The sequential implementation in Java (Apache Commons library) is sensibly less efficient than the C++ one (ALGLIB), with a latency more than three times higher. Both C++ implementations of the parallel library exhibit a higher latency with one worker with respect to the sequential application. The situation improves for higher parallelism degrees up to  $p = 16$ , after which the computation starts to take longer again. The STL allocator version of the C++ library performs better than the slab allocator version, especially for  $p \geq 16$ . The Java parallel implementation with a single worker is nearly on par with the sequential version. After that it behaves like the other two implementations, peaking at  $p = 16$  and degrading for higher parallelism degrees.

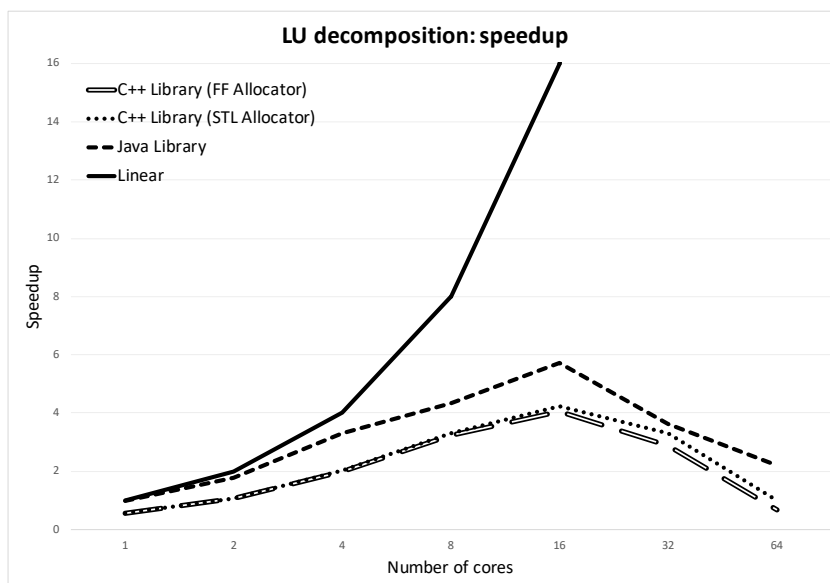


Figure 5.19: Speedup for the matrix LU decomposition benchmark.

The speedup metrics, by definition, strongly depends on the performance of the sequential program. Since the C++ one took less than a third of the time needed for its Java counterpart, it should come to no surprise that the plot of Figure 5.19 shows the Java parallel library on top. All three implementations follow a similar trend, slowly increasing for  $p \leq 16$  and then rapidly decreasing thereafter. The two C++ implementations are basically on par with each other until the peak at  $p = 16$ , after which the slab allocator version falls behind.

Figure 5.20 shows the scalability plot. Since the Java library performs nearly the same as the sequential version for  $p = 1$ , its scalability is practically on par with its speedup. Conversely, the C++ implementations fare better here, as they are not compared with the highly efficient ALGLIB sequential program anymore.

Finally, the efficiency for this benchmark can be found in Figure 5.21. The same considerations we made for the speedup apply here: the Java implementation “looks better” due to the relatively poor performance of the sequential version. The C++ library variants behave the same, with the STL allocator having a slight edge. All three implementations reach less than 0.1 efficiency for  $p \geq 32$ . This poor performance may be explained by the exponentially higher amount of communication needed for the algorithm as  $p$  increases, especially due to the two-phase broadcast detailed before.




Figure 5.20: Scalability for the matrix LU decomposition benchmark.




Figure 5.21: Efficiency for the matrix LU decomposition benchmark.

### 5.3.5 Other programs

Some other programs were implemented for this work. The BSPedupack contains a “benchmark” of sorts, called `bspbench`, that stresses both communication and computation aspects of the machine on which it runs. It can also be used to derive the parameters  $g, h, L$  of the BSPC implemented in the library. We provided this program (`BSPbench.cpp`), along with a custom-designed communication stress test (`commstress.cpp`) and a battery of unit tests (`unit_tests.cpp`), the latter two for checking the correctness of the implementation.

The `bspbench` implementation reported the following values on the machine used for the performance evaluation.

<b>p</b>	<b>r<sup>A</sup></b>	<b>g</b>	<b>l</b>
1	1170.29	597.108	3666.18
2	1170.29	868.933	35265.2
4	1170.29	3591.81	91652.2
8	1170.29	6183.59	215230
16	1170.29	12593.1	465453
32	1170.29	51731.2	1916450

<sup>A</sup> Mflops/s

Figure 5.22: BSP parameters for the Phi machine as reported by `bspbench`.

## 5.4 Summary

In this Chapter we measured the performance of the library implemented in this work on a variety of commonly-used applications, both numerical and non-numerical. The experimental results are very promising: while the absolute performance depends on the algorithm, the BSP implementations that used our library showed a decent-to-good increase in performance at any parallelism degree. The other main “competitor” in the field of object-oriented modern BSP implementations, the MulticoreBSP for Java library, is consistently outranked by our solution. Although the native difference in performance between the Java and C++ languages helped achieving good results, we consider our implementation choices to have played a bigger role in determining the final difference in performance.

## Chapter 6

# Conclusion

In this thesis we introduced the design and implementation of a C++ library for the Bulk Synchronous Parallel paradigm, targeting shared memory multicore machines. The architectural design draws its inspiration from Tiskin’s BSPRAM model [55], but maintains the communication and the shared memory as separate concepts. The implementation uses the compositional capabilities of the FastFlow framework to build its internal structure, which happens to be a plain farm template, actually. The library can also be considered an extension of FastFlow, as the whole BSP computation ultimately takes place in a FastFlow node which can be freely used inside any other parallel pattern.

Both the design and implementation aspects of the library varied a lot during the development of this work. At first, we thought about providing an all-new programming style for BSP computations: each superstep would be represented as a C++ object `superstep<Tin, Tout>`, with clearly defined input and output types. These objects were to be given, in the desired order, to a BSP “executor” that would regulate the superstep flow and perform the requested communications. We soon realized that this programming style was not necessarily easier to use, instead proving to be hugely limiting on the communication capabilities (each node could only send either a single object or a collection of same-type objects per superstep).

After this “superstep-centric” draft, we briefly reasoned upon an architecture where nodes were arranged in a fully or partially connected mesh, in order to remove the need for an entity that manages communications. The channels between nodes were to be implemented using FastFlow SPSC queues (see Section 2.2). This approach was quickly dismissed, as a fully connected mesh would require a huge number of queues even for modest parallelism degrees (with 64 nodes there need to be 4096 channels), and a partially connected mesh would not improve communication times too much. Moreover, this approach didn’t fully utilize FastFlow’s parallel patterns, instead relying on its low-level mechanisms. While this is certainly allowed by the framework, we wanted to exploit the easiness of use provided by its algorithmic skeletons.

Finally, we converged to a satisfying architectural model and implementation scheme, the ones we discussed in Chapters 3 and 4. Of course there was still a lot to improve: particular focus was given on finding the most efficient representation of shared objects in the memory, due to the difficulties of having to work with heterogeneous data to be stored inside a single container.

Another aspect that required particular attention was the requests mechanism, as it put a lot of stress on the number of end-of-superstep operations to be executed, leading to a lower performance. A common pattern we noticed in many BSP computations was that nodes often performed a high number of communication requests per superstep, but each communication only regarded one or few data elements (which, we remind, must be copied into the request object). The communicator therefore had to continuously allocate and deallocate many small request objects. We reasoned that the standard C++ allocator could be less efficient in this scenario, and sought after different allocation schemes which better suited our needs. Fortunately, FastFlow provided a custom solution based on the slab allocator (Section 2.2.1), which was therefore used in the library.

At the end of this implementation phase we obtained a final product that matched the goals we set for ourselves. The library provides support for BSP programs with a simple and clean API. The run-time support relieves users from having to manage memory operations and shared variable registration and deregistration, unlike many other BSP libraries. Programmers can code in a lean and modern way, relying on the object-oriented paradigm and supporting all C++11 (and above) features. The BSP computation can be performed over any type of data, including user-defined classes and objects. As we will see in a moment, the library's performance is sound, as it shows good scalability for basically any parallelism degree. The library is also integrated in the FastFlow framework, as the BSP computation can use input data received from other nodes and can send output to them.

We extensively tested our solution, both during and after the final stages of development. For conducting performance analysis, we chose four BSP algorithms to be used as benchmark and implemented them. Each one of those programs feature a different communication scheme and a different partitioning of the input between nodes. Three of the four programs are taken from Bisseling's BSPedupack [9], a collection of BSP algorithms for solving common numerical problems. The last one is a sorting algorithm, a crucial operation in basically any field of computer science. Each program was run multiple times using a different number of worker units, to check how the library performed at different parallelism degrees. In addition to the "normal" library, we also ran tests over a variant of the final implementation that used the standard C++ STL allocator, to test if the slab allocator effectively improved performance. We compared our results with identical tests for the MulticoreBSP for Java library. We used the BSPedupack for Java implementation of the three edupack programs and wrote our own BSP sorting program for Java.

The experimental results proved that our implementation is sound. We consistently achieved lower latencies than the Java library, with a large enough margin that allows us to ignore the inherent gap between the two languages. In some cases, our implementation was over ten times faster than the equivalent Java version. With few exceptions, our implementation also presented better scalability for any parallelism degree up to the machine's number of cores (in this case, 64). Lastly, we obtained good results also for the speedup and efficiency parameters, after factoring that C++ sequential versions of the programs were faster than Java ones, therefore "polluting" these two last performance indices. To our surprise, both the STL allocator version and the slab (FastFlow) allocator one produced similar results. We expected the slab alloca-






Figure 6.1: FastFlow graph of a possible implementation of the MultiBSP model. The solid line nodes compose the MultiBSP tree.

tor version to gain a more consistent edge in the abovementioned scenario of many requests with small payloads.

## 6.1 Future work

As mentioned above, the fact that the STL allocator provide similar results to the slab allocator (even outperforming it in some scenarios) can be studied more in-depth. A reasonable approach would be to allow the programmer to choose the desired allocation strategy, potentially even letting him/her specify a custom allocator.

An interesting direction for future work is the implementation of the MultiBSP model (introduced in [60]) via exploitation of FastFlow’s compositional capabilities. Since the whole BSP computation is already a FastFlow node, a nested BSP run can be achieved simply by allowing the computation node to act as a BSP node, therefore becoming part of the broader composition. The MultiBSP tree in this case corresponds to the FastFlow graph, minus the farm emitter and collector entities (Figure 6.1).

We could, lastly, exploit FastFlow’s support for different architectures in order to provide other BSP variants. For example, we can use FastFlow `ff_dnodes` [4] to run BSP programs on a distributed computer, therefore “closing the circle” by providing support for the same architecture described in Valiant’s seminal work. Another interesting topic, in the same vein, could be to provide support for offloading work to GPU accelerators. Both these directions also work well within the abovementioned MultiBSP paradigm.

While these features weren’t implemented due to time constraints, we ultimately feel that they concern a slightly broader scope than the one we con-

sidered in this work, namely to realize an efficient and simple-to-use BSP library for shared-memory multicores integrated in the FastFlow framework. We think that the implementation we provided in this thesis matched our goals.

# Appendix A

## Library quick start


This Appendix details the API documentation for the library, i.e. public methods and fields that the programmer can use to build a BSP program. A complete documentation, which includes private and internally-used methods and fields, is available as Doxygen comments in the source code. The closing Sections provide additional clarifications about some mechanisms of the `bsp_variable` and `bsp_array` classes.

### A.1 API documentation

#### A.1.1 `bsp_program` Class Reference

```
#include <bsp_program.hpp>
```

Inheritance diagram for `bsp_program`:



#### Public Member Functions

- `bsp_program(std::vector<std::unique_ptr<bsp_node>>&& _processors, std::function<void(void)> _pre = nullptr, std::function<void(void)> _post = nullptr)`
- `void start(void* in = nullptr)`
- `void* svc(void* in) override`

#### Class Description

Implements the Bulk Synchronous Parallel pattern as a FastFlow node.

## Constructor Documentation

### bsp\_program()

```
explicit bsp_program(std::vector<std::unique_ptr<bsp_node>>&& _processors,  
                    std::function<void(void)> _pre = nullptr,  
                    std::function<void(void)> _post = nullptr)
```

Constructor for the bsp\_program object.

### Parameters

_processors	vector of BSP nodes for the computation
_pre	optional function to be executed before the BSP computation
_post	optional function to be executed after the BSP computation

## Member Function Documentation

### start(void\*)

```
void start(void* in = nullptr)
```

Creates the FastFlow inner graph and executes the BSP computation.

### Parameters

in	optional, the input token received from the preceding FastFlow node
----	---

---

### svc(void\*)

```
void* svc(void* in) override
```

Service function for the BSP program node. Starts the BSP computation.

### Parameters

in	input token received from the preceding FastFlow node
----	---

### Returns


GO\_ON, a special FastFlow token to continue the broader computation

---

## A.1.2 bsp\_node Class Reference

```
#include <bsp_node.hpp>
```

Inheritance diagram for bsp\_node:



## Public Member Functions

- `void*` `svc(void* in)` `final`

## Protected Member Functions

- `void` `emit_output(void* payload)`
- `template <typename T>`  
`bsp_variable<T> get_variable(const T& initial_value)`
- `template <typename T>`  
`bsp_array<T> get_array(const std::vector<T>& initial_value)`
- `template <typename T>`  
`bsp_array<T> get_array(std::vector<T>* handle)`
- `template <typename T>`  
`bsp_array<T> get_empty_array(int size)`
- `int` `bsp_pid()`
- `int` `bsp_nprocs()`
- `void` `bsp_sync()`
- `virtual void` `parallel_function() = 0`

## Protected Attributes

- `const void*` `fastflow_input`  
*Pointer to the FastFlow input token*

## Class Description

Specialization of a `ff_node` to work as the unit of computation in the BSP model.

## Member Function Documentation

### `bsp_nprocs()`

`int` `bsp_nprocs()`

Returns the number of nodes in the current BSP computation.

### Returns

the number of nodes in the current BSP computation.

---

### `bsp_pid()`

`int` `bsp_pid()`

Returns the ID for this node.

### Returns

this node's ID.

---

### bsp\_sync()

```
void bsp_sync()
```

Terminates the current superstep and waits for the other nodes to sync.

---

### emit\_output(void\*)

```
void emit_output(void* payload)
```

Forwards an output token to the next stage in the FastFlow graph.

### Parameters

payload	the token to forward
---------	----------------------

---

### get\_array(const std::vector<T>&)

```
template <typename T>
bsp_array<T> get_array(const std::vector<T>& initial_value)
```

Requests a new array with elements of type T from the communicator, initializing it with the copy a given vector.

### Template Parameters

T	the type of elements of the requested array
---	---

### Parameters

initial_value	value to copy inside this node's private copy of the array
---------------	--

### Returns

a handle to this node's private copy of the shared array

---

### get\_array(std::vector<T>\*)

```
template <typename T>
bsp_array<T> get_array(std::vector<T>* handle)
```

Requests a new array with elements of type T from the communicator, initializing it with the pointer to a vector. Any modifications done to the initializing vector after the call to this method are inherently BSP unsafe.

### Template Parameters

T	the type of elements of the requested array
---	---

### Parameters

handle	a pointer to the vector that will become this node's private copy of the shared array
--------	---

### Returns

a handle to this node's private copy of the shared array

---

### get\_empty\_array(int)

```
template <typename T>
bsp_array<T> get_empty_array(int size)
```

Requests a new array with elements of type T from the communicator, initializing it with an empty vector of given size.

### Template Parameters

T	the type of elements of the requested array
---	---

### Parameters

size	the size of the empty vector that will become this node's private copy of the shared array
------	--

### Returns

a handle to this node's private copy of the shared array

---

### get\_variable(const T&)

```
template <typename T>
bsp_variable<T> get_variable(const T& initial_value)
```

Requests a new variable of type T from the communicator.

### Template Parameters

T	the type of the requested variable
---	------------------------------------

### Parameters

initial_value	value to copy inside this node's private copy of the variable
---------------	---

### Returns

a handle to this node's private copy of the shared variable

---

### parallel\_function()

```
virtual void parallel_function() = 0;
```

Function to be overwritten as the main parallel execution.

---

### svc(void\*)

```
void* svc(void* in) final
```

The FastFlow node service method. Implementations of this class cannot redefine it.

### Parameters

in	the input token (in this case, a special value ENDCOMP)
----	---

### Returns


the same token as the input

---

## A.1.3 bsp\_variable<T> Template Class Reference

```
#include <bsp_variable.hpp>
```

Inheritance diagram for bsp\_variable<T>:



### Public Member Functions

- T get()
- void bsp\_put(const T& elem, int destination)
- void bsp\_put(const bsp\_variable<T>& other)
- void bsp\_put(const bsp\_variable<T>& other, int destination)
- void bsp\_put(int destination)
- void bsp\_get(int source) **override**
- T bsp\_direct\_get(int source)
- T& BSPunsafe\_access()

### Class Description

A variable data structure, private to a single worker but with support for communication with other similar entities.

### Template Parameters

T	type of the variable
---	----------------------

### Member Function Documentation

#### bsp\_direct\_get(int source)

T bsp\_direct\_get(int source)

Returns a copy of a node's private element of this shared variable. This method will return immediately, without waiting for a superstep sync.



### Parameters

source	ID of the source node
--------	-----------------------

### Returns

a copy of the desired element

---

### **bsp\_get(int source)**

**void** bsp\_get(**int** source) **override**

Replaces this node's private element with another node's private element of this shared variable.

### Parameters

source	ID of the source node
--------	-----------------------

---

### **bsp\_put(const bsp\_variable<T>&)**

**void** bsp\_put(**const** bsp\_variable<T>& other)

Replaces this node's private element with this node's private element of another shared variable.

### Parameters

other	the handle to the other shared variable
-------	---

---

### **bsp\_put(const bsp\_variable<T>&, int)**

**void** bsp\_put(**const** bsp\_variable<T>& other, **int** destination)

Replaces another node's private element with this node's private element of another shared variable.

### Parameters

other	the handle to the other shared variable
destination	ID of the destination node

---

### **bsp\_put(const T&, int)**

**void** bsp\_put(**const** T& elem, **int** destination)

Replaces another node's private element of the shared variable.

### Parameters

elem	element to be copied
destination	ID of the destination node

---

### **bsp\_put(int)**

`void bsp_put(int destination)`

Replaces another node's private element with this node's private element of the shared variable.

#### **Parameters**

destination	ID of the destination node
-------------	----------------------------

---

### **BSPunsafe\_access()**

`T& BSPunsafe_access()`

Returns a handle to this node's private element of the shared variable. The returned object can be modified at will, without waiting for superstep syncs. This function is **BSP unsafe**.

#### **Returns**

a reference to the node's private element of the shared variable

---

### **get()**

`T get()`

Returns a copy of the node's private element of the shared variable.

#### **Returns**


a copy of the desired element

---

## **A.1.4 bsp\_array<T> Template Class Reference**

`#include <bsp_array.hpp>`

Inheritance diagram for `bsp_array<T>`:



## Public Member Functions

- `T get()`
- `std::vector<T> get()`
- `int size()`
- `void BSPunsafe_put(const T& elem, int pos)`
- `void bsp_put(const T& elem, int pos)`
- `void bsp_put(const bsp_variable<T>& elem, int pos)`
- `void bsp_put(const T& elem, int destination, int pos)`
- `void bsp_put(const bsp_variable<T>& elem, int destination, int pos)`
- `void bsp_put(const std::vector<T>& array)`
- `void bsp_put(const bsp_array<T>& other)`
- `void bsp_put(const std::vector<T>& array, int destination)`
- `void bsp_put(const bsp_array<T>& other, int destination)`
- `void bsp_put(const std::vector<T>& array, int src_offset, int dest_offset, int length)`
- `void bsp_put(const bsp_array<T>& other, int src_offset, int dest_offset, int length)`
- `void bsp_put(const std::vector<T>& array, int destination, int src_offset, int dest_offset, int length)`
- `void bsp_put(const bsp_array<T>& other, int destination, int src_offset, int dest_offset, int length)`
- `void bsp_get(int from) override`
- `void bsp_get(int from, int from_offset, int to_offset, int length)`
- `T bsp_direct_get(int source, int pos)`
- `std::vector<T> bsp_direct_get(int source)`
- `std::vector<T>& BSPunsafe_access()`

## Class Description

An array structure, private to a single worker but with support for communication with other similar entities.

## Template Parameters

<b>T</b>	type of the elements contained in the array
----------	---

## Member Function Documentation

### `bsp_direct_get(int)`

`std::vector<T> bsp_direct_get(int source)`

Returns a copy of another node's private array. This method will return immediately, without waiting for a superstep sync.

## Parameters

<code>source</code>	ID of the source node
---------------------	-----------------------

### Returns

the source's private copy of the shared array

---

### **bsp\_direct\_get(int, int)**

T bsp\_direct\_get(int source, int pos)

Returns a copy of the element at the desired position from another node's copy of the array. This method will return immediately, without waiting for a super-step sync.

### Parameters

source	ID of the source node
pos	position of the desired element

### Returns

the element of the source's private copy at the desired position

---

### **bsp\_get(int)**

void bsp\_get(int from) override

Replaces the node's private copy of the array with another node's private copy.

### Parameters

from	ID of the source node
------	-----------------------

---

### **bsp\_get(int, int, int, int)**

void bsp\_get(int from, int from\_offset, int to\_offset, int length)

Replaces a portion of the node's private copy of the array with a portion of another node's private copy.

### Parameters

from	ID of the source node
from_offset	starting position of the replacing portion
to_offset	starting position of the portion to be replaced
length	length of the replacing portion

---

### **bsp\_put(const bsp\_array<T>&)**

void bsp\_put(const bsp\_array<T>& other)

Replaces the node's private copy of the array with the vector contained inside another bsp\_array.

### Parameters

other	other bsp_array containing the vector that will replace the current one
-------	---

---

### bsp\_put(const bsp\_array<T>&, int)

**void** bsp\_put(**const** bsp\_array<T>& other, **int** destination)

Replaces another node's private copy of the array with this node's private copy of the vector contained inside another bsp\_array.

### Parameters

other	other bsp_array containing the vector that will replace the current one
destination	ID of the destination node

---

### bsp\_put(const bsp\_array<T>&, int, int, int, int)

**void** bsp\_put(**const** bsp\_array<T>& other, **int** destination, **int** src\_offset, **int** dest\_offset, **int** length)

Replaces a portion of another node's private copy of the array with a portion of this node's private copy of the vector contained inside another bsp\_array.

### Parameters

other	other bsp_array containing the vector with the desired portion
destination	ID of the destination node
src_offset	starting position of the replacing portion
dest_offset	starting position of the portion to be replaced
length	length of the replacing portion

---

### bsp\_put(const bsp\_array<T>&, int, int, int)

**void** bsp\_put(**const** bsp\_array<T>& other, **int** src\_offset, **int** dest\_offset, **int** length)

Replaces a portion of this node's private copy of the array with a portion of this node's private copy of the vector contained inside another bsp\_array.

### Parameters

other	other bsp_array containing the vector with the desired portion
src_offset	starting position of the replacing portion
dest_offset	starting position of the portion to be replaced
length	length of the replacing portion

**bsp\_put(const bsp\_variable<T>&, int, int)**

**void** bsp\_put(**const** bsp\_variable<T>& elem, **int** destination, **int** pos)

Puts an element (contained in a bsp\_variable) into the given position in another node's private copy of the array.

**Parameters**

elem	bsp_variable that contains the element to be inserted
destination	ID of the destination node
pos	position in the array

---

**bsp\_put(const bsp\_variable<T>&, int)**

**void** bsp\_put(**const** bsp\_variable<T>& elem, **int** pos)

Puts an element (contained in a bsp\_variable) into the given position in this node's private copy of the array.

**Parameters**

elem	bsp_variable that contains the element to be inserted
pos	position in the array

---

**bsp\_put(const std::vector<T>&)**

**void** bsp\_put(**const** std::vector<T>& array)

Replaces the node's private copy of the array with the given vector.

**Parameters**

array	vector that will replace the current private copy of the shared array
-------	---

---

**bsp\_put(const std::vector<T>&, int)**

**void** bsp\_put(**const** std::vector<T>& array, **int** destination)

Replaces another node's private copy of the array with the given vector.

**Parameters**

array	vector that will replace the current private copy of the shared array
destination	ID of the destination node

---

**bsp\_put(const std::vector<T>&, int, int, int, int)**

```
void bsp_put(const std::vector<T>& array, int destination,
            int src_offset, int dest_offset, int length)
```

Replaces a portion of another node's private copy of the array with a portion of the given vector.

**Parameters**

array	vector that will replace the current private copy of the shared array
destination	ID of the destination node
src_offset	starting position of the replacing portion
dest_offset	starting position of the portion to be replaced
length	length of the replacing portion

---

**bsp\_put(const std::vector<T>&, int, int, int)**

```
void bsp_put(const std::vector<T>& array, int src_offset,
            int dest_offset, int length)
```

Replaces a portion of this node's private copy of the array with a portion of the given vector.

**Parameters**

array	vector that will replace the current private copy of the shared array
src_offset	starting position of the replacing portion
dest_offset	starting position of the portion to be replaced
length	length of the replacing portion

---

**bsp\_put(const T&, int, int)**

```
void bsp_put(const T& elem, int destination, int pos)
```

Puts an element into the given position in another node's private copy of the array.

**Parameters**

elem	element to be inserted
destination	ID of the destination node
pos	position in the array

---

**bsp\_put(const T&, int)**

```
void bsp_put(const T& elem, int pos)
```

Puts an element into the given position in this node's private copy of the array.

### Parameters

elem	element to be inserted
pos	position in the array

---

### BSPunsafe\_access()

```
std::vector<T>& BSPunsafe_access()
```

Returns a handle to this node's private element of the shared array. The returned object can be modified at will, without waiting for superstep syncs. This function is **BSP unsafe**.

### Returns

a reference to the node's private copy of the array

---

### BSPunsafe\_put(const T&, int)

```
void BSPunsafe_put(const T& elem, int pos)
```

Puts an element into the contained array in the desired position. This function is **BSP unsafe**.

### Parameters

elem	element to be inserted
pos	position in the array

---

### get()

```
std::vector<T> get()
```

Returns a copy of the node's private element of the shared array.

### Returns

a copy of the array

---

### get(int)

```
T get(int position)
```

Returns a copy of the element in a given position of the node's private copy of the shared array.

### Parameters

position	position in the array
----------	-----------------------



**Returns**

a copy the array element in the desired position

---

**size()**

```
int size()
```

Returns the size of the array.

**Returns**

the size of the array

## A.2 Variable/array request mechanism

The main way to perform communication between BSP processors is using `bsp_variables` (or their specialization `bsp_arrays`). `bsp_variable<T>`s are special containers that hold an element of type `T` for each BSP node, i.e. each node will refer to a different - private - element using the same `bsp_variable` handle.

`bsp_arrays` are specializations of `bsp_variables` for variable-sized array types. A `bsp_array<T>` is basically equivalent to a `bsp_variable<std::vector<T>>`, except it also provides functions to efficiently get and put single elements or portions of the array instead of having to work with the whole object. In the following, any mention of `bsp_variable` mechanisms is also valid for `bsp_arrays`, unless where noted.

`bsp_variable` objects cannot be built using the standard initialization techniques; instead, they are created by the appropriate `get_variable<T>` method of the `bsp_node` class. In reality, `bsp_variable` objects are mere handles to the actual containers, which have different lifecycles. A new underlying container for variables of type `T` is created by the system when a BSP node requests one more `bsp_variable<T>` than the other BSP nodes; when the other BSP nodes "catch up" and make one more request for another `bsp_variable<T>`, the system gives them a handle to this freshly-created container. In short, for all nodes it holds the following:

**Proposition 2.** The `bsp_variable<T>` object obtained by the  $n$ -th consecutive call to `get_variable<T>` in a given superstep is a handle for the  $n$ -th container of objects of type `T`, regardless of everything else such as variable names.

The following example (listing A.1) helps understand these mechanisms. Suppose we have a BSP program with two nodes that execute this program:

```
void parallel_function() override {
    int id = bsp_pid();
    if (id == 0) {
        auto v1 = get_variable<int>(0);
        sleep(1000);
        auto v2 = get_variable<int>(0);
        auto v3 = get_variable<double>(0.0);
    }
    if (id == 1) {
        sleep(100);
        auto v1 = get_variable<double>(1.1);
        auto v2 = get_variable<int>(1);
        auto v3 = get_variable<int>(1);
    }
    [...]
}
```

Listing A.1: The code for the example.

The nodes request the same number of variables for each type, but in a different order. (The `sleep` functions are to emphasize the order of execution.) What happens is that the two nodes will call the same shared object with different names.

1. Node 0 requests a variable of type `int`. This is request number 1 for this node and this type, but globally the system has reserved 0 containers for type `int`. The system thus creates a container for an `int` variable: this container will reserve space for two `ints`. The system then sets the first `int` to 0 (as requested by Node 0) and returns a `bsp_variable<int>`, which is an handle to the aforementioned container and in particular to the element assigned to Node 0. This `bsp_variable<int>` has the name `v1` in the environment of Node 0.
2. Node 1 requests a variable of type `double`. The system creates a new container of `double` elements, sets the element corresponding to Node 1 to 1.1, and returns the handle to this element to Node 1 as a `bsp_variable<double>` named `v1`. Note that Node 0's `v1` and Node 1's `v1` point to two different containers of different types.
3. Node 1 requests a variable of type `int`. This is the first request done by Node 1 for such a variable, and globally the system has already reserved a container of type `int`. No new containers are created: the second `int` of the existing container is set to 1, as requested by Node 1, and a `bsp_variable<int>` which refers to this element is returned to Node 1. This `bsp_variable<int>` has the name `v2` in the environment of Node 1, but it actually refers to the same container as Node 0's `v1`! If Node 0 were to call `v1.bsp_put(5, 1)` (see next subsection), this operation would actually change the element referred by Node 1 with the name `v2` after a superstep sync (i.e. `v2.get()` would return 5 after a superstep sync).

4. Node 1 requests a variable of type `int`. This is the second request for such a variable, but globally there's only one `int` container, so the system creates a new one. Node 1's handle for this container is a `bsp_variable<int>` named `v3`.
5. Node 0 requests a variable of type `int`. It's the second request for such a variable and there are already two containers of type `int`, so the system returns a handle to the second container, which is a `bsp_variable<int>` named `v2`.
6. Node 0 requests a variable of type `double`. It's the first request for such a variable and there's already one container of type `double`, so the system returns a handle to this container, which is a `bsp_variable<double>` named `v3`. Node 0's `v3` and Node 1's `v1` refer to the same container.




Figure A.1: An example of the difference between `bsp_variable`s and their underlying containers.

### A.3 A note on BSPunsafe methods

Both `bsp_variable` and `bsp_array` classes feature a few methods, identified by the prefix `BSPunsafe`, that allows the user to directly access the node's private copy of the shared object. Another way to gaining "direct access" to a `bsp_array` is by requesting it via the `get_array<T>(pointer)` method; the vector pointed by `pointer` is the same entity as the element in the container. This type of access to data is "BSP-unsafe" in the sense that it doesn't fit well with BSP directives (`get` and `put`). For example, if node  $n$  modifies its element with `BSPunsafe_access()` during superstep  $k$  and another node  $m$  performs a `bsp_put()` on the same element in the same superstep, at the beginning of superstep  $k + 1$  the value held in  $n$ 's element will be the one written by  $m$  (as `bsp_put()` writes are performed at the end of the superstep, overwriting any modification previously made). If the data will be used in a read-only way, it's advisable to get a copy using the `bsp_direct_get()` methods.

The BSP-unsafe methods can nevertheless be useful when properly managed; some of the test programs (i.e. the Fast Fourier Transform and LU decomposition) use them for easier access to private copies, saving some superstep synchronizations. Other BSP implementations (e.g. `MulticoreBSP` for Java) also provide similar methods.

# Appendix B

## Source code

### B.1 BSP Library for FastFlow

#### bsp\_array.hpp

```
1  #ifndef FF_BSP_BSP_ARRAY_HPP
2  #define FF_BSP_BSP_ARRAY_HPP
3
4  #include <type_traits>
5  #include <vector>
6  #include "bsp_variable.hpp"
7
8  /**
9   * An array structure, private to a single worker but with support
10  * for communication with other similar entities.
11  *
12  * @tparam T type of the elements contained in the array.
13  */
14  template<typename T>
15  class bsp_array : bsp_container {
16
17      // The template type must be copy-constructible and copy-assignable
18      static_assert(std::is_copy_constructible<T>::value &&
19                  std::is_copy_assignable<T>::value,
20                  "Type of elements of bsp_array must be copy-constructible "
21                  " and copy-assignable");
22
23  private:
24
25      // bsp_communicator can access private fields of this class
26      friend class bsp_communicator;
27
28      /**
29       * Pointer to the actual vector data.
30       */
31      std::vector<T>* arr;
32
33      /**
34       * Default constructor.
35       */
36      bsp_array() = default;
37
38      /**
```

```

39     * Builds a bsp_array object that holds a node's private copy of a shared
40     * array.
41     *
42     * @param _ref ID of the shared array
43     * @param _hold ID of the requesting node
44     * @param comm pointer to the communicator component
45     * @param ptr pointer to the node's private copy
46     */
47     bsp_array(int _ref, int _hold, bsp_communicator* comm,
48             std::vector<T>* ptr) :
49             bsp_container(_ref, _hold, comm), arr{ptr} {
50     }
51
52     /**
53     * Returns the shared object type (in this case, an array).
54     * @return the C++ vartype value for arrays
55     */
56     vartype var_type() final {
57         return vartype::array;
58     }
59
60     public:
61
62     /**
63     * Returns a copy of the element in position <tt>position</tt> of the
64     * node's private copy of the shared array.
65     * @param position position of the desired element
66     * @return a copy of the array element in the desired position
67     */
68     T get(int position) {
69         return arr->at(position);
70     }
71
72     /**
73     * Returns a copy of the node's private copy of the shared array.
74     * @return a copy of the array
75     */
76     std::vector<T> get() {
77         return *arr;
78     }
79
80     /**
81     * Returns the size of the array.
82     * @return the size of the array
83     */
84     int size() {
85         return arr->size();
86     }
87
88     /**
89     * Puts an element into the contained array in the desired
90     * position. This function is <b>BSP unsafe</b>.
91     * @param elem element to be copied
92     * @param pos position in the array
93     */
94     void BSPunsafe_put(const T& elem, int pos) {
95         arr->at(pos) = elem;
96     }

```

```

97
98  /**
99   * Puts an element into the given position in this node's private copy of
100  * the array.
101  * @param elem element to be inserted
102  * @param pos position in the array
103  */
104  void bsp_put(const T& elem, int pos) {
105      bsp_put(elem, holder, pos);
106  }
107
108  /**
109   * Puts an element (contained in a <tt>bsp_variable</tt>) into the given
110   * position in this node's private copy of the array.
111   * @param elem |c bsp_variable that contains the element to be inserted
112   * @param pos position in the array
113   */
114  void bsp_put(const bsp_variable<T>& elem, int pos) {
115      const auto& t = *(elem.element);
116      bsp_put(t, pos);
117  }
118
119  /**
120   * Puts an element into the given position in another node's private copy of
121   * the array.
122   * @param elem element to be inserted
123   * @param destination ID of the destination node
124   * @param pos position in the array
125   */
126  void bsp_put(const T& elem, int destination, int pos) {
127      comm->array_put(reference, holder, destination, pos, elem);
128  }
129
130  /**
131   * Puts an element (contained in a <tt>bsp_variable</tt>) into the given
132   * position in another node's private copy of the array.
133   * @param elem |c bsp_variable that contains the element to be inserted
134   * @param destination ID of the destination node
135   * @param pos position in the array
136   */
137  void bsp_put(const bsp_variable<T>& elem, int destination, int pos) {
138      const auto& t = *(elem.element);
139      bsp_put(t, destination, pos);
140  }
141
142  /**
143   * Replaces the node's private copy of the array with the given vector.
144   * @param array vector that will replace the current private copy of the
145   * array
146   */
147  void bsp_put(const std::vector<T>& array) {
148      comm->array_put(reference, holder, 0, holder, 0, 0, array);
149  }
150
151  /**
152   * Replaces the node's private copy of the array with the vector
153   * contained inside another <tt>bsp_array</tt>.
154   * @param other <tt>bsp_array</tt> containing the vector that will

```

```

155     * replace the current one
156     */
157 void bsp_put(const bsp_array<T>& other) {
158     const auto& t = *(other.arr);
159     bsp_put(t);
160 }
161
162 /**
163  * Replaces another node's private copy of the array with the given vector.
164  * @param array vector that will replace the current private copy of the
165  * array
166  * @param destination ID of the destination node
167  */
168 void bsp_put(const std::vector<T>& array, int destination) {
169     comm->array_put(reference, holder, 0, destination, 0, 0, array);
170 }
171
172 /**
173  * Replaces another node's private copy of the array with the vector
174  * contained inside the given <tt>bsp_array</tt>.
175  * @param other <tt>bsp_array</tt> containing the vector that will
176  * replace the current private copy of the array
177  * @param destination ID of the destination node
178  */
179 void bsp_put(const bsp_array<T>& other, int destination) {
180     const auto& t = *(other.arr);
181     bsp_put(t, destination);
182 }
183
184 /**
185  * Replaces a portion of this node's private copy of the array with a
186  * portion of the given vector.
187  * @param array vector that contains the desired portion
188  * @param src_offset starting position of the replacing portion
189  * @param dest_offset starting position of the portion to be replaced
190  * @param length length of the replacing portion
191  */
192 void bsp_put(const std::vector<T>& array, int src_offset, int dest_offset,
193             int length) {
194     comm->array_put(reference, holder, src_offset, holder, dest_offset,
195                   length, array);
196 }
197
198 /**
199  * Replaces a portion of this node's private copy of the array with a
200  * portion of the vector contained inside the given <tt>bsp_array</tt>.
201  * @param other <tt>bsp_array</tt> containing the vector that contains
202  * the desired portion
203  * @param src_offset starting position of the replacing portion
204  * @param dest_offset starting position of the portion to be replaced
205  * @param length length of the replacing portion
206  */
207 void bsp_put(const bsp_array<T>& other, int src_offset, int dest_offset,
208             int length) {
209     const auto& t = *(other.arr);
210     bsp_put(t, src_offset, dest_offset, length);
211 }
212

```

```

213  /**
214   * Replaces a portion of another node's private copy of the array with a
215   * portion of the given vector.
216   * @param array vector that contains the desired portion
217   * @param destination ID of the destination node
218   * @param src_offset starting position of the replacing portion
219   * @param dest_offset starting position of the portion to be replaced
220   * @param length length of the replacing portion
221   */
222  void bsp_put(const std::vector<T>& array, int destination, int src_offset,
223             int dest_offset, int length) {
224      comm->array_put(reference, holder, src_offset, destination, dest_offset,
225                    length, array);
226  }
227
228  /**
229   * Replaces a portion of another node's private copy of the array with a
230   * portion of the vector contained inside the given <tt>bsp_array</tt>.
231   * @param other <tt>bsp_array</tt> containing the vector that contains
232   * the desired portion
233   * @param destination ID of the destination node
234   * @param src_offset starting position of the replacing portion
235   * @param dest_offset starting position of the portion to be replaced
236   * @param length length of the replacing portion
237   */
238  void bsp_put(const bsp_array<T>& other, int destination, int src_offset,
239             int dest_offset, int length) {
240      const auto& t = *(other.arr);
241      bsp_put(t, destination, src_offset, dest_offset, length);
242  }
243
244  /**
245   * Replaces the node's private copy of the array with another node's
246   * private copy.
247   * @param from ID of the source node
248   */
249  void bsp_get(int from) override {
250      comm->array_get<T>(reference, from, 0, holder, 0, 0);
251  }
252
253  /**
254   * Replaces a portion of the node's private copy of the array with a
255   * portion of another node's private copy.
256   * @param from ID of the source node
257   * @param from_offset starting position of the replacing portion
258   * @param to_offset starting position of the portion to be replaced
259   * @param length length of the replacing portion
260   */
261  void bsp_get(int from, int from_offset, int to_offset, int length) {
262      comm->array_get<T>(reference, from, from_offset, holder, to_offset,
263                    length);
264  }
265
266  /**
267   * Returns a copy of the element at the desired position from another node's
268   * copy of the array. This method will return immediately, without
269   * waiting for a superstep sync.
270   * @param source ID of the source node

```



```

271     * @param pos position of the desired element
272     * @return the element of the source's private copy at the desired position
273     */
274     T bsp_direct_get(int source, int pos) {
275         return comm->array_direct_get<T>(reference, source, pos);
276     }
277
278     /**
279     * Returns a copy of another node's private array. This method will return
280     * immediately, without waiting for a superstep sync.
281     * @param source ID of the source node
282     * @return the source's private copy of the shared array
283     */
284     std::vector<T> bsp_direct_get(int source) {
285         return comm->array_direct_get<T>(reference, source);
286     }
287
288     /**
289     * Returns a handle to this node's private copy of the shared array. The
290     * returned object can be modified at will, without waiting for superstep
291     * syncs.
292     * This function is <b>BSP unsafe</b>.
293     * @return a reference to the node's private copy of the array
294     */
295     std::vector<T>& BSPunsafe_access() {
296         return *arr;
297     }
298
299 };
300
301
302 #endif //FF_BSP_BSP_ARRAY_HPP

```

## bsp\_barrier.hpp

```

1  #ifndef FF_BSP_BSP_BARRIER_HPP
2  #define FF_BSP_BSP_BARRIER_HPP
3
4  #include <mutex>
5
6  /**
7   * A simple reusable barrier.
8   */
9  class bsp_barrier {
10 private:
11
12     /** Mutex for access to the barrier
13     std::mutex mutex;
14     /** Condition variable for the wait/notify mechanism
15     std::condition_variable cond_var;
16     /** The number of workers that need to access the barrier before
17     /** unlocking it.
18     int threshold = 0;
19     /** Current count of workers at the barrier
20     int count = 0;
21     /** Number of times the barrier has been used
22     int generation = 0;

```

```

23
24 public:
25
26     /**
27      * Deleted default constructor.
28      */
29     bsp_barrier() = delete;
30
31     /**
32      * Creates a barrier for a certain number of threads.
33      * @param size the number of threads that will access the barrier
34      */
35     explicit bsp_barrier(int size) :
36         threshold{size},
37         count{size},
38         generation{0} {
39     }
40
41     /**
42      * Copy constructor.
43      * @param other the other <tt>bsp_barrier</tt> object to copy.
44      */
45     bsp_barrier(const bsp_barrier& other) :
46         threshold{other.threshold},
47         count{other.count},
48         generation{other.generation} {
49     }
50
51     /**
52      * Waits on the barrier until all peers do the same.
53      */
54     void wait() {
55         std::unique_lock<std::mutex> lock{mutex};
56         auto lastgen = generation;
57         if (!(--count)) {
58             generation++;
59             count = threshold;
60             cond_var.notify_all();
61         } else {
62             cond_var.wait(lock,
63                 [this, lastgen]() { return lastgen != generation; });
64         }
65     }
66 };
67
68 #endif //FF_BSP_BSP_BARRIER_HPP

```

### bsp\_communicator.hpp

```

1 #ifndef FF_BSP_BSP_COMMUNICATOR_HPP
2 #define FF_BSP_BSP_COMMUNICATOR_HPP
3
4 #include <set>
5 #include <iostream>
6 #include "bsp_array.hpp"
7
8 #ifdef STL_ALLOC

```

```

9  #define m_shared_ptr(U, what) \
10     std::allocate_shared<U, std::allocator<U>> (std::allocator<U>(), (what))
11 #define alloc(U) std::allocator<U>
12 #else
13 #define m_shared_ptr(U, what) \
14     std::allocate_shared<U, ff::FF_Allocator<U>> (ff::FF_Allocator<U>(), (what))
15 #define alloc(U) ff::FF_Allocator<U>
16 #endif
17
18 /**
19  * Implementation of the communicator entity.
20  */
21
22 // Requests management
23
24 /**
25  * Inserts a request for replacing a node's private copy of a shared variable
26  * with the provided value.
27  * @tparam T the type of the variable
28  * @param what ID of the variable to be modified
29  * @param source ID of the node that performs the request
30  * @param destination ID of the node that holds the private copy of the
31  * variable that will be modified
32  * @param element the value to be copied inside the destination
33  */
34 template<typename T>
35 void bsp_communicator::variable_put(int what, int source, int destination,
36                                     const T& element) {
37     if (what == 0) return;
38     mutexes[destination].lock();
39     requests[destination].emplace_back(request_type::var_put, what, source,
40                                       destination, 0, 0, 0,
41                                       m_shared_ptr(T, element));
42     mutexes[destination].unlock();
43 }
44
45 /**
46  * Inserts a request for replacing a node's private copy of a shared variable
47  * with the requestor's private value of the same shared variable.
48  * @tparam T the type of the variable
49  * @param what ID of the variable to be modified
50  * @param source ID of the node that performs the request
51  * @param dest ID of the node that holds the private copy of the
52  * variable that will be modified
53  */
54 template<typename T>
55 void bsp_communicator::variable_put(int what, int source, int dest) {
56     if (what == 0) return;
57     mutexes[dest].lock();
58     auto elem = (static_cast<T*>(variables_storage.at(what).element[source]));
59     requests[dest].emplace_back(request_type::var_put, what, source, dest, 0, 0,
60                                0, m_shared_ptr(T, *elem));
61     mutexes[dest].unlock();
62 }
63
64 /**
65  * Inserts a request for replacing an element of a node's private copy of a
66  * shared array with the provided value.

```

```

67 * @tparam T the type of the elements of the array
68 * @param what ID of the array to be modified
69 * @param source ID of the node that performs the request
70 * @param dest ID of the node that holds the private copy of the
71 * array that will be modified
72 * @param pos position of the element to be modified
73 * @param elem the value to be copied inside the destination
74 */
75 template<typename T>
76 void bsp_communicator::array_put(int what, int source, int dest, int pos,
77                                 const T& elem) {
78     mutexes[dest].lock();
79     requests[dest].emplace_back(request_type::arr_put_el, what, source, dest, 0,
80                                 pos, 0, m_shared_ptr(T, elem));
81     mutexes[dest].unlock();
82 }
83
84 /**
85 * Inserts a request for replacing a portion of a node's private copy of a
86 * shared array with a portion of the provided array.
87 * @tparam T the type of the elements of the arrays
88 * @param what ID of the array to be modified
89 * @param src ID of the node that performs the request
90 * @param src_off starting position of the portion of the replacing array
91 * @param dest ID of the node that holds the private copy of the
92 * array that will be modified
93 * @param dest_off starting position of the portion to be replaced
94 * @param len length of the portion of the replacing array
95 * @param v the replacing array
96 */
97 template<typename T>
98 void bsp_communicator::array_put(int what, int src, int src_off, int dest,
99                                 int dest_off, int len,
100                                const std::vector<T>& v) {
101     mutexes[dest].lock();
102     requests[dest].emplace_back(request_type::arr_put, what, src, dest,
103                                 src_off, dest_off, len,
104                                 m_shared_ptr(std::vector<T>, v));
105     mutexes[dest].unlock();
106 }
107
108 /**
109 * Inserts a request for replacing a portion of a node's private copy of a
110 * shared array with a portion of the private copy of the requesting node
111 * @tparam T the type of the elements of the arrays
112 * @param what ID of the array to be modified
113 * @param src ID of the node that performs the request
114 * @param src_off starting position of the portion of the requesting node's
115 * private copy of the array
116 * @param dest ID of the node that holds the private copy of the
117 * array that will be modified
118 * @param dest_off starting position of the portion to be replaced
119 * @param len length of the portion of the replacing array
120 */
121 template<typename T>
122 void bsp_communicator::array_get(int what, int from, int from_off, int to,
123                                 int to_off, int len) {
124     mutexes[to].lock();

```

```

125     auto arr = static_cast<std::vector<T*>>(arrays_storage.at(
126         what).element[from]);
127     requests[to].emplace_back(request_type::arr_get, what, from, to, from_off,
128         to_off, len, m_shared_ptr(std::vector<T>, *arr));
129     mutexes[to].unlock();
130 }
131
132 /**
133  * Returns the value of a node's private copy of a shared variable,
134  * without waiting for a superstep sync.
135  * @tparam T the type of the variable
136  * @param what ID of the variable to be queried
137  * @param source ID of the node that holds the requested private copy
138  * @return a copy of the desired value
139  */
140 template<typename T>
141 T bsp_communicator::variable_direct_get(int what, int source) {
142     if (what == 0) return fastflow_input;
143     return *(static_cast<T*>(variables_storage.at(what).element.at(source)));
144 }
145
146 /**
147  * Returns the value of an element of a node's private copy of a shared array,
148  * without waiting for a superstep sync.
149  * @tparam T the type of the elements of the array
150  * @param what ID of the variable to be queried
151  * @param src ID of the node that holds the requested private copy
152  * @param pos position of the desired element
153  * @return a copy of the desired value
154  */
155 template<typename T>
156 T bsp_communicator::array_direct_get(int what, int src, int pos) {
157     auto arptr = static_cast<std::vector<T*>>(arrays_storage.at(
158         what).element.at(src));
159     return arptr->at(pos);
160 }
161
162 /**
163  * Returns a node's private copy of a shared array, without waiting for a
164  * superstep sync.
165  * @tparam T the type of the elements of the array
166  * @param what ID of the variable to be queried
167  * @param src ID of the node that holds the requested private copy
168  * @return a copy of the desired array
169  */
170 template<typename T>
171 std::vector<T> bsp_communicator::array_direct_get(int what, int src) {
172     return *(static_cast<std::vector<T*>>(arrays_storage.at(what).element.at(
173         src)));
174 }
175
176 /**
177  * Requests the handle to a shared variable of type |c T. Will create a new
178  * shared variable if the requesting node already has the handle for all
179  * already-present shared variables of type T. Sets the private copy of the
180  * returned shared variable to an initial value.
181  * @tparam T the type of the requested shared variable
182  * @param holder ID of the node that requests the handle

```

```

183 * @param initial_val pointer to the value to be copied inside the requesting
184 * node's private copy
185 * @return a handle of the shared variable, wrapped in a \c bsp_variable.
186 */
187 template<typename T>
188 bsp_variable<T> bsp_communicator::get_variable(int holder, T* initial_val) {
189     auto tname = typeid(T).name();
190     // no. of variables of type T requested by the current worker
191     int get_count = 0;
192     try {
193         get_count = var_count[holder].at(std::string(tname));
194     } catch (const std::out_of_range&) {
195         var_count[holder].insert({std::string(tname), 0});
196     }
197     // hash on type, superstep and number of vars
198     int hash = get_hash(typeid(T).name(), (generation * 5000000) + get_count);
199     int ref;
200     // try to find a variable w/ the same hash
201     //(i.e. another worker has already requested the creation of the var)
202     try {
203         // multiple readers-single writer pattern
204         std::shared_lock lock(var_mutex);
205         ref = variable_dict.at(hash);
206     } catch (const std::out_of_range&) {
207         std::unique_lock lock(var_mutex);
208         // try again, in case another thread created the variable
209         // while this one waited
210         auto iter = variable_dict.find(hash);
211         // if the variable is still not present...
212         if (iter == variable_dict.end()) {
213             // create the variable...
214             inner_var var{nprocs, [] (void* el, void* other) {
215                 auto tptr = static_cast<T*>(el);
216                 *tptr = *(static_cast<T*>(other));
217             }, [] (void* el) {
218                 auto tptr = static_cast<T*>(el);
219                 delete tptr;
220             }};
221             ref = variables_storage.size();
222             // ..and store it
223             variables_storage.push_back(var);
224             // together with its hash
225             variable_dict.insert({hash, ref});
226         } else {
227             // the variable is present
228             ref = iter->second;
229         }
230     }
231     // initialize with the value requested by the worker
232     variables_storage[ref].element.at(holder) = initial_val;
233     var_count[holder].at(std::string(tname))++;
234     return bsp_variable<T>{ref, holder, this, initial_val};
235 }
236
237 /**
238 * Requests the handle to a shared array with elements of type \c T. Will
239 * create a new shared array if the requesting node already has the handle for
240 * all already-present shared arrayss of type T. Sets the private copy of the

```

```

241 * returned shared array to an initial value.
242 * @tparam T the type of the requested shared variable
243 * @param holder ID of the node that requests the handle
244 * @param initial_arr pointer to the value to be copied inside the requesting
245 * node's private copy
246 * @param to_delete flag that indicates whether the array can be deleted
247 * safely when performing cleanup at the end of the computation
248 * @return a handle of the shared array, wrapped in a lc_bsp_array.
249 */
250 template<typename T>
251 bsp_array<T>
252 bsp_communicator::get_array(int holder, std::vector<T*> initial_arr,
253                             bool to_delete) {
254     auto tname = typeid(T).name();
255     int get_count = 0;
256     try {
257         get_count = arr_count[holder].at(std::string(tname));
258     } catch (const std::out_of_range&) {
259         arr_count[holder].insert({std::string(tname), 0});
260     }
261     int hash = get_hash(typeid(T).name(), (generation * 5000000) + get_count);
262     int ref;
263     try {
264         std::shared_lock lock(arr_mutex);
265         ref = array_dict.at(hash);
266     } catch (const std::out_of_range&) {
267         std::unique_lock lock(arr_mutex);
268         auto iter = array_dict.find(hash);
269         if (iter == array_dict.end()) {
270             inner_array arr{nprocs, [] (void* el, void* other, int pos) {
271                 auto arrptr = static_cast<std::vector<T*>>(el);
272                 arrptr->at(pos) = *(static_cast<T*>(other));
273             }, [] (void* el, int srcof, int dstof, int size, void* toput) {
274                 auto arrptr = static_cast<std::vector<T*>>(el);
275                 auto otherptr = static_cast<std::vector<T, alloc(T)>>(tput);
276                 if (size == 0) {
277                     arrptr->resize(otherptr->size());
278                     for (size_t idx{0}; idx < otherptr->size(); ++idx) {
279                         arrptr->at(idx) = otherptr->at(idx);
280                     }
281                 } else {
282                     std::copy_n(otherptr->begin() + srcof, size,
283                                 arrptr->begin() + dstof);
284                 }
285             }, to_delete ? [] (void* el) {
286                 auto tptr = static_cast<std::vector<T*>>(el);
287                 delete tptr;
288             } : [] (void*) {});
289             ref = arrays_storage.size();
290             arrays_storage.push_back(arr);
291             array_dict.insert({hash, ref});
292         } else {
293             ref = iter->second;
294         }
295     }
296     arrays_storage[ref].element.at(holder) = initial_arr;
297     arr_count[holder].at(std::string(tname))++;
298     return bsp_array<T>{ref, holder, this, initial_arr};

```

```

299 }
300
301 /**
302  * During the superstep synchronization, processes all requests with a
303  * certain node as destination. The last node to finish this operation will also
304  * advance the superstep count.
305  * @param id ID of the node that will process the requests.
306  */
307 void bsp_communicator::process_requests(int id) {
308     // Request filtering by worker id
309     for (const auto& req: requests[id]) {
310         switch (req.t) {
311             case request_type::var_put: {
312                 auto ptr = variables_storage.at(
313                     req.reference).element[req.destination];
314                 variables_storage.at(req.reference).swap(ptr,
315                                         req.element.get());
316                 break;
317             }
318             case request_type::arr_put_el: {
319                 auto ptr = arrays_storage.at(
320                     req.reference).element[req.destination];
321                 arrays_storage.at(req.reference).put(ptr, req.element.get(),
322                                         req.dest_offset);
323                 break;
324             }
325             case request_type::arr_put: {
326                 auto ptr = arrays_storage.at(
327                     req.reference).element[req.destination];
328                 arrays_storage.at(req.reference).replace(ptr, req.src_offset,
329                                         req.dest_offset,
330                                         req.length,
331                                         req.element.get());
332                 break;
333             }
334             case request_type::arr_get: {
335                 auto ptr = arrays_storage.at(
336                     req.reference).element[req.destination];
337                 arrays_storage.at(req.reference).replace(ptr, req.src_offset,
338                                         req.dest_offset,
339                                         req.length,
340                                         req.element.get());
341                 break;
342             }
343         }
344     }
345     // The last worker to finish managing its requests will advance the
346     // computation by increasing the superstep count and resetting all the
347     // superstep-specific data structures used by the communicator
348     if (++process_count == nprocs) {
349         generation++;
350         delete[] arr_count;
351         arr_count = new std::map<std::string, int>[nprocs]();
352         delete[] var_count;
353         var_count = new std::map<std::string, int>[nprocs]();
354         delete[] requests;
355         requests = new std::vector<request, alloc(request)>[nprocs]();
356         process_count = 0;

```



```

357     }
358 }
359
360 /**
361  * Makes the FastFlow input token available to all BSP nodes.
362  * @param in pointer to the input token
363  */
364 void bsp_communicator::set_fastflow_input(const void* in) {
365     fastflow_input = in;
366 }
367
368 /**
369  * Returns the pointer to the FastFlow input token.
370  * @return a \c const pointer to the FastFlow input token.
371  */
372 const void* bsp_communicator::get_fastflow_input() {
373     return fastflow_input;
374 }
375
376 /**
377  * Deallocates all the data structures used in the communicator and all shared
378  * arrays and variables.
379  */
380 void bsp_communicator::end() {
381
382     for (auto& var: variables_storage) {
383         auto del = var.free_el;
384         for (auto& el: var.element) {
385             del(el);
386         }
387     }
388
389     for (auto& arr: arrays_storage) {
390         auto del = arr.free_el;
391         for (auto& el: arr.element) {
392             del(el);
393         }
394     }
395
396     delete[] arr_count;
397     arr_count = nullptr;
398     delete[] var_count;
399     var_count = nullptr;
400     delete[] requests;
401     requests = nullptr;
402     delete[] mutexes;
403     mutexes = nullptr;
404 }
405
406 #endif //FF_BSP_BSP_COMMUNICATOR_HPP

```

## bsp\_internals.hpp

```

1 #ifndef FF_BSP_INTERNALS_HPP
2 #define FF_BSP_INTERNALS_HPP
3
4 #include <utility>

```

```

5  #include <vector>
6  #include <memory>
7  #include <shared_mutex>
8  #include <map>
9  #include <algorithm>
10 #include <atomic>
11 #include "stl_allocator.hpp"
12
13 #ifndef STL_ALLOC
14     #define alloc(T) std::allocator<T>
15 #else
16     #define alloc(T) ff::FF_Allocator<T>
17 #endif
18
19 /**
20  * Contains definitions and class prototypes used by the communicator and
21  * variable/array classes
22  */
23
24 /**
25  * Represents the type of a shared object (variable or array)
26  */
27 typedef enum {
28     array,
29     variable
30 } vartype;
31
32 // Forward declaration
33 class bsp_communicator;
34
35 /**
36  * Base class for BSP variables and arrays, i.e. private copies of shared
37  * objects.
38  */
39 class bsp_container {
40 protected:
41     ///! ID of the shared object
42     int reference;
43     ///! ID of the node that owns this private copy
44     int holder;
45     ///! Pointer to the communicator object
46     bsp_communicator* comm;
47
48     ///! Method that will return the object's type (variable or array)
49     virtual vartype var_type() = 0;
50
51     ///
52     * Base constructor for the class.
53     * @param _reference ID of the shared object
54     * @param holder_pid ID of the node that owns this private copy
55     * @param _comm Pointer to the communicator object
56     */
57     bsp_container(int _reference, int holder_pid, bsp_communicator* _comm) :
58         reference{_reference}, holder{holder_pid}, comm{_comm} {};
59
60 public:
61
62     virtual void bsp_get(int) = 0;

```

```

63 };
64
65
66 // Forward declaration
67 template<typename T>
68 class bsp_variable;
69 // Forward declaration
70 template<typename T>
71 class bsp_array;
72
73 // Class declaration for the communicator
74 class bsp_communicator {
75 private:
76     /**
77      * Represent the possible types for a request object.
78      */
79     enum request_type {
80         var_put,
81         arr_put_el,
82         arr_put,
83         arr_get
84     };
85
86     /**
87      * Representation of a put/get request in the communicator.
88      */
89     struct request {
90         request_type t;
91         int reference;
92         int source;
93         int destination;
94         int src_offset;
95         int dest_offset;
96         int length;
97         std::shared_ptr<void> element;
98
99         request(request_type _t, int _ref, int _src, int _dest,
100                int _src_of, int _dst_of, int _len, std::shared_ptr<void> _el) :
101             t{_t}, reference{_ref}, source{_src}, destination{_dest},
102             src_offset{_src_of}, dest_offset{_dst_of}, length{_len},
103             element{std::move(_el)} {};
104     };
105
106     /**
107      * Shared variable object in the shared memory.
108      */
109     struct inner_var {
110         ///! Vector of private copies of the variable
111         std::vector<void*, alloc(void*)> element;
112         ///! Bookkeeping function needed to work with <tt>void*</tt>
113         ///! Replaces a variable with another value
114         void (* swap)(void* el,
115                      void* other);
116         ///! Bookkeeping function needed to work with <tt>void*</tt>
117         ///! Safely frees memory occupied by an element
118         void (* free_el)(
119             void* el);
120

```

```

121     /**
122      * Constructor for the shared variable object.
123      * @param nprocs number of BSP nodes in the computation
124      * @param swapfun pointer to function for element swapping
125      * @param free_elfun pointer to function for safe delete
126      */
127     inner_var(int nprocs, void (* swapfun)(void*, void*),
128              void (* free_elfun)(void*)) : swap{swapfun},
129                                         free_el{free_elfun} {
130         element.resize(nprocs);
131     }
132 };
133
134 /**
135  * Shared array object in the shared memory.
136  */
137 struct inner_array {
138     ///! Vector of private copies of the array
139     std::vector<void*, alloc(void*)> element;
140     ///! Bookkeeping function needed to work with <tt>void*</tt>
141     ///! Replaces an element of the a private copy of the array with another
142     ///! value
143     void (* put)(void* el, void* toput,
144                 int pos);
145     ///! Bookkeeping function needed to work with <tt>void*</tt>
146     ///! Replaces a portion of a private copy of the array
147     void (* replace)(void* el, int srcof, int dstof, int len,
148                     void* toput);
149     ///! Bookkeeping function needed to work with <tt>void*</tt>
150     ///! Safely frees memory occupied by an element
151     void (* free_el)(
152                 void* el);
153
154     /**
155      * Constructor for the shared variable object.
156      * @param nprocs number of BSP nodes in the computation
157      * @param putfun pointer to function for element swapping
158      * @param replacefun pointer to function for replacing a portion
159      * @param free_elfun pointer to function for safe delete
160      */
161     inner_array(int nprocs, void (* putfun)(void*, void*, int),
162                void (* replacefun)(void*, int, int, int, void*),
163                void (* free_elfun)(void*)) :
164         put{putfun}, replace{replacefun}, free_el{free_elfun} {
165         element.resize(nprocs);
166     }
167 };
168
169 /**
170  * Hashes a string and a number according to the dbj2 hash.
171  * @param s the string to be hashed
172  * @param seed an initial value for the hashing function
173  * @return a (hopefully) unique number that represents the input pair
174  */
175 static int get_hash(const char* s, int seed) {
176     unsigned int hash = seed + 5381;
177     while (*s) {
178         hash = hash * 33 ^ (*s++);

```

```

179     }
180     return hash;
181 }
182
183 ///! Number of BSP nodes
184 int nprocs;
185 ///! Number of current superstep
186 int generation = 1;
187
188 ///! Counts the variables requested by each node in this superstep
189 std::map<std::string, int>* var_count;
190 ///! Counts the arrays requested by each node in this superstep
191 std::map<std::string, int>* arr_count;
192
193 ///! Mutex for multiple readers-single writer access to the variables
194 mutable std::shared_mutex var_mutex;
195 ///! Mutex for multiple readers-single writer access to the arrays
196 mutable std::shared_mutex arr_mutex;
197
198 ///! Dictionary for quick retrieving of variables based on their hash
199 std::map<int, int> variable_dict;
200 ///! Dictionary for quick retrieving of arrays based on their hash
201 std::map<int, int> array_dict;
202
203 ///! Container that stores variables
204 std::vector<inner_var> variables_storage;
205 ///! Container that stores arrays
206 std::vector<inner_array> arrays_storage;
207
208 ///! Array of mutexes to access request queues
209 std::mutex* mutexes;
210 ///! Requests of variable/array modifications submitted during this superstep
211 std::vector<request, alloc(request)>* requests;
212
213 ///! Counts the nodes that requested a sync during this superstep
214 std::atomic_int process_count{0};
215
216 ///! Pointer to the input token received by the FastFlow node
217 const void* fastflow_input = nullptr;
218
219 public:
220
221 /**
222  * Constructor for the communicator.
223  * @param _nprocs number of BSP nodes in this computation.
224  */
225 explicit bsp_communicator(int _nprocs) : nprocs[_nprocs] {
226     var_count = new std::map<std::string, int>[_nprocs]();
227     arr_count = new std::map<std::string, int>[_nprocs]();
228     mutexes = new std::mutex[_nprocs]();
229     requests = new std::vector<request, alloc(request)>[_nprocs]();
230     inner_var invar{nprocs, [](void*, void*) {}, [](void*) {}};
231     variables_storage.push_back(invar);
232 };
233
234 // Refer to bsp_communicator.hpp
235
236 void set_fastflow_input(const void*);

```

```

237
238     const void* get_fastflow_input();
239
240     template<typename T>
241     void variable_put(int, int, int, const T&);
242
243     template<typename T>
244     void variable_put(int, int, int);
245
246     template<typename T>
247     T variable_direct_get(int, int);
248
249     template<typename T>
250     void array_put(int, int, int, int, const T&);
251
252     template<typename T>
253     void array_put(int, int, int, int, int, int, const std::vector<T>&);
254
255     template<typename T>
256     void array_get(int, int, int, int, int, int);
257
258     template<typename T>
259     T array_direct_get(int, int, int);
260
261     template<typename T>
262     std::vector<T> array_direct_get(int, int);
263
264     template<typename T>
265     bsp_variable<T> get_variable(int holder, T* initial_val);
266
267     template<typename T>
268     bsp_array<T>
269     get_array(int holder, std::vector<T>* initial_arr, bool to_delete = true);
270
271     void process_requests(int id);
272
273     void end();
274
275 };
276
277 #endif //FF_BSP_BSP_VARIABLE_INTERNAL_HPP

```

## bsp\_node.hpp

```

1  #ifndef FF_BSP_BSP_NODE_HPP
2  #define FF_BSP_BSP_NODE_HPP
3
4  #include <ff/ff.hpp>
5  #include "bsp_communicator.hpp"
6  #include "bsp_barrier.hpp"
7
8
9  /**
10 * Specialization of a ff_node to work as the unit of computation in the BSP
11 * model.
12 */
13 class bsp_node: public ff::ff_node {

```

```

14 private:
15
16     ///! Number of BSP nodes in this computation
17     int nprocs;
18     ///! ID for this node
19     int id = -1;
20     ///! Pointer to the communicator; implementation of the Communication Channel
21     bsp_communicator* comm;
22     ///! Pointer to the barrier; implementation of the Synchronization Channel
23     bsp_barrier* barrier;
24
25     /// bsp_program can access fields of this class
26     friend class bsp_program;
27
28 protected:
29
30     ///! Pointer to the FastFlow input token
31     const void* fastflow_input;
32
33     /**
34      * Forwards an output token to the next stage in the FastFlow graph.
35      * @param payload the token to forward
36      */
37     void emit_output(void* payload) {
38         ff_send_out(payload);
39     }
40
41     /**
42      * Requests a new variable of type T from the communicator.
43      * @tparam T the type of the requested variable
44      * @param initial_value value to copy inside this node's private copy of
45      * the variable
46      * @return a handle to this node's private copy of the shared variable
47      */
48     template <typename T>
49     bsp_variable<T> get_variable(const T& initial_value) {
50         T* val = new T(initial_value);
51         return comm->get_variable<T>(id, val);
52     }
53
54     /**
55      * Requests a new array with elements of type T from the communicator,
56      * initializing it with the copy a given vector.
57      * @tparam T the type of elements of the requested array
58      * @param initial_value value to copy inside this node's private copy of
59      * the array
60      * @return a handle to this node's private copy of the shared variable
61      */
62     template <typename T>
63     bsp_array<T> get_array(const std::vector<T>& initial_value) {
64         auto val = new std::vector<T>(initial_value);
65         return comm->get_array(id, val);
66     }
67
68     /**
69      * Requests a new array with elements of type T from the communicator,
70      * initializing it with the pointer of a vector. Any modifications done
71      * to the initializing vector after the call to this method is inherently

```

```

72     * BSP unsafe.
73     * @tparam T the type of elements of the requested array
74     * @param handle a pointer to the vector that will become this node's
75     * private copy of the shared array
76     * @return a handle to this node's private copy of the shared variable
77     */
78     template <typename T>
79     bsp_array<T> get_array(std::vector<T>* handle) {
80         return comm->get_array(id, handle, false);
81     }
82
83     /**
84     * Requests a new array with elements of type T from the communicator,
85     * initializing it with an empty vector of given size.
86     * @tparam T the type of elements of the requested array
87     * @param size the size of the empty vector that will become this node's
88     * private copy of the shared array
89     * @return a handle to this node's private copy of the shared variable
90     */
91     template <typename T>
92     bsp_array<T> get_empty_array(int size) {
93         auto val = new std::vector<T>(size);
94         return comm->get_array(id, val);
95     }
96
97     /**
98     * Returns the ID for this node.
99     * @return this node's ID.
100    */
101    int bsp_pid() {
102        return id;
103    }
104
105    /**
106    * Returns the number of nodes in the current BSP computation.
107    * @return the number of nodes in the current BSP computation.
108    */
109    int bsp_nprocs() {
110        return nprocs;
111    }
112
113    /**
114    * Terminates the current superstep and waits for the other nodes to sync.
115    */
116    void bsp_sync() {
117        barrier->wait();
118        comm->process_requests(id);
119        barrier->wait();
120    }
121
122    /**
123    * Function to be overwritten as the main parallel execution.
124    */
125    virtual void parallel_function() = 0;
126
127    public:
128
129    /**

```



```

130     * The FastFlow node service method.
131     *
132     * Implementations of this class cannot redefine it.
133     * @param in the input token (in this case, a special value \c ENDCOMP)
134     * @return the same token as the input
135     */
136     void* svc (void* in) final {
137         fastflow_input = comm->get_fastflow_input();
138         parallel_function();
139         return in;
140     }
141 };
142
143 #endif //FF_BSP_BSP_NODE_HPP

```

### bsp\_program.hpp

```

1  #ifndef FF_BSP_BSP_PROGRAM_HPP
2  #define FF_BSP_BSP_PROGRAM_HPP
3
4  #include <ff/ff.hpp>
5  #include <memory>
6  #include <iostream>
7  #include "bsp_node.hpp"
8
9  /**
10     * Implements the Bulk Synchronous Parallel pattern as a FastFlow node.
11     */
12     class bsp_program : public ff::ff_node {
13
14     private:
15
16         /**
17          * The emitter of the master-worker scheme; it allows the user to perform a
18          * (sequential) function before the execution of the main parallel part.
19          */
20         struct emitter : ff::ff_node {
21             /// Value to end the computation after every node finishes its job
22             void* ENDCOMP = (void*) ((unsigned long long) ff::FF_TAG_MIN - 1);
23
24             /// Optional function to be executed before the BSP computation
25             std::function<void(void)> preprocessing;
26             /// Number of BSP nodes
27             int emitter_nprocs;
28
29             /**
30              * Constructor for the farm emitter.
31              * @param pre optional function to be executed before the BSP
32              * computation
33              * @param size number of BSP nodes
34              */
35             emitter(std::function<void(void)> pre, int size) :
36                 preprocessing{std::move(pre)},
37                 emitter_nprocs{size} {
38             }
39
40             /**

```

```

41     * Service function of the farm emitter.
42     * @param in input FastFlow token
43     * @return the End-of-Stream token
44     */
45     void* svc(void* in) override {
46         if (preprocessing != nullptr) preprocessing();
47         for (int i = 0; i < emitter_nprocs; i++) {
48             ff_send_out(ENDCOMP);
49         }
50         return EOS;
51     }
52 };
53
54 /**
55  * The collector of the master-worker scheme; it allows the user to perform
56  * a (sequential) function after the execution of the main parallel part.
57  */
58 struct collector : ff::ff_node {
59     /** Value to end the computation after every node finishes its job
60     void* ENDCOMP = (void*) ((unsigned long long) ff::FF_TAG_MIN - 1);
61     /** Optional function to be executed after the BSP computation
62     std::function<void(void)> postprocessing;
63     /** Number of nodes in the BSP computation
64     int threshold;
65     /** Nodes that have finished their local computation
66     int count;
67     /** Pointer to the the outer class
68     bsp_program* master;
69
70     /**
71     * Constructor for the farm collector.
72     * @param post optional function to be executed after the BSP
73     * computation
74     * @param size number of BSP nodes
75     */
76     collector(std::function<void(void)> post, int size) :
77         postprocessing{std::move(post)},
78         threshold{size},
79         count{0} {
80     };
81
82     /**
83     * Service function of the farm emitter.
84     * @param in input FastFlow token
85     * @return the input token, or GO_ON
86     */
87     void* svc(void* in) override {
88         if (in == ENDCOMP) {
89             if (++count == threshold) {
90                 if (postprocessing != nullptr) postprocessing();
91             }
92             return GO_ON;
93         } else {
94             master->forward(in);
95         }
96         return in;
97     }
98 };

```

```

99
100  /**
101   * Forwards the received token to the succeeding FastFlow node.
102   * @param token the token to be forwarded.
103   */
104  void forward(void* token) {
105      ff_send_out(token);
106  }
107
108  ///! Number of BSP nodes (workers of the FastFlow farm pattern)
109  int nprocs;
110  ///! Vector of pointers to BSP nodes
111  std::vector<std::unique_ptr<bsp_node>> processors;
112  ///! The entity responsible for communication between nodes
113  bsp_communicator comm;
114  ///! A barrier for synchronization between supersteps
115  bsp_barrier barr;
116  ///! The farm emitter
117  emitter E;
118  ///! The farm collector
119  collector C;
120
121  public:
122
123  /**
124   * Constructor for the bsp_program object. Pushes necessary information
125   * into the relevant entities (emitter, workers, collector)
126   * @param _processors vector of BSP nodes for the computation
127   * @param _pre optional function to be executed before the BSP computation
128   * @param _post optional function to be executed after the BSP computation
129   */
130  explicit bsp_program(std::vector<std::unique_ptr<bsp_node>>&& _processors,
131                      std::function<void(void)> _pre = nullptr,
132                      std::function<void(void)> _post = nullptr) :
133      nprocs{static_cast<int>(_processors.size())},
134      comm{nprocs},
135      barr{nprocs},
136      E{std::move(_pre), nprocs},
137      processors{std::move(_processors)},
138      C{std::move(_post), nprocs} {
139      for (size_t i{0}; i < nprocs; ++i) {
140          processors[i]->nprocs = nprocs;
141          processors[i]->id = i;
142          processors[i]->barrier = &barr;
143          processors[i]->comm = &comm;
144          C.master = this;
145      }
146  };
147
148  /**
149   * Creates the FastFlow inner graph and executes the BSP computation.
150   * @param in optional, the input token received from the preceding FastFlow
151   * node
152   */
153  void start(void* in = nullptr) {
154      std::vector<std::unique_ptr<ff::ff_node>> workers;
155      for (size_t i{0}; i < nprocs; ++i) {
156          auto d = static_cast<ff_node*>(processors[i].release());

```

```

157         workers.emplace_back(std::unique_ptr<ff_node>(d));
158     }
159
160     comm.set_fastflow_input(in);
161
162     ff::ff_Farm<> farm(std::move(workers), E, C);
163
164     if (farm.run_and_wait_end() < 0)
165         std::cout << "error in running farm" << std::endl;
166
167     comm.end();
168 }
169
170 /**
171  * Service function for the BSP program node. Starts the BSP computation.
172  * @param in input token received from the preceding FastFlow node
173  * @return GO_ON
174  */
175 void* svc(void* in) override {
176     start(in);
177     return GO_ON;
178 }
179 };
180
181 #endif //FF_BSP_BSP_PROGRAM_HPP

```

## bsp\_variable.hpp

```

1  #ifndef FF_BSP_BSP_VARIABLE_HPP
2  #define FF_BSP_BSP_VARIABLE_HPP
3
4  #include <type_traits>
5  #include "bsp_internals.hpp"
6
7  /**
8   * A variable data structure, private to a single worker but with support for
9   * communication with other similar entities.
10   *
11   * @tparam T type of the variable
12   */
13  template<typename T>
14  class bsp_variable : public bsp_container {
15
16     // The template type must be copy-constructible and copy-assignable
17     static_assert(std::is_copy_constructible<T>::value &&
18                 std::is_copy_assignable<T>::value,
19                 "Type of bsp_variable must be copy-constructible "
20                 "and copy-assignable");
21
22 private:
23
24     // bsp_array can access private fields of this class
25     template<typename E>
26     friend class bsp_array;
27
28     // bsp_communicator can access private fields of this class
29     friend class bsp_communicator;

```

```

30
31  /**
32   * Pointer to the actual data element.
33   */
34  T* element;
35
36  /**
37   * Default constructor.
38   */
39  bsp_variable() = default;
40
41  /**
42   * Builds a bsp_variable object that holds a node's private copy of a shared
43   * variable.
44   * @param _ref ID of the shared variable
45   * @param _hold ID of the requesting node
46   * @param comm pointer to the communicator component
47   * @param ptr pointer to the node's private copy
48   */
49  explicit bsp_variable(int _ref, int _hold, bsp_communicator* comm, T* ptr) :
50      bsp_container(_ref, _hold, comm), element{ptr} {
51  }
52
53  /**
54   * Returns the shared object type (in this case, a variable).
55   * @return the |c| vartype value for variables
56   */
57  vartype var_type() final {
58      return vartype::variable;
59  }
60
61  public:
62
63  /**
64   * Returns a copy of the node's private element of the shared variable.
65   * @return a copy of the desired element
66   */
67  T get() {
68      return *element;
69  }
70
71  /**
72   * Replaces another node's private element of the shared variable.
73   * @param elem element to be copied
74   * @param destination ID of the destination node
75   */
76  void bsp_put(const T& elem, int destination) {
77      comm->variable_put<T>(reference, holder, destination, elem);
78  }
79
80  /**
81   * Replaces this node's private element with this node's private element of
82   * another shared variable.
83   * @param other the handle to the other shared variable
84   */
85  void bsp_put(const bsp_variable<T>& other) {
86      const auto& t = *(other.element);
87      bsp_put(t, holder);

```

```

88     }
89
90     /**
91     * Replaces another node's private element with this node's private
92     * element of another shared variable.
93     * @param other the handle to the other shared variable
94     * @param destination ID of the destination node
95     */
96     void bsp_put(const bsp_variable<T>& other, int destination) {
97         const auto& t = *(other.element);
98         bsp_put(t, destination);
99     }
100
101     /**
102     * Replaces another node's private element with this node's private
103     * element of this shared variable.
104     * @param destination ID of the destination node
105     */
106     void bsp_put(int destination) {
107         comm->variable_put<T>(reference, holder, destination);
108     }
109
110     /**
111     * Replaces this node's private element with another node's private
112     * element of this shared variable.
113     * @param source ID of the source node
114     */
115     void bsp_get(int source) override {
116         comm->variable_put<T>(reference, source, holder);
117     }
118
119     /**
120     * Returns a copy of a node's private element of this shared variable.
121     * This method will return immediately, without waiting for a superstep
122     * sync.
123     * @param source ID of the source node
124     * @return a copy of the desired element
125     */
126     T bsp_direct_get(int source) {
127         return comm->variable_direct_get<T>(reference, source);
128     }
129
130     /**
131     * Returns a handle to this node's private element of the shared variable.
132     * The returned object can be modified at will, without waiting for
133     * superstep syncs.
134     * This function is <b>BSP unsafe</b>.
135     * @return a reference to the node's private element of the shared variable
136     */
137     T& BSPunsafe_access() {
138         return *element;
139     }
140 };
141
142 #endif //FF_BSP_BSP_VARIABLE_HPP

```

## stl\_allocator.hpp

```
1  #ifndef FF_BSP_STL_ALLOCATOR_HPP
2  #define FF_BSP_STL_ALLOCATOR_HPP
3
4  #include <ff/allocator.hpp>
5
6  namespace ff {
7
8      /**
9       * STL-compliant wrapper for FastFlow's custom allocator.
10     * @tparam T the type of objects to be allocated or deallocated.
11     */
12     template<typename T>
13     class FF_Allocator {
14     public:
15         using value_type = T;
16         using propagate_on_container_move_assignment = std::true_type;
17         using is_always_equal = std::true_type;
18
19         /**
20          * Default constructor.
21          */
22         FF_Allocator() noexcept = default;
23
24         /**
25          * Empty copy constructor.
26          * @tparam U the type of objects of the other allocator object
27          * @param other the other allocator object
28          */
29         template<class U>
30         explicit FF_Allocator(const FF_Allocator<U>& other) noexcept {};
31
32         /**
33          * Allocates n * sizeof(T) bytes of uninitialized storage by calling
34          * the FastFlow allocator instance's |c malloc function.
35          * @param n the number of objects to allocate storage for
36          * @return the pointer to the first newly-allocated object
37          */
38         value_type* allocate(std::size_t n) {
39             return static_cast<value_type*>(FF_Allocator::instance()
40                 ->malloc(n * sizeof(value_type)));
41         }
42
43         /**
44          * Deallocates the storage referenced by the pointer p, which must be
45          * a pointer obtained by an earlier call to allocate().
46          * @param ptr pointer to the object to deallocate
47          */
48         void deallocate(value_type* ptr, std::size_t) noexcept {
49             FF_Allocator::instance()->free(ptr);
50         }
51     };
52
53     /**
54     * Checks that two allocator objects can be considered the same.
55     * @tparam T type of the first allocator
56     * @tparam U type of the second allocator
```

```

57     * @return true
58     */
59     template<class T, class U>
60     bool
61     operator==(const FF_Allocator<T>&, const FF_Allocator<U>&) { return true; }
62
63     /**
64     * Checks that two allocator objects can be considered different.
65     * @tparam T type of the first allocator
66     * @tparam U type of the second allocator
67     * @return false
68     */
69     template<class T, class U>
70     bool
71     operator!=(const FF_Allocator<T>&, const FF_Allocator<U>&) { return false; }
72 }
73
74 #endif //FF_BSP_STL_ALLOCATOR_HPP

```

## B.2 Parallel test programs

### BSPinprod.cpp

```

1  #include <bsp_program.hpp>
2
3  struct BSPinprod : public bsp_node {
4
5      int problem_size;
6
7      explicit BSPinprod(int n) : problem_size{n} {};
8
9      inline int nloc(int p, int s, int n) const {
10         return (n + p - s - 1) / p;
11     }
12
13     double bspip(int p, int s, int n,
14                 const std::vector<double>& x,
15                 const std::vector<double>& y) {
16
17         auto Inprod = get_empty_array<double>(p);
18
19         double inprod = 0.0;
20
21         for (int i{0}; i < nloc(p, s, n); ++i) {
22             inprod += x.at(i) * y.at(i);
23         }
24         for (int t{0}; t < p; ++t) {
25             Inprod.bsp_put(inprod, t, s);
26         }
27         bsp_sync();
28
29         const auto& Inprod_arr = Inprod.BSPunsafe_access();
30         double alpha = 0.0;
31         for (int t{0}; t < p; ++t) {
32             alpha += Inprod_arr.at(t);
33         }
34         return alpha;

```



```

35     }
36
37     void parallel_function() override {
38         int n = problem_size;
39         int p = bsp_nprocs();
40         int s = bsp_pid();
41         int nl = nloc(p, s, n);
42
43         std::vector<double> x(nl);
44
45         for (int i{0}; i < nl; ++i) {
46             x[i] = i * p + s + 1;
47         }
48
49         bsp_sync();
50         std::cout << "taking time 1" << std::endl;
51         auto t1 = std::chrono::high_resolution_clock::now();
52         double alpha = bspip(p, s, n, x, x);
53         auto t2 = std::chrono::high_resolution_clock::now();
54         auto time = std::chrono::duration_cast<std::chrono::milliseconds>(
55             t2 - t1);
56         if (s == 0) {
57             std::cout << "Processor " << s << ": sum of squares up to "
58                 << n << "*" << n << " is " << alpha << std::endl;
59             std::cout << "Processor " << s << ": local time taken is "
60                 << time.count()
61                 << std::endl;
62         }
63     }
64 };
65
66 int main(int argc, char* argv[]) {
67     if (argc == 1) {
68         std::cerr << "Usage: " << argv[0] << " <P> (n)" << std::endl;
69         std::cerr << "         where <P> is the number of processors used"
70             << std::endl;
71         std::cerr << "         and (n) is the size of input vector, "
72             << "optional (default is 500000000)" << std::endl;
73         return 1;
74     }
75     int p = std::stoi(argv[1]);
76     int n = (argc >= 3) ? std::stoi(argv[2]) : 500000000;
77     std::vector<std::unique_ptr<bsp_node>> nodes;
78     for (int i{0}; i < p; ++i) {
79         nodes.emplace_back(std::make_unique<BSPinprod>(n));
80     }
81     auto checksum = [n]() {
82         auto dn = static_cast<double>(n);
83         dn *= n + 1.0;
84         dn *= 2.0 * n + 1.0;
85         dn /= 6.0;
86         std::cout << "Checksum: " << dn << std::endl;
87     };
88 };
89 bsp_program computation(std::move(nodes), nullptr, checksum);
90 computation.start();
91 return 0;
92 }

```

## BSPpsrs.cpp

```
1  #include <random>
2  #include <chrono>
3  #include <bsp_program.hpp>
4
5  struct BSPpsrs : public bsp_node {
6
7      int n;
8      int seed;
9
10     explicit BSPpsrs(int problem_size, int random_seed = -1) :
11         n{problem_size},
12         seed{random_seed} {
13     };
14
15     void parallel_function() override {
16
17         bsp_array<int> to_sort = get_empty_array<int>(1);
18         int p = bsp_nprocs();
19         int s = bsp_pid();
20         auto t = std::chrono::high_resolution_clock::now();
21
22         if (s == 0) {
23             if (seed == -1) {
24                 std::random_device x;
25                 seed = x();
26             }
27             std::mt19937 mtw(seed);
28             std::vector<int> data(n);
29             std::iota(data.begin(), data.end(), 0);
30             std::shuffle(data.begin(), data.end(), mtw);
31
32             #ifdef DEBUG
33             std::ofstream printl{"generated_array.log"};
34             for (const auto& el: data) printl << el << std::endl;
35             printl.close();
36             #endif
37
38             auto t2 = std::chrono::duration_cast<std::chrono::milliseconds>(
39                 std::chrono::high_resolution_clock::now() - t).count();
40             std::cout << "Ended vector generation and shuffling (spent " << t2
41                 << " ms)" <<
42                 "\nstarting parallel part" << std::endl;
43
44             t = std::chrono::high_resolution_clock::now();
45             int numels = n / p;
46             int count = 0;
47             for (int i{0}; i < n; i += numels) {
48                 auto last = std::min(n, i + numels);
49                 to_sort.bsp_put(
50                     std::vector<int>(data.begin() + i, data.begin() + last),
51                     count++);
52             }
53
54             bsp_sync();
55
56             bsp_array<std::vector<int>> ps_array = get_empty_array<std::vector<int>>(
```

```

57         p); // array of vectors!
58
59     auto& vec = to_sort.BSPunsafe_access();
60
61     std::sort(vec.begin(), vec.end());
62
63     #ifdef DEBUG
64         std::ofstream print1a{"distributed_array-" + std::to_string(s) + ".log"};
65         for (const auto& el: vec) print1a << el << std::endl;
66         print1a.close();
67     #endif
68
69     std::vector<int> primary_samples;
70
71     size_t samplesize = vec.size() / p;
72     size_t i;
73
74     for (i = 0; i < vec.size(); i += samplesize) {
75         primary_samples.emplace_back(vec.at(i));
76     }
77
78     if (i != vec.size() - 1) primary_samples.emplace_back(vec.at(i - 1));
79
80     #ifdef DEBUG
81         std::ofstream print2{"primary_samples-" + std::to_string(s) + ".log"};
82         for (const auto& el: primary_samples) print2 << el << std::endl;
83         print2.close();
84     #endif
85
86     for (i = 0; i < p; ++i) {
87         ps_array.bsp_put(primary_samples, i, s);
88     }
89
90     bsp_sync();
91
92     bsp_array<std::vector<int>> portion = get_empty_array<std::vector<int>>(
93         p);
94
95     std::vector<int> ps_all;
96     std::vector<int> secondary_samples;
97     const auto& psref = ps_array.BSPunsafe_access();
98
99     for (const auto& vecs: psref) {
100         ps_all.insert(ps_all.end(), vecs.begin(), vecs.end());
101     }
102
103     std::sort(ps_all.begin(), ps_all.end());
104
105     samplesize = ps_all.size() / p;
106
107     for (i = 0; i < ps_all.size(); i += samplesize) {
108         secondary_samples.emplace_back(ps_all.at(i));
109     }
110
111     if (i == ps_all.size())
112         secondary_samples.emplace_back(ps_all.at(i - 1));
113
114     #ifdef DEBUG

```

```

115     std::ofstream print3{"secondary_samples-" + std::to_string(s) + ".log"};
116     for (const auto& el: secondary_samples) print3 << el << std::endl;
117     print3.close();
118 #endif
119
120     int upperbound;
121     int count = 1;
122
123     do {
124         upperbound = secondary_samples.at(count++);
125     } while (vec.at(0) > upperbound);
126
127     count--;
128     std::vector<int> temp;
129     for (i = 0; i < vec.size(); ++i) {
130         if (vec.at(i) > upperbound) {
131             portion.bsp_put(temp, count - 1, s);
132             temp.clear();
133             upperbound = secondary_samples.at(++count);
134         }
135         temp.emplace_back(vec.at(i));
136     }
137
138     portion.bsp_put(temp, count - 1, s);
139     bsp_sync();
140
141     bsp_array<std::vector<int>> final_arr =
142         get_empty_array<std::vector<int>>(p);
143
144     std::vector<int> secondary_block;
145
146     auto& pbls_ref = portion.BSPunsafe_access();
147
148     size_t totalsz = 0;
149     for (const auto& pbl: pbls_ref) totalsz += pbl.size();
150
151     secondary_block.reserve(totalsz);
152     for (auto& pbl: pbls_ref) {
153         secondary_block.insert(
154             secondary_block.end(),
155             std::make_move_iterator(pbl.begin()),
156             std::make_move_iterator(pbl.end()));
157     }
158
159     std::sort(secondary_block.begin(), secondary_block.end());
160
161 #ifdef DEBUG
162     std::ofstream print4{"secondary_block-" + std::to_string(s) + ".log"};
163     for (const auto& el: secondary_block) print4 << el << std::endl;
164     print4.close();
165 #endif
166
167     final_arr.bsp_put(secondary_block, 0, s);
168     bsp_sync();
169
170     if (s == 0) {
171         auto& sbls_ref = final_arr.BSPunsafe_access();
172

```

```

173     std::vector<int> final;
174     for (auto& sbl: sbls_ref) {
175         final.insert(final.end(), std::make_move_iterator(sbl.begin()),
176                     std::make_move_iterator(sbl.end()));
177     }
178
179     #ifdef DEBUG
180     std::ofstream print5{"final.log"};
181     for (const auto& el: final) print5 << el << std::endl;
182     print5.close();
183     #endif
184
185     bool passed = true;
186     for (int j{0}; j < n; j++) {
187         if (final[j] != j) {
188             passed = false;
189             break;
190         }
191     }
192
193     std::cout << "Check " << (passed ? "passed" : "failed")
194               << std::endl;
195     auto t1 = std::chrono::high_resolution_clock::now();
196     std::cout << "Parallel part took "
197               << std::chrono::duration_cast<std::chrono::milliseconds>(
198                 t1 - t).count() << " ms." << std::endl;
199     }
200
201     bsp_sync(); // Needed to keep all threads alive!
202 }
203
204 };
205
206 int main(int argc, char* argv[]) {
207     if (argc < 3) {
208         std::cerr << "Usage: " << argv[0] << " <N> <P> (seed)" << std::endl;
209         std::cerr << " where N is the problem size (power of 2)" << std::endl;
210         std::cerr << "       P is the number of threads used (power of 2)"
211                   << std::endl;
212         std::cerr << "       N must be >= P^3" << std::endl;
213         std::cerr
214             << "       seed is an optional seed for permutations "
215             << "(leave blank to randomize it)"
216             << std::endl;
217         return -1;
218     }
219     int n, p, s;
220     n = std::stoi(argv[1]);
221     p = std::stoi(argv[2]);
222     if (n < p * p * p) {
223         std::cerr << "N must be >= P^3" << std::endl;
224         return -1;
225     }
226     s = (argc >= 4 ? std::stoi(argv[3]) : -1);
227     std::vector<std::unique_ptr<bsp_node>> nodes;
228     for (size_t i{0}; i < p; ++i) {
229         nodes.push_back(std::make_unique<BSPpsrs>(n, s));
230     }

```

```

231
232     bsp_program mbsp(std::move(nodes));
233     mbsp.start();
234     return 0;
235 }

```

## BSPfft.cpp

```

1  #include <cmath>
2  #include <fstream>
3  #include <bsp_program.hpp>
4
5  struct BSPfft : public bsp_node {
6
7      std::vector<double> global_x;
8      int global_n;
9
10     explicit BSPfft(int _n, std::vector<double> glob) : global_n{_n}, global_x{
11         std::move(glob)} {
12     }
13
14     constexpr static double PI = 3.141592653589793;
15
16     static void ufft(std::vector<double>& x, int offset, int n, bool sign,
17         const std::vector<double>& w) {
18
19         for (int k = 2; k <= n; k *= 2) {
20             int nk = n / k;
21             for (int r = 0; r < nk; ++r) {
22                 int rk = 2 * r * k;
23                 for (int j = 0; j < k; j += 2) {
24                     double wr = w[j * nk];
25                     double wi;
26                     if (sign) {
27                         wi = w[j * nk + 1];
28                     } else {
29                         wi = -w[j * nk + 1];
30                     }
31
32                     int j0 = rk + j + offset;
33                     int j1 = j0 + 1;
34                     int j2 = j0 + k;
35                     int j3 = j2 + 1;
36
37                     double taur = wr * x[j2] - wi * x[j3];
38                     double tau1 = wi * x[j2] + wr * x[j3];
39
40                     x[j2] = x[j0] - taur;
41                     x[j3] = x[j1] - tau1;
42                     x[j0] += taur;
43                     x[j1] += tau1;
44                 }
45             }
46         }
47     }
48
49     static void ufft_init(int n, std::vector<double>& w) {

```

```

50     assert(w.size() == n);
51
52     if (n == 1) return;
53
54     w[0] = 1.0;
55     w[1] = 0.0;
56
57     if (n == 4) {
58         w[2] = 0.0;
59         w[3] = -1.0;
60     } else if (n >= 8) {
61         double theta = -2.0 * PI / static_cast<double>(n);
62         for (int j = 1; j <= n / 8; j++) {
63             w[2 * j] = std::cos(j * theta);
64             w[2 * j + 1] = std::sin(j * theta);
65         }
66         for (int j = 0; j < n / 8; j++) {
67             int n4j = n / 4 - j;
68             w[2 * n4j] = -w[2 * j + 1];
69             w[2 * n4j + 1] = -w[2 * j];
70         }
71         for (int j = 1; j < n / 4; j++) {
72             int n2j = n / 2 - j;
73             w[2 * n2j] = -w[2 * j];
74             w[2 * n2j + 1] = w[2 * j + 1];
75         }
76     }
77 }
78
79 static void twiddle(std::vector<double>& x, int length, bool sign,
80                   const std::vector<double>& w, int offset) {
81     for (int jo = 0; jo < 2 * length; jo += 2) {
82         int j = jo;
83         int j1 = j + 1;
84         double wr = w[offset + j];
85         double wi;
86
87         if (sign) {
88             wi = w[offset + j1];
89         } else {
90             wi = -w[offset + j1];
91         }
92
93         double xr = x[j];
94         double xi = x[j1];
95         x[j] = wr * xr - wi * xi;
96         x[j1] = wi * xr + wr * xi;
97     }
98 }
99
100 static void twiddle_init(int n, double alpha, const std::vector<int>& rho,
101                          std::vector<double>& w, int offset) {
102     double theta = -2.0 * PI * alpha / static_cast<double>(n);
103     for (int j = 0; j < n; ++j) {
104         double rt = static_cast<double>(rho[j]) * theta;
105         w[offset + 2 * j] = std::cos(rt);
106         w[offset + 2 * j + 1] = std::sin(rt);
107     }

```

```

108     }
109
110     static void
111     permute(std::vector<double>& x, int n, const std::vector<int>& sigma,
112            int line) {
113         assert(x.size() / 2 == sigma.size());
114
115         for (int j = 0; j < n; ++j) {
116             if (j < sigma[j]) {
117                 int j0 = 2 * j;
118                 int j1 = j0 + 1;
119                 int j2 = 2 * sigma[j];
120                 int j3 = j2 + 1;
121                 double tmpr = x[j0];
122                 double tmpi = x[j1];
123                 x[j0] = x[j2];
124                 x[j1] = x[j3];
125                 x[j2] = tmpr;
126                 x[j3] = tmpi;
127             }
128         }
129     }
130
131     static void bitrev_init(std::vector<int>& rho) {
132         int n = rho.size();
133
134         auto binary_len = static_cast<int>(std::ceil(
135             std::log(static_cast<double>(n)) / std::log(2.0)));
136         std::vector<bool> bits(binary_len);
137         std::vector<int> pwr2(binary_len);
138         pwr2[0] = 1;
139         for (int j = 1; j < binary_len; ++j) {
140             pwr2[j] = pwr2[j - 1] * 2;
141         }
142         int j = 0;
143         while (j < n - 1) {
144             j++;
145             int lastbit = 0;
146             while (bits[lastbit]) {
147                 bits[lastbit] = false;
148                 lastbit++;
149             }
150             bits[lastbit] = true;
151             int val = 0;
152             for (int k = 0; k < binary_len; ++k) {
153                 if (bits[k]) {
154                     val += pwr2[binary_len - k - 1];
155                 }
156             }
157             rho[j] = val;
158         }
159     }
160
161     static int k1_init(int n, int p) {
162         assert(p < n);
163
164         int np = n / p;
165         int c;

```



```

166     for (c = 1; c < p; c *= np);
167     return n / c;
168 }
169
170 static void bspfft_init(int n, int p, int s, std::vector<double>& w0,
171                       std::vector<double>& w,
172                       std::vector<double>& tw, std::vector<int>& rho_np,
173                       std::vector<int>& rho_p) {
174     int np = n / p;
175     bitrev_init(rho_np);
176
177     if (p > 1) {
178         bitrev_init(rho_p);
179     }
180     int k1 = k1_init(n, p);
181     ufft_init(k1, w0);
182     ufft_init(np, w);
183
184     int ntw = 0;
185     for (int c = k1; c <= p; c *= np) {
186         double alpha = static_cast<double>(s % c) / static_cast<double>(c);
187         twiddle_init(np, alpha, rho_np, tw, 2 * ntw * np);
188         ntw++;
189     }
190 }
191
192 static void calcError(const std::vector<double>& xlocal,
193                     const std::vector<double>& xarr, int n, int p,
194                     int s) {
195     double error = 0.0;
196     int c = 0;
197     for (; c <= n / p; c++) {
198         double lerror = std::abs(xlocal[c] - xarr[c]);
199         error += lerror;
200     }
201     std::cout << s << ": local error is "
202               << (error / static_cast<double>(n)) << std::endl;
203 }
204
205 void bspredistr(bsp_array<double>& x, int n, int p, int s, int c0, int c1,
206                bool rev, const std::vector<int>& rho_p) {
207     assert(1 <= c0);
208     assert(c0 <= c1);
209     assert(c1 <= p);
210
211     auto xarr = x.BSPunsafe_access();
212     int np = (int) xarr.size() / 2;
213     int ratio = c1 / c0;
214     int size = std::max(np / ratio, 1);
215     int npackets = np / size;
216     std::vector<double> tmp(2 * size);
217
218     assert (p <= n);
219
220     int j0, j2;
221     if (rev) {
222         j0 = rho_p[s] % c0;
223         j2 = rho_p[s] / c0;

```

```

224     } else {
225         j0 = s % c0;
226         j2 = s / c0;
227     }
228
229     for (int j = 0; j < npackets; j++) {
230         int jglob = j2 * c0 * np + j * c0 + j0;
231         int destproc = (jglob / (c1 * np)) * c1 + jglob % c1;
232         int destindex = (jglob % (c1 * np)) / c1;
233         for (int r = 0; r < size; ++r) {
234             int tr = 2 * r;
235             int tjrr = 2 * (j + r * ratio);
236             tmp[tr] = xarr[tjrr];
237             tmp[tr + 1] = xarr[tjrr + 1];
238         }
239         assert(destproc <= p);
240         assert(destindex < xarr.size() / 2);
241         x.bsp_put(tmp, destproc, 0, 2 * destindex, 2 * size);
242     }
243     bsp_sync();
244 }
245
246 void bspfft(bsp_array<double>& x, int n, int p, int s, bool sign,
247            const std::vector<double>& w0, const std::vector<double>& w,
248            const std::vector<double>& tw,
249            const std::vector<int>& rho_np, const std::vector<int>& rho_p) {
250     int np = n / p;
251     int k1 = k1_init(n, p);
252     permute(x.BSPunsafe_access(), np, rho_np, __LINE__);
253     bool rev = true;
254
255     for (int r = 0; r < np / k1; r++) {
256         fft(x.BSPunsafe_access(), 2 * r * k1, k1, sign, w0);
257     }
258
259     int c0 = 1;
260     int ntw = 0;
261
262     for (int c = k1; c <= p; c *= np) {
263         bspredistr(x, n, p, s, c0, c, rev, rho_p);
264         rev = false;
265         twiddle(x.BSPunsafe_access(), np, sign, tw, 2 * ntw * np);
266         fft(x.BSPunsafe_access(), 0, np, sign, w);
267         c0 = c;
268         ntw++;
269     }
270
271     if (!sign) {
272         auto& xarr = x.BSPunsafe_access();
273         double ninv = 1.0 / static_cast<double>(n);
274         for (int j = 0; j < 2 * np; ++j) {
275             xarr[j] *= ninv;
276         }
277     }
278 }
279
280 std::vector<double>& fft(const std::vector<double>& xlocal) {
281     int s = bsp_pid();

```

```

282
283     int n = global_n / 2;
284     int p = bsp_nprocs();
285     int k1 = k1_init(n, p);
286     std::vector<double> w0(k1);
287     std::vector<double> w(n / p);
288     std::vector<double> tw(2 * n / p);
289     std::vector<int> rho_np(n / p);
290     std::vector<int> rho_p(p);
291
292     bsp_array<double> x = get_array(xlocal);
293
294     auto time = std::chrono::high_resolution_clock::now();
295     if (p == 1) {
296         bspfft_init(n, p, s, w0, w, tw, rho_np, rho_p);
297         permute(x.BSPUnsafe_access(), n, rho_np, __LINE__);
298         ufft(x.BSPUnsafe_access(), 0, n, true, w);
299         permute(x.BSPUnsafe_access(), n, rho_np, __LINE__);
300         ufft(x.BSPUnsafe_access(), 0, n, false, w);
301         auto& xi = x.BSPUnsafe_access();
302         double ninv = 1.0 / static_cast<double>(n);
303         for (int j = 0; j < 2 * n; ++j) {
304             xi[j] *= ninv;
305         }
306         calcError(x.BSPUnsafe_access(), global_x, n, p, s);
307         auto t = std::chrono::high_resolution_clock::now();
308         auto tc = std::chrono::duration_cast<std::chrono::milliseconds>(
309             t - time).count();
310         std::cout << "Parallel part took " << tc << " ms." << std::endl;
311         return x.BSPUnsafe_access();
312     }
313
314     bsp_sync();
315
316     bspfft_init(n, p, s, w0, w, tw, rho_np, rho_p);
317     std::cout << s << ": calling bspfft with n=" << n << ", p=" << p
318         << ", s=" << s << std::endl;
319     bspfft(x, n, p, s, true, w0, w, tw, rho_np, rho_p);
320     bsp_sync();
321
322     std::cout << s << ": calling bspfft (inv) with n=" << n << ", p=" << p
323         << ", s=" << s << std::endl;
324     bspfft(x, n, p, s, false, w0, w, tw, rho_np, rho_p);
325
326     if (s == 0) {
327         auto t = std::chrono::high_resolution_clock::now();
328         auto tc = std::chrono::duration_cast<std::chrono::milliseconds>(
329             t - time).count();
330         std::cout << "Parallel part took " << tc << " ms." << std::endl;
331     }
332
333     calcError(x.BSPUnsafe_access(), xlocal, n, p, s);
334     return x.BSPUnsafe_access();
335 }
336
337 void parallel_function() override {
338
339     bsp_array<double> local_x = get_empty_array<double>(1);

```

```

340     if (bsp_pid() == 0) {
341         auto p = bsp_nprocs();
342         for (int s = 0; s < p; ++s) {
343             std::vector<double> xlocal(global_n / p);
344             int c = 0;
345             int n1 = global_n / 2;
346             for (int i = 0; i < 2 * n1; ++i) {
347                 if ((i / 2) % p == s) {
348                     xlocal[c++] = global_x[i];
349                 }
350             }
351             local_x.bsp_put(xlocal, s);
352         }
353     }
354
355     bsp_sync();
356     fft(local_x.BSPunsafe_access());
357
358 }
359 };
360
361 int main(int argc, char* argv[]) {
362     if (argc < 3) {
363         std::cerr << "Usage: " << argv[0] << " <N> <P>" << std::endl;
364         std::cerr << " where N is the problem size (power of 2)" << std::endl;
365         std::cerr << "         P is the number of threads used (power of 2)"
366                 << std::endl;
367         return -1;
368     }
369     int n, p;
370     n = std::stoi(argv[1]);
371     p = std::stoi(argv[2]);
372     std::vector<double> global_x(n);
373     for (int i = 0; i < n; i += 2) {
374         global_x[i] = static_cast<double>(i) / 2.0;
375     }
376
377     std::vector<std::unique_ptr<bsp_node>> nodes;
378     for (size_t i{0}; i < p; ++i) {
379         nodes.push_back(std::make_unique<BSPfft>(n, global_x));
380     }
381
382     bsp_program mbsp(std::move(nodes));
383     mbsp.start();
384 }

```

## BSPlu.cpp

```

1  #include <bsp_program.hpp>
2  #include <cmath>
3  #include <limits>
4  #include <random>
5
6  class BSPlu : public bsp_node {
7
8  private:
9

```

```

10  int param_n, param_M, param_N;
11  std::mt19937 rand;
12  std::uniform_real_distribution<double> dist{-32.768, 32.768};
13
14  template<typename T>
15  using matrix = std::vector<std::vector<T>>;
16
17  using clock = std::chrono::high_resolution_clock;
18  using milliseconds = std::chrono::milliseconds;
19
20  inline int nloc(int p, int s, int n) const {
21      return (n + p - s - 1) / p;
22  }
23
24  void bsplu(int M, int N, int n, int s, int t,
25            std::vector<int>& raw_pi, matrix<double>& a) {
26
27      int p = bsp_pid();
28      int P = bsp_nprocs();
29      int p_i = p / N;
30      int p_j = p % N;
31
32      assert(p_i == s);
33      assert(p_j == t);
34
35      auto pi = get_array<int>(&raw_pi);
36
37      std::vector<bsp_array<double>> A;
38      for (int i{0}; i < a.size(); ++i) {
39          A.emplace_back(get_array<double>(&a.at(i)));
40      }
41      std::vector<double> raw_pivots(M);
42      auto pivots = get_array<double>(&raw_pivots);
43      std::vector<int> raw_pivot_info(3);
44      auto pivot_info = get_array<int>(&raw_pivot_info);
45      std::vector<double> raw_pivot_row(n / N);
46      auto pivot_row = get_array<double>(&raw_pivot_row);
47      std::vector<double> raw_pivot_col(n / M);
48      auto pivot_col = get_array<double>(&raw_pivot_col);
49
50      bsp_sync();
51
52
53      for (int i{0}; i < nloc(P, p, n); ++i) {
54          raw_pi.at(i) = p + i * P;
55      }
56
57      for (int k{0}; k < n - 1; ++k) {
58          int kdN = k / N;
59          raw_pivot_info.at(0) = nloc(M, p_i, k);
60          int kdM = raw_pivot_info.at(0);
61          double temp_m;
62          double absmax = 0.0;
63          if (p_j == k % N) {
64              for (int i{raw_pivot_info.at(0)}; i < a.size(); ++i) {
65                  if (std::abs(a.at(i).at(kdN)) > absmax) {
66                      raw_pivot_info.at(0) = i;
67                      absmax = std::abs(a.at(i).at(kdN));

```

```

68     }
69   }
70   temp_m = (raw_pivot_info.at(0) == a.size()) ?
71     -std::numeric_limits<double>::infinity() :
72     a.at(raw_pivot_info.at(0)).at(kdN);
73   for (int i{0}; i < M; ++i) {
74     pivots.bsp_put(temp_m, i * N + p_j, p_i);
75   }
76 }
77
78 bsp_sync();
79
80 raw_pivot_info.at(1) = p_i;
81 raw_pivot_info.at(2) = k % N;
82 if (p_j == raw_pivot_info.at(2)) {
83   temp_m = raw_pivots.at(0);
84   raw_pivot_info.at(1) = 0;
85   for (int i{1}; i < M; ++i) {
86     double temp = raw_pivots.at(i);
87     if (temp > temp_m) {
88       temp_m = temp;
89       raw_pivot_info.at(1) = i;
90     }
91   }
92   auto el = pivot_info.bsp_direct_get(
93     raw_pivot_info.at(1) * N + raw_pivot_info.at(2), 0);
94   pivot_info.BSPunsafe_access().at(0) = el;
95   for (int i{kdM}; i < a.size(); ++i) {
96     a.at(i).at(kdN) /= temp_m;
97   }
98   if (p_i == raw_pivot_info.at(1)) {
99     a.at(raw_pivot_info.at(0)).at(kdN) = temp_m;
100  }
101
102  for (int j{0}; j < N; j++) {
103    pivots.bsp_put(raw_pivots, p_i * N + j,
104      raw_pivot_info.at(1), raw_pivot_info.at(1),
105      1);
106    pivot_info.bsp_put(raw_pivot_info, p_i * N + j);
107  }
108 }
109
110 bsp_sync();
111
112 temp_m = raw_pivots.at(raw_pivot_info.at(1));
113
114 assert(temp_m != 0.0);
115
116 int temp_m_rho =
117   (M * raw_pivot_info.at(0) + raw_pivot_info.at(1)) % P;
118 int temp_l_cyc =
119   (M * raw_pivot_info.at(0) + raw_pivot_info.at(1)) / P;
120
121 if (p == temp_m_rho) {
122   pi.bsp_get(k % P, k / P, temp_l_cyc, 1);
123 }
124 if (p == k % P) {
125   pi.bsp_get(temp_m_rho, temp_l_cyc, k / P, 1);

```

```

126     }
127     if (p_i == raw_pivot_info.at(1)) {
128         auto arr = A.at(k / M).bsp_direct_get((k % M) * N + p_j);
129         if (arr.size() == 2)
130             std::cout << "size 2" << std::endl;
131         A.at(raw_pivot_info.at(0))
132             .bsp_put(arr);
133     }
134     if (p_i == k % M) {
135         auto arr = A.at(raw_pivot_info.at(0)).bsp_direct_get(
136             raw_pivot_info.at(1) * N + p_j);
137         if (arr.size() == 2)
138             std::cout << "size 2" << std::endl;
139         A.at(k / M)
140             .bsp_put(arr);
141     }
142
143     if (p_i == raw_pivot_info.at(1)) {
144         for (int j{0}; j < a.at(0).size(); ++j) {
145             raw_pivot_row.at(j) = a.at(raw_pivot_info.at(0)).at(j);
146         }
147         for (int i{0}; i < M; ++i) {
148             if (i == p_i) continue;
149             pivot_row.bsp_put(raw_pivot_row, i * N + p_j);
150         }
151     }
152
153     if (p_j == raw_pivot_info.at(2)) {
154         for (int i{0}; i < a.size(); ++i) {
155             raw_pivot_col.at(i) = a.at(i).at(k / N);
156         }
157         for (int j{0}; j < N; ++j) {
158             if (j == p_j) continue;
159             pivot_col.bsp_put(raw_pivot_col, p_i * N + j);
160         }
161     }
162
163     bsp_sync();
164
165     if (p_i == raw_pivot_info.at(1)) {
166         auto el = A.at(raw_pivot_info.at(0)).bsp_direct_get(
167             p_i * N + k % N);
168         std::copy_n(el.begin() + k / N, 1,
169             raw_pivot_col.begin() + raw_pivot_info.at(0));
170         //pivot_col.bsp_put(el, k/N, raw_pivot_info.at(0), 1);
171     }
172
173     bsp_sync();
174
175     int  istart = k / M;
176     int  jstart = k / N;
177     // a.at(istart) = A.at(istart).BSPunsafe_access();
178     if (p_i <= k % M) istart++;
179     if (p_j <= k % N) jstart++;
180     auto& raw_pivot_row2 = pivot_row.BSPunsafe_access();
181     auto& raw_pivot_col2 = pivot_col.BSPunsafe_access();
182     for (int i{istart}; i < a.size(); i++) {
183         //a.at(i) = A.at(i).BSPunsafe_access();

```

```

184         for (int j{jstart}; j < a.at(i).size(); ++j) {
185             a.at(i).at(j) -=
186                 raw_pivot_row2.at(j) * raw_pivot_col2.at(i);
187         }
188     }
189     bsp_sync();
190 }
191 }
192
193 public:
194
195     BSPlu(int n, int M, int N, unsigned int seed) : param_n{n}, param_M{M},
196                                                    param_N{N}, rand{seed} {
197
198     };
199
200     void parallel_function() override {
201         auto time = clock::now();
202
203         int n = param_n;
204         int M = param_M;
205         int N = param_N;
206         int s = bsp_pid() / N;
207         int t = bsp_pid() % N;
208         int nlr = nloc(M, s, n);
209         int nlc = nloc(N, t, n);
210         matrix<double> a(nlr);
211         for (int i{0}; i < nlr; ++i) {
212             a.at(i).resize(nlc);
213             for (int j{0}; j < nlc; ++j) {
214                 a.at(i).at(j) = dist(rand);
215             }
216         }
217
218         std::vector<int> pi(nloc(bsp_nprocs(), bsp_pid(), n));
219         auto timec = std::chrono::duration_cast<milliseconds>(
220             clock::now() - time).count();
221         std::cout << "Processor " << bsp_pid() << ": init took " << timec
222             << " ms." << std::endl;
223
224         time = clock::now();
225         bsplu(M, N, n, s, t, pi, a);
226         timec = std::chrono::duration_cast<milliseconds>(
227             clock::now() - time).count();
228         if (bsp_pid() == 0)
229             std::cout << "Processor " << bsp_pid() << ": LU took " << timec
230                 << " ms." << std::endl;
231     }
232 };
233
234 int main(int argc, char* argv[]) {
235     if (argc < 4) {
236         std::cerr << "Usage: " << argv[0] << " <n> <M> <N>" << std::endl;
237         std::cerr << "           will start LU decomposition on a n by n matrix"
238             << std::endl;
239         std::cerr << "           using M times N threads." << std::endl;
240         return 1;
241     }

```



```

242
243     int n = std::stoi(argv[1]);
244     int M = std::stoi(argv[2]);
245     int N = std::stoi(argv[3]);
246
247     std::random_device rd;
248     std::vector<std::unique_ptr<bsp_node>> nodes;
249     for (int i{0}; i < M * N; ++i) {
250         nodes.emplace_back(std::make_unique<BSPlu>(n, M, N, rd()));
251     }
252     auto t0 = std::chrono::high_resolution_clock::now();
253     auto f0 = [&t0]() {
254         t0 = std::chrono::high_resolution_clock::now();
255     };
256
257     auto f1 = [&t0]() {
258         auto t1 = std::chrono::high_resolution_clock::now();
259         auto diff = std::chrono::duration_cast<std::chrono::milliseconds>(
260             t1 - t0);
261         std::cout << "Main part took " << diff.count() << " ms." << std::endl;
262     };
263
264     bsp_program computation(std::move(nodes), f0, f1);
265     computation.start();
266     return 0;
267 }

```

## B.3 Sequential test programs

### sequential\_inprod.cpp

```

1  #include <string>
2  #include <chrono>
3  #include <iostream>
4
5  int main(int argc, char* argv[]) {
6
7      if (argc < 2) return 1;
8
9      int n = std::stoi(argv[1]);
10
11     auto x = new double[n]();
12
13     for (int i{0}; i < n; ++i) {
14         x[i] = i + 1;
15     }
16
17     double inprod = 0.0;
18     auto t1 = std::chrono::high_resolution_clock::now();
19     for (int i{0}; i < n; ++i) {
20         inprod += x[i] * x[i];
21     }
22     auto t2 = std::chrono::high_resolution_clock::now() - t1;
23     std::cout << "Processor 0: sum of squares up to " << n << "*" << n << " is "
24         << inprod << std::endl;
25     std::cout << "Processor 0: local time taken is "
26         << std::chrono::duration_cast<std::chrono::milliseconds>(

```

```

27         t2).count() << std::endl;
28
29     auto dn = static_cast<double>(n);
30     dn *= n + 1.0;
31     dn *= 2.0 * n + 1.0;
32     dn /= 6.0;
33     std::cout << "Checksum: " << dn << std::endl;
34
35     delete[] x;
36
37     return 0;
38 }

```

### sequential\_sorting.cpp

```

1  #include <random>
2  #include <vector>
3  #include <iostream>
4  #include <algorithm>
5  #include <numeric>
6  #include <chrono>
7
8  int main(int argc, char* argv[]) {
9      if (argc < 2) {
10         std::cerr << "Usage: " << argv[0] << " <N> (seed)" << std::endl;
11         std::cerr << " where N is the problem size (power of 2)" << std::endl;
12         std::cerr
13             << "         seed is an optional seed for permutations "
14             << "(leave blank to randomize it)"
15             << std::endl;
16         return -1;
17     }
18     int n, s;
19     n = std::stoi(argv[1]);
20     s = (argc >= 3 ? std::stoi(argv[2]) : -1);
21     auto t = std::chrono::high_resolution_clock::now();
22     if (s == -1) {
23         std::random_device x;
24         s = x();
25     }
26     std::mt19937 mtw(s);
27     std::vector<int> data(n);
28     std::iota(data.begin(), data.end(), 0);
29     std::shuffle(data.begin(), data.end(), mtw);
30
31     int* dat2 = new int[n]();
32
33     for (int i = 0; i < n; ++i) dat2[i] = data.at(i);
34
35     auto t1 = std::chrono::duration_cast<std::chrono::milliseconds>(
36         std::chrono::high_resolution_clock::now() - t).count();
37     std::cout << "Ended vector generation and shuffling (spent " << t1 << " ms)"
38         << "\nstarting sequential part" << std::endl;
39
40     auto t2 = std::chrono::high_resolution_clock::now();
41
42     std::sort(data.begin(), data.end());

```

```

43
44     auto t3p = std::chrono::high_resolution_clock::now();
45
46     std::sort(dat2, dat2 + n);
47
48     auto t4 = std::chrono::high_resolution_clock::now();
49
50     bool passed = true;
51     for (int j{0}; j < n; j++) {
52         if (data[j] != j) passed = false;
53     }
54
55     auto t3 = std::chrono::duration_cast<std::chrono::milliseconds>(
56         t3p - t2).count();
57     auto t4c = std::chrono::duration_cast<std::chrono::milliseconds>(
58         t4 - t3p).count();
59
60     std::cout << "Check " << (passed ? "passed" : "failed") << std::endl;
61     std::cout << "Parallel part took " << t3 << " ms." << std::endl;
62     std::cout << "Plain array part took " << t4c << " ms." << std::endl;
63
64     delete[] dat2;
65
66 }

```

## sequential\_fft.cpp

```

1  #include <complex>
2  #include <iostream>
3  #include <vector>
4  #include <chrono>
5  #include <cassert>
6
7  constexpr static double PI = 3.141592653589793;
8
9  void ufft(std::vector<double>& x, int offset, int n, bool sign,
10          const std::vector<double>& w) {
11
12     for (int k = 2; k <= n; k *= 2) {
13         int nk = n / k;
14         for (int r = 0; r < nk; ++r) {
15             int rk = 2 * r * k;
16             for (int j = 0; j < k; j += 2) {
17                 double wr = w[j * nk];
18                 double wi;
19                 if (sign) {
20                     wi = w[j * nk + 1];
21                 } else {
22                     wi = -w[j * nk + 1];
23                 }
24
25                 int j0 = rk + j + offset;
26                 int j1 = j0 + 1;
27                 int j2 = j0 + k;
28                 int j3 = j2 + 1;
29
30                 double taur = wr * x[j2] - wi * x[j3];

```

```

31         double tau_i = w_i * x[j2] + w_r * x[j3];
32
33         x[j2] = x[j0] - tau_r;
34         x[j3] = x[j1] - tau_i;
35         x[j0] += tau_r;
36         x[j1] += tau_i;
37     }
38 }
39 }
40 }
41
42 void ufft_init(int n, std::vector<double>& w) {
43     assert(w.size() == n);
44
45     if (n == 1) return;
46
47     w[0] = 1.0;
48     w[1] = 0.0;
49
50     if (n == 4) {
51         w[2] = 0.0;
52         w[3] = -1.0;
53     } else if (n >= 8) {
54         double theta = -2.0 * PI / static_cast<double>(n);
55         for (int j = 1; j <= n / 8; j++) {
56             w[2 * j] = std::cos(j * theta);
57             w[2 * j + 1] = std::sin(j * theta);
58         }
59         for (int j = 0; j < n / 8; j++) {
60             int n4j = n / 4 - j;
61             w[2 * n4j] = -w[2 * j + 1];
62             w[2 * n4j + 1] = -w[2 * j];
63         }
64         for (int j = 1; j < n / 4; j++) {
65             int n2j = n / 2 - j;
66             w[2 * n2j] = -w[2 * j];
67             w[2 * n2j + 1] = w[2 * j + 1];
68         }
69     }
70 }
71
72
73 void twiddle_init(int n, double alpha, const std::vector<int>& rho,
74                 std::vector<double>& w, int offset) {
75     double theta = -2.0 * PI * alpha / static_cast<double>(n);
76     for (int j = 0; j < n; ++j) {
77         double rt = static_cast<double>(rho[j]) * theta;
78         w[offset + 2 * j] = std::cos(rt);
79         w[offset + 2 * j + 1] = std::sin(rt);
80     }
81 }
82
83 void permute(std::vector<double>& x, int n, const std::vector<int>& sigma) {
84     assert(x.size() / 2 == sigma.size());
85
86     for (int j = 0; j < n; ++j) {
87         if (j < sigma[j]) {
88             int j0 = 2 * j;

```

```

89         int j1 = j0 + 1;
90         int j2 = 2 * sigma[j];
91         int j3 = j2 + 1;
92         double tmpr = x[j0];
93         double tmpi = x[j1];
94         x[j0] = x[j2];
95         x[j1] = x[j3];
96         x[j2] = tmpr;
97         x[j3] = tmpi;
98     }
99 }
100 }
101
102 void bitrev_init(std::vector<int>& rho) {
103     int n = rho.size();
104
105     auto binary_len = static_cast<int>(std::ceil(
106         std::log(static_cast<double>(n)) / std::log(2.0)));
107     std::vector<bool> bits(binary_len);
108     std::vector<int> pwrs(binary_len);
109     pwrs[0] = 1;
110     for (int j = 1; j < binary_len; ++j) {
111         pwrs[j] = pwrs[j - 1] * 2;
112     }
113     int j = 0;
114     while (j < n - 1) {
115         j++;
116         int lastbit = 0;
117         while (bits[lastbit]) {
118             bits[lastbit] = false;
119             lastbit++;
120         }
121         bits[lastbit] = true;
122         int val = 0;
123         for (int k = 0; k < binary_len; ++k) {
124             if (bits[k]) {
125                 val += pwrs[binary_len - k - 1];
126             }
127         }
128         rho[j] = val;
129     }
130 }
131
132
133 void fft_init(int n, std::vector<double>& w,
134             std::vector<double>& tw, std::vector<int>& rho_np) {
135     bitrev_init(rho_np);
136     ufft_init(n, w);
137
138     twiddle_init(n, 0, rho_np, tw, 0);
139 }
140
141 static void
142 calcError(const std::vector<double>& xlocal, const std::vector<double>& xarr) {
143     double error = 0.0;
144     for (int c{0}; c < xlocal.size(); c++) {
145         double lerror = std::abs(xlocal[c] - xarr[c]);
146         error += lerror;

```

```

147     }
148     std::cout << "local error is "
149             << (error / static_cast<double>(xlocal.size())) << std::endl;
150 }
151
152 int main(int argc, char* argv[]) {
153     if (argc < 2) {
154         std::cerr << "Usage: " << argv[0] << " <N>" << std::endl;
155         std::cerr << " where N is the problem size (power of 2)" << std::endl;
156         return -1;
157     }
158     int n;
159     n = std::stoi(argv[1]);
160     std::vector<double> data(n);
161
162     n /= 2;
163     for (int i = 0; i < n; i += 2) {
164         data[i] = static_cast<double>(i) / 2.0;
165     }
166
167     std::vector<double> old(data);
168
169     std::vector<double> w(n);
170     std::vector<double> tw(2 * n);
171     std::vector<int> rho_np(n);
172
173     auto t = std::chrono::high_resolution_clock::now();
174     fft_init(n, w, tw, rho_np);
175     // forward fft
176     permute(data, n, rho_np);
177     ufft(data, 0, n, true, w);
178
179     // inverse fft
180     permute(data, n, rho_np);
181     ufft(data, 0, n, false, w);
182
183     double ninv = 1.0 / static_cast<double>(n);
184     for (int j = 0; j < 2 * n; ++j) {
185         data[j] *= ninv;
186     }
187
188     calcError(data, old);
189     auto t1 = std::chrono::high_resolution_clock::now();
190
191     std::cout << "Sequential part took "
192             << std::chrono::duration_cast<std::chrono::milliseconds>(
193                 t1 - t).count() << " ms." << std::endl;
194
195     calcError(data, old);
196
197     return 0;
198 }
199
200

```

## sequential\_lu.cpp

```
1 #include <iostream>
2 #include <chrono>
3 #include "lib/src/linalg.h"
4
5 int main(int argc, char* argv[]) {
6     if (argc < 2) {
7         std::cerr << "Usage: " << argv[0] << " <n>" << std::endl;
8         std::cerr << "          will start LU decomposition on a n by n matrix."
9             << std::endl;
10        return 1;
11    }
12
13    int n = std::stoi(argv[1]);
14
15    alglib::real_2d_array a;
16    alglib::integer_1d_array piv;
17
18    a.setlength(n, n);
19    piv.setlength(n);
20
21    for (int i{0}; i < n; ++i) {
22        for (int j{0}; j < n; ++j) {
23            a[i][j] = (alglib::randomreal() * 65.536) - 32.768;
24        }
25    }
26
27    auto t1 = std::chrono::high_resolution_clock::now();
28    alglib::rmatrixlu(a, n, n, piv);
29    auto t2 = std::chrono::high_resolution_clock::now();
30
31    auto time = std::chrono::duration_cast<std::chrono::milliseconds>(t2 - t1);
32    std::cout << "Sequential computation took " << time.count() << " ms."
33        << std::endl;
34 }
```

## B.4 Miscellaneous

### commstress.cpp

```
1 #include <bsp_program.hpp>
2 #include <random>
3 #include <iostream>
4
5
6 /**
7  * benchmarks the communication aspects of the bsp pattern
8  */
9 struct communication_stress : public bsp_node {
10     int NITERS = 10000;
11
12     void parallel_function() override {
13         using clock = std::chrono::high_resolution_clock;
14         using duration = std::chrono::nanoseconds;
15
16         std::random_device rd;
```

```

17     std::mt19937 mersenne(rd());
18     std::uniform_int_distribution uid(0, NITERS - 1);
19
20     long long int rng_duration = 0;
21     if (bsp_pid() == 0) {
22         // benchmarking the time spent on an RNG call
23         int x = 0;
24         auto rng_start = clock::now();
25         for (int i{0}; i < NITERS; ++i) {
26             if (i % 2) x += uid(mersenne);
27             else x -= uid(mersenne);
28         }
29         auto rng_stop = clock::now();
30         auto dur = std::chrono::duration_cast<duration>(
31             rng_stop - rng_start).count();
32         rng_duration = dur / NITERS;
33         // Use the value
34         std::cout << "RNG benchmark: a single RNG call takes "
35                 << rng_duration << " ns (value is " << x
36                 << "), whole op takes "
37                 << dur << "ns" << std::endl;
38     }
39
40     // Test #1: insertion of random values inside random workers' variables
41
42     bsp_variable<int> var = get_variable<int>(0);
43
44     auto start = clock::now();
45     for (size_t i{0}; i < NITERS; ++i) {
46         size_t next = uid(mersenne) % bsp_nprocs();
47         var.bsp_put(uid(mersenne), next);
48         bsp_sync();
49     }
50
51     if (bsp_pid() == 0) {
52         auto time = std::chrono::duration_cast<duration>(
53             clock::now() - start).count() - (2 * NITERS * rng_duration);
54         std::cout << "--- PHASE 1 - VARIABLES ---" << std::endl;
55         std::cout << "Spent " << time << "ns" << std::endl;
56         std::cout << "That is, " << time / NITERS << "ns per operation"
57                 << std::endl;
58     }
59
60     // Test #2: replacement of arrays inside random worker's memory
61
62     bsp_array<int> arr1 = get_empty_array<int>(NITERS);
63     std::vector<int> swap(NITERS);
64
65     auto gen = [&]() { return uid(mersenne); };
66     std::generate(swap.begin(), swap.end(), gen);
67
68     start = clock::now();
69     for (size_t i{0}; i < NITERS; ++i) {
70         size_t next = uid(mersenne) % bsp_nprocs();
71         arr1.bsp_put(swap, next);
72         bsp_sync();
73     }
74

```



```

75     if (bsp_pid() == 0) {
76         auto time = std::chrono::duration_cast<duration>(
77             clock::now() - start).count() - (NITERS * rng_duration);
78         std::cout << "--- PHASE 2 - ARR_SWAP ---" << std::endl;
79         std::cout << "Spent " << time << "ns" << std::endl;
80         std::cout << "That is, " << time / NITERS << "ns per operation"
81             << std::endl;
82     }
83
84     // Test #3: replacement of portion of arrays inside
85     // random positions in a random worker's memory
86
87     bsp_array<int> arr2 = get_empty_array<int>(NITERS);
88     std::vector<int> empl(10);
89     std::generate(empl.begin(), empl.end(), gen);
90
91     start = clock::now();
92     for (size_t i{0}; i < NITERS; ++i) {
93         size_t next = uid(mersenne) % bsp_nprocs();
94         size_t pos = uid(mersenne) % (NITERS - 10);
95         arr2.bsp_put(empl, next, pos, 10);
96         bsp_sync();
97     }
98
99     if (bsp_pid() == 0) {
100         auto time = std::chrono::duration_cast<duration>(
101             clock::now() - start).count() - (2 * NITERS * rng_duration);
102         std::cout << "--- PHASE 3 - ARR_EMPLACE ---" << std::endl;
103         std::cout << "Spent " << time << "ns" << std::endl;
104         std::cout << "That is, " << time / NITERS << "ns per operation"
105             << std::endl;
106     }
107
108     // Test #4: replacement of elements in random positions
109     // in arrays inside random worker's memory
110
111     bsp_array<int> arr3 = get_empty_array<int>(NITERS);
112
113     start = clock::now();
114     for (size_t i{0}; i < NITERS; ++i) {
115         size_t next = uid(mersenne) % bsp_nprocs();
116         for (size_t j{0}; j < 50; ++j) {
117             size_t pos = uid(mersenne) % NITERS;
118             arr3.bsp_put(bsp_pid(), next, pos);
119         }
120         bsp_sync();
121     }
122
123     if (bsp_pid() == 0) {
124         auto time = std::chrono::duration_cast<duration>(
125             clock::now() - start).count() - (NITERS * rng_duration) -
126             (50 * NITERS * rng_duration);
127         std::cout << "--- PHASE 4 - ARR_PUT ---" << std::endl;
128         std::cout << "Spent " << time << "ns" << std::endl;
129         std::cout << "That is, " << time / (NITERS * 50)
130             << "ns per operation" << std::endl;
131     }
132 }

```

```

133 };
134
135 int main() {
136     std::vector<std::unique_ptr<bsp_node>> nodes;
137     for (size_t i{0}; i < 4; ++i) {
138         nodes.push_back(std::make_unique<communication_stress>());
139     }
140     bsp_program mbsp(std::move(nodes));
141     mbsp.start();
142 }

```

## BSPbench.cpp

```

1  #include <array>
2  #include <cmath>
3  #include <bsp_program.hpp>
4
5  struct bspbench : public bsp_node {
6
7      int maxH, maxN, niters;
8
9      const double MEGA = 1000000.0;
10
11     bspbench(int _maxH, int _maxN, int _niters) :
12         maxH{_maxH}, maxN{_maxN}, niters{_niters} {};
13
14     std::array<double, 2>
15     leastsquares(int h0, int h1, const std::vector<double>& t) const {
16         std::array<double, 2> ret{};
17         auto nh = static_cast<double>(h1 - h0 + 1);
18
19         double sumt = 0.0;
20         double sumth = 0.0;
21
22         for (int h = h0; h <= h1; ++h) {
23             sumt += t[h];
24             sumth += t[h] * h;
25         }
26         double sumh = static_cast<double>(h1 * h1 - h0 * h0 + h1 + h0) / 2;
27         double sumhh = static_cast<double>(h1 * (h1 + 1) * (2 * h1 + 1) -
28             (h0 - 1) * h0 * (2 * h0 - 1)) / 6;
29
30         if (std::abs(nh) > std::abs(sumh)) {
31             double a = sumh / nh;
32             ret[0] = (sumth - a * sumt) / (sumhh - a * sumh);
33             ret[1] = (sumt - sumh * ret[0]) / nh;
34         } else {
35             double a = nh / sumh;
36             ret[0] = (sumt - a * sumth) / (sumh - a * sumhh);
37             ret[1] = (sumth - sumhh * ret[0]) / sumh;
38         }
39
40         return ret;
41     }
42
43     void parallel_function() override {
44

```

```

45     double r = 0.0;
46
47     std::vector<double> x(maxN);
48     std::vector<double> y(maxN);
49     std::vector<double> z(maxN);
50
51     std::vector<int> destproc(maxH);
52     std::vector<int> destindex(maxH);
53     std::vector<double> src(maxH);
54
55     std::vector<double> t(maxH + 1);
56
57     int p = bsp_nprocs();
58     int s = bsp_pid();
59
60     bsp_array<double> Time = get_empty_array<double>(p);
61     bsp_array<double> Dest = get_empty_array<double>(2 * maxH + p);
62
63     for (int n = 1; n <= maxN; n *= 2) {
64
65         double alpha = 1.0 / 3.0;
66         double beta = 4.0 / 9.0;
67         for (int i = 0; i < n; ++i) {
68             z[i] = x[i] = y[i] = i;
69         }
70
71         auto time0 = std::chrono::high_resolution_clock::now();
72         for (int iter = 0; iter < niters; ++iter) {
73             for (int i = 0; i < n; ++i) {
74                 y[i] += alpha * x[i];
75             }
76             for (int i = 0; i < n; ++i) {
77                 z[i] -= beta * x[i];
78             }
79         }
80         auto time1 = std::chrono::high_resolution_clock::now();
81         auto time = std::chrono::duration_cast<std::chrono::milliseconds>(
82             time1 - time0).count() / 1000.0;
83         Time.bsp_put(time, 0, s);
84         bsp_sync();
85
86         if (s == 0) {
87             auto time_arr = Time.BSPunsafe_access();
88             double mintime = time_arr[0];
89             double maxtime = time_arr[0];
90             for (int s1 = 1; s1 < p; ++s1) {
91                 mintime = std::min(mintime, time_arr[s1]);
92                 maxtime = std::max(maxtime, time_arr[s1]);
93             }
94             if (mintime > 0.0) {
95                 int nflops = 4 * niters * n;
96                 r = 0.0;
97                 for (int s1 = 0; s1 < p; ++s1) {
98                     r += static_cast<double>(nflops) / time_arr[s1];
99                 }
100                r /= static_cast<double>(p);
101                std::cout << "n= "
102                    << n << " min= "

```

```

103         << nflops / (maxtime * MEGA) << " max= "
104         << nflops / (mintime * MEGA) << " av= "
105         << r / MEGA << " Mflop/s" << std::endl;
106     std::cout << " fool=" << y[n - 1] + z[n - 1] << std::endl;
107     } else {
108         std::cout << "minimum time is 0" << std::endl;
109     }
110 }
111 }
112
113 for (int h = 0; h <= maxH; ++h) {
114     for (int i = 0; i < h; ++i) {
115         src[i] = i;
116         if (p == 1) {
117             destproc[i] = 0;
118             destindex[i] = i;
119         } else {
120             destproc[i] = (s + 1 + i % (p - 1)) % p;
121             destindex[i] = s + (i / (p - 1)) * p;
122         }
123     }
124
125     auto time0 = std::chrono::high_resolution_clock::now();
126     for (int iter = 0; iter < niters; ++iter) {
127         for (int i = 0; i < h; i++) {
128             Dest.bsp_put(src[i], destproc[i], destindex[i]);
129         }
130         bsp_sync();
131     }
132     auto time1 = std::chrono::high_resolution_clock::now();
133     auto time = std::chrono::duration_cast<std::chrono::milliseconds>(
134         time1 - time0).count() / 1000.0;
135
136     if (s == 0) {
137         t[h] = (time * r) / static_cast<double>(niters);
138         std::cout << "Time of " << h << "-relation= "
139             << time / niters << " sec= " << t[h] << " flops"
140             << std::endl;
141     }
142 }
143
144 if (s == 0) {
145     auto temp = leastsquares(0, p, t);
146     std::cout << "Range h=0 to p: g= " << temp[0] << ", l= " << temp[1]
147         << std::endl;
148     temp = leastsquares(p, maxH, t);
149     std::cout << "Range h=p to HMAX: g= " << temp[0] << ", l= "
150         << temp[1] << std::endl;
151
152     std::cout << "The bottom line for this BSP computer is:"
153         << std::endl;
154     std::cout << "p= " << p << ", r= " << r / MEGA << " Mflop/s, g= "
155         << temp[0] << ", l= " << temp[1] << std::endl;
156 }
157 }
158 }
159
160 };

```

```

161
162 int main(int argc, char* argv[]) {
163
164     int nprocs, maxh, maxn, niters;
165
166     if (argc == 1) {
167         std::cerr << "Usage: " << argv[0] << " <P> (NITERS) (MAXN) (MAXH)"
168             << std::endl;
169         std::cerr << "<.> are obligatory parameters" << std::endl;
170         std::cerr << "(..) are optional" << std::endl;
171         return -1;
172     }
173
174     nprocs = std::stoi(argv[1]);
175
176     if (argc >= 3) {
177         niters = std::stoi(argv[2]);
178     } else {
179         niters = 10000;
180     }
181
182     if (argc >= 4) {
183         maxn = std::stoi(argv[3]);
184     } else {
185         maxn = 1024;
186     }
187
188     if (argc >= 5) {
189         maxh = std::stoi(argv[4]);
190     } else {
191         maxh = 128;
192     }
193
194     std::vector<std::unique_ptr<bsp_node>> nodes;
195     for (size_t i{0}; i < nprocs; ++i) {
196         nodes.push_back(std::make_unique<bspbench>(maxh, maxn, niters));
197     }
198     bsp_program mbsp(std::move(nodes));
199     mbsp.start();
200 }

```

## ff\_example.cpp

```

1  #include <iostream>
2  #include <iterator>
3  #include <algorithm>
4  #include "bsp_program.hpp"
5
6  struct ff_example : public bsp_node {
7      void parallel_function() override {
8          int s = bsp_pid();
9          auto v1 = (const std::vector<long>*) fastflow_input;
10         long count = 0;
11         int portion_size = v1->size() / bsp_nprocs();
12         for (int i{portion_size * s}; i < portion_size * (s + 1); ++i)
13             count += v1->at(i);
14         bsp_array<long> counts = get_empty_array<long>(bsp_nprocs());

```

```

15     counts.bsp_put(count, 0, s);
16     bsp_sync();
17     if (s == 0) {
18         long total = 0;
19         for (long l: counts.BSPunsafe_access()) {
20             total += l;
21         }
22         emit_output((void*) total);
23     }
24 }
25 };
26
27 struct generator : public ff::ff_node {
28     void* svc(void* in) override {
29         auto source = new std::vector<long>;
30         for (int i{0}; i < 10; ++i) source->push_back(i);
31         ff_send_out(source);
32         return EOS;
33     }
34 };
35
36 struct checker : public ff::ff_node {
37     void* svc(void* in) override {
38         if (in != GO_ON && in != EOS) {
39             auto val = (long) in;
40             if (val == 45) std::cout << "OK" << std::endl;
41             else std::cout << "KO" << std::endl;
42         }
43         return GO_ON;
44     }
45 };
46
47 int main() {
48     std::vector<std::unique_ptr<bsp_node>> nodes;
49     for (int i{0}; i < 2; ++i) {
50         nodes.push_back(std::make_unique<ff_example>());
51     }
52     auto pre_fun = []() {
53         std::cout << "Before BSP computation" << std::endl;
54     };
55     auto post_fun = []() {
56         std::cout << "After BSP computation" << std::endl;
57     };
58     bsp_program computation(std::move(nodes), pre_fun, post_fun);
59     generator g;
60     checker c;
61
62     ff::ff_Pipe<> pipe(g, computation, c);
63
64     if (pipe.run_and_wait_end() < 0)
65         std::cout << "Error in running pipe" << std::endl;
66     return 0;
67 }

```

## mwe.cpp

```
1 #include <iostream>
2 #include <iterator>
3 #include <algorithm>
4 #include <thread>
5 #include "bsp_program.hpp"
6
7 struct my_bsp_comp : public bsp_node {
8     void parallel_function() override {
9         int s = bsp_pid();
10        int n = bsp_nprocs();
11        auto v1 = get_variable<int>(10);
12        v1.bsp_put(s, (s + 1) % n);
13        bsp_sync();
14        std::vector<int> init{s, s, s, s};
15        auto a1 = get_array<int>(init);
16        a1.bsp_put(init, (s + 1) % n, 0, 2, 2);
17        bsp_sync();
18        auto a = a1.bsp_direct_get(s);
19        std::this_thread::sleep_for(std::chrono::seconds(s));
20        std::cout << s << ": v1 is " << v1.bsp_direct_get(s) << std::endl;
21        std::cout << s << ": a1 is ";
22        std::copy(a.begin(),
23                a.end(),
24                std::ostream_iterator<int>(std::cout, " "));
25        std::cout << std::endl;
26    }
27 };
28
29 int main() {
30     std::vector<std::unique_ptr<bsp_node>> nodes;
31     for (int i{0}; i < 2; ++i) {
32         nodes.push_back(std::make_unique<my_bsp_comp>());
33     }
34     auto pre_fun = []() {
35         std::cout << "Before BSP computation" << std::endl;
36     };
37     auto post_fun = []() {
38         std::cout << "After BSP computation" << std::endl;
39     };
40     bsp_program computation(std::move(nodes), pre_fun, post_fun);
41     computation.start();
42     return 0;
43 }
```

## unit\_tests.cpp

```
1 #include <bsp_program.hpp>
2
3 static bool ok = true;
4 static int total_tests = 35;
5
6 static int mod(int k, int n) {
7     return ((k %= n) < 0) ? k + n : k;
8 }
9
```

```

10 static void assertint(int a, int b, int id, int line) {
11     std::string s;
12     if (a != b) {
13         s += "Node " + std::to_string(id) + ": mismatch at line " +
14             std::to_string(line)
15             + " (expected " + std::to_string(b) + ", obtained " +
16             std::to_string(a) + ")";
17         std::cerr << s << std::endl;
18         ok = false;
19     }
20 }
21
22 static void assertint(int a, int b, int c, int id, int line) {
23     std::string s;
24     if (a != b && a != c) {
25         s += "Node " + std::to_string(id) + ": mismatch at line " +
26             std::to_string(line)
27             + " (expected " + std::to_string(b) + " or " + std::to_string(c) +
28             ", obtained " + std::to_string(a) + ")";
29         std::cerr << s << std::endl;
30         ok = false;
31     }
32 }
33
34 static void
35 assertvec(const std::vector<int>& a, const std::vector<int>& b, int id,
36           int line) {
37     std::string s;
38     auto printvect = [](const std::vector<int>& v) {
39         std::string s{"["};
40         for (const auto& el: v) {
41             s += std::to_string(el) + " ";
42         }
43         s.pop_back();
44         s += "];";
45         return s;
46     };
47     if (a != b) {
48         s += "Node " + std::to_string(id) + ": mismatch at line " +
49             std::to_string(line)
50             + " (expected " + printvect(b) + ", obtained " + printvect(a) +
51             ")";
52         std::cerr << s << std::endl;
53         ok = false;
54     }
55 }
56
57 static void assertvec(const std::vector<int>& a, const std::vector<int>& b,
58                       const std::vector<int>& c, int id, int line) {
59     std::string s;
60     auto printvect = [](const std::vector<int>& v) {
61         std::string s{"["};
62         for (const auto& el: v) {
63             s += std::to_string(el) + " ";
64         }
65         s.pop_back();
66         s += "];";
67         return s;

```



```

68     };
69     if (a != b && a != c) {
70         s += "Node " + std::to_string(id) + ": mismatch at line " +
71             std::to_string(line)
72             + " (expected " + printvect(b) + " or " + printvect(c) +
73             ", obtained " + printvect(a) + ")";
74         std::cerr << s << std::endl;
75         ok = false;
76     }
77 }
78
79 struct unit_tests : public bsp_node {
80
81     void parallel_function() override {
82
83         auto endfun = [this](bool end = false) {
84             static int passed_tests = 0;
85             if (end) {
86                 std::cout << "Number of tests executed successfully: "
87                     << passed_tests << "/" << total_tests << std::endl;
88                 if (passed_tests == total_tests)
89                     std::cout << "ALL TESTS COMPLETED SUCCESSFULLY!"
90                         << std::endl;
91                 else
92                     std::cerr
93                         << "Some tests failed, "
94                         << "refer to the logs for more information"
95                         << std::endl;
96             } else {
97                 if (ok) {
98                     passed_tests++;
99                     std::cout << "Test passed";
100                 } else std::cout << "Test failed";
101                 std::cout << " (" << passed_tests << "/" << total_tests << ")"
102                     << std::endl;
103                 ok = true;
104             }
105         };
106
107
108         if (bsp_pid() == 0) {
109             std::cout
110                 << "PART 1 : VARIABLES (SEQ)"
111                 << std::endl;
112             std::cout
113                 << "Test 1: bsp_put(const T& elem, int id)"
114                 << std::endl;
115         }
116
117         bsp_variable<int> v1 = get_variable(-1);
118
119         v1.bsp_put(bsp_pid(), bsp_pid());
120
121         bsp_sync();
122
123         assertint(v1.BSPunsafe_access(), bsp_pid(), bsp_pid(), __LINE__);
124
125         if (bsp_pid() == 0) {

```

```

126         endfun();
127         std::cout
128             << "Test 2 : bsp_put(const bsp_variable<T>& other)"
129             << std::endl;
130     }
131
132     bsp_variable<int> v2 = get_variable(-1);
133
134     v2.bsp_put(v1);
135
136     bsp_sync();
137
138     assertint(v2.BSPunsafe_access(), bsp_pid(), bsp_pid(), __LINE__);
139
140     if (bsp_pid() == 0) {
141         endfun();
142         std::cout
143             << "Test 3 : bsp_put(const bsp_variable<T>& other, int id)"
144             << std::endl;
145     }
146
147     bsp_variable<int> v3 = get_variable(-1);
148
149     v3.bsp_put(v2, mod((bsp_pid() + 1), bsp_nprocs()));
150     bsp_sync();
151
152     assertint(v3.BSPunsafe_access(), mod((bsp_pid() - 1), bsp_nprocs()),
153             bsp_pid(), __LINE__);
154
155     if (bsp_pid() == 0) {
156         endfun();
157         std::cout
158             << "Test 4: bsp_put(int destination)"
159             << std::endl;
160     }
161
162     bsp_variable<int> v4 = get_variable(bsp_pid());
163
164     int nextnode = mod(bsp_pid() + 1, bsp_nprocs());
165     v4.bsp_put(nextnode);
166     bsp_sync();
167
168     assertint(v4.BSPunsafe_access(), mod((bsp_pid() - 1), bsp_nprocs()),
169             bsp_pid(), __LINE__);
170
171     if (bsp_pid() == 0) {
172         endfun();
173         std::cout
174             << "Test 5: bsp_get(int destination)"
175             << std::endl;
176     }
177
178     bsp_variable<int> v5 = get_variable(bsp_pid());
179     bsp_sync(); // needed to make sure all nodes have time to create the variable
180
181     v5.bsp_get(nextnode);
182     bsp_sync();
183

```

```

184     assertint(v5.BSPunsafe_access(), nextnode, bsp_pid(), __LINE__);
185
186     if (bsp_pid() == 0) {
187         endfun();
188         std::cout
189             << "PART 2 : VARIABLES (PAR)"
190             << std::endl;
191         std::cout
192             << "Test 6: bsp_put(const T& elem, int id)"
193             << std::endl;
194     }
195
196     bsp_variable<int> v6 = get_variable(bsp_pid());
197
198     if (bsp_pid() != 0) v6.bsp_put(bsp_pid(), 0);
199     bsp_sync();
200
201     if (bsp_pid() == 0) {
202         assertint(v6.BSPunsafe_access(), 1, 2, bsp_pid(), __LINE__);
203         endfun();
204         std::cout
205             << "Test 7 : bsp_put(const bsp_variable<T>& other, int id)"
206             << std::endl;
207     }
208
209     bsp_variable<int> v7 = get_variable(-1);
210     if (bsp_pid() != 0) v7.bsp_put(v6, 0);
211     bsp_sync();
212
213     if (bsp_pid() == 0) {
214         assertint(v7.BSPunsafe_access(), 1, 2, bsp_pid(), __LINE__);
215         endfun();
216         std::cout
217             << "Test 8: bsp_put(int destination)"
218             << std::endl;
219     }
220
221     bsp_variable<int> v8 = get_variable(bsp_pid());
222     if (bsp_pid() != 0) v8.bsp_put(0);
223     bsp_sync();
224
225     if (bsp_pid() == 0) {
226         assertint(v8.BSPunsafe_access(), 1, 2, bsp_pid(), __LINE__);
227         endfun();
228         std::cout
229             << "PART 3: ARRAYS (SEQ)"
230             << std::endl;
231         std::cout
232             << "Test 9 : bsp_put(const T& elem, int pos)"
233             << std::endl;
234     }
235
236     std::vector<int> base_array{9, 8, 7, 6, 5, 4, 3, 2, 1, 0};
237     std::vector<int> compare_array1{bsp_pid(), 8, 7, 6, 5, 4, 3, 2, 1, 0};
238
239     bsp_array<int> a1 = get_array(base_array);
240     a1.bsp_put(bsp_pid(), 0);
241     bsp_sync();

```

```

242
243     assertvec(a1.BSPunsafe_access(), compare_array1, bsp_pid(), __LINE__);
244
245     if (bsp_pid() == 0) {
246         endfun();
247         std::cout
248             << "Test 10: bsp_put(const bsp_variable<T>& elem, int pos)"
249             << std::endl;
250     }
251
252     bsp_array<int> a2 = get_array(base_array);
253     bsp_variable<int> v9 = get_variable(bsp_pid());
254     a2.bsp_put(v9, 0);
255     bsp_sync();
256
257     assertvec(a2.BSPunsafe_access(), compare_array1, bsp_pid(), __LINE__);
258
259     if (bsp_pid() == 0) {
260         endfun();
261         std::cout
262             << "Test 11 : bsp_put(T elem, int dest, int pos)"
263             << std::endl;
264     }
265
266     int prevnode = mod(bsp_pid() - 1, bsp_nprocs());
267     std::vector<int> compare_array2{prevnode, 8, 7, 6, 5, 4, 3, 2, 1, 0};
268     bsp_array<int> a3 = get_array(base_array);
269
270     a3.bsp_put(bsp_pid(), nextnode, 0);
271     bsp_sync();
272
273     assertvec(a3.BSPunsafe_access(), compare_array2, bsp_pid(), __LINE__);
274     if (bsp_pid() == 0) {
275         endfun();
276         std::cout
277             << "Test 12 : bsp_put(bsp_variable<T>"
278             << "elem, int dest, int pos)"
279             << std::endl;
280     }
281
282     bsp_array<int> a4 = get_array(base_array);
283     bsp_variable<int> v10 = get_variable(bsp_pid());
284     a4.bsp_put(v10, nextnode, 0);
285     bsp_sync();
286
287     assertvec(a4.BSPunsafe_access(), compare_array2, bsp_pid(), __LINE__);
288
289     if (bsp_pid() == 0) {
290         endfun();
291         std::cout
292             << "Test 13 : bsp_put(std::vector<T> other)"
293             << std::endl;
294     }
295
296
297     std::vector<int> replace_array{bsp_pid(), 1, 2, 3, 4};
298     bsp_array<int> a5 = get_array(base_array);
299

```

```

300     a5.bsp_put(replace_array);
301     bsp_sync();
302
303     assertvec(a5.BSPunsafe_access(), replace_array, bsp_pid(), __LINE__);
304
305     if (bsp_pid() == 0) {
306         endfun();
307         std::cout
308             << "Test 14 : bsp_put(bsp_array<T> other)"
309             << std::endl;
310     }
311
312     bsp_array<int> a6 = get_array(base_array);
313     bsp_array<int> a7 = get_array(replace_array);
314
315     a6.bsp_put(a7);
316     bsp_sync();
317
318     assertvec(a6.BSPunsafe_access(), replace_array, bsp_pid(), __LINE__);
319
320     if (bsp_pid() == 0) {
321         endfun();
322         std::cout
323             << "Test 15 : bsp_put(std::vector<T> other, int dest)"
324             << std::endl;
325     }
326
327     auto control_array1 = std::vector<int>{prevnode, 1, 2, 3, 4};
328
329     bsp_array<int> a8 = get_array(base_array);
330
331     a8.bsp_put(replace_array, nextnode);
332     bsp_sync();
333
334     assertvec(a8.BSPunsafe_access(), control_array1, bsp_pid(), __LINE__);
335
336     if (bsp_pid() == 0) {
337         endfun();
338         std::cout
339             << "Test 16 : bsp_put(bsp_array<T> other, int dest)"
340             << std::endl;
341     }
342
343     bsp_array<int> a9 = get_array(base_array);
344     bsp_array<int> a10 = get_array(replace_array);
345
346     a9.bsp_put(a10, nextnode);
347     bsp_sync();
348
349     assertvec(a9.BSPunsafe_access(), control_array1, bsp_pid(), __LINE__);
350
351     if (bsp_pid() == 0) {
352         endfun();
353         std::cout
354             << "Test 17 : bsp_put(vector<T>, src_off, dst_off, len)"
355             << std::endl;
356     }
357

```

```

358     std::vector<int> substitution{0, 1, 2, 3, 4};
359     for (auto& el: substitution) el += bsp_pid();
360     std::vector<int> control_array2(base_array);
361     std::copy_n(substitution.begin() + 1, 3, control_array2.begin() +
362               2);
363     // {9, n+1, n+2, n+3, 5, 4, 3, 2, 1, 0}
364
365     bsp_array<int> a11 = get_array(base_array);
366     a11.bsp_put(substitution, 1, 2, 3);
367     bsp_sync();
368
369     assertvec(a11.BSPunsafe_access(), control_array2, bsp_pid(), __LINE__);
370
371     if (bsp_pid() == 0) {
372         endfun();
373         std::cout
374             << "Test 18: bsp_put(bsp_array<T>, src_off, dst_off, len)"
375             << std::endl;
376     }
377
378     bsp_array<int> a12 = get_array(base_array);
379     bsp_array<int> a13 = get_array(substitution);
380
381     a12.bsp_put(a13, 1, 2, 3);
382     bsp_sync();
383
384     assertvec(a12.BSPunsafe_access(), control_array2, bsp_pid(), __LINE__);
385
386     if (bsp_pid() == 0) {
387         endfun();
388         std::cout
389             << "Test 19 : bsp_put(vector<T>, dest, "
390             << "src_off, dst_off, len)"
391             << std::endl;
392     }
393
394     std::vector<int> substitution2{0, 1, 2, 3, 4};
395     for (auto& el: substitution2) el += prevnode;
396     std::vector<int> control_array3(base_array);
397     std::copy_n(substitution2.begin() + 1, 3, control_array3.begin() + 2);
398
399     bsp_array<int> a14 = get_array(base_array);
400     a14.bsp_put(substitution, nextnode, 1, 2, 3);
401     bsp_sync();
402
403     assertvec(a14.BSPunsafe_access(), control_array3, bsp_pid(), __LINE__);
404
405     if (bsp_pid() == 0) {
406         endfun();
407         std::cout
408             << "Test 20: bsp_put(bsp_array<T>, dest, "
409             << "src_off, dst_off, len)"
410             << std::endl;
411     }
412
413     bsp_array<int> a15 = get_array(base_array);
414     bsp_array<int> a16 = get_array(substitution);
415

```

```

416     a15.bsp_put(a16, nextnode, 1, 2, 3);
417     bsp_sync();
418
419     assertvec(a15.BSPunsafe_access(), control_array3, bsp_pid(), __LINE__);
420
421     if (bsp_pid() == 0) {
422         endfun();
423         std::cout
424             << "Test 21 : bsp_get(int destination)"
425             << std::endl;
426     }
427
428     std::vector<int> fillpids(5, bsp_pid());
429     bsp_array<int> a17 = get_array(fillpids);
430     bsp_sync();
431
432     a17.bsp_get(nextnode);
433     bsp_sync();
434
435     assertvec(a17.BSPunsafe_access(), std::vector<int>(5, nextnode),
436             bsp_pid(), __LINE__);
437
438     if (bsp_pid() == 0) {
439         endfun();
440         std::cout
441             << "Test 22 : bsp_get(source, src_offs, dest_offs, length)"
442             << std::endl;
443     }
444
445     std::vector<int> base_arr2{0, 1, 2, 3, 4};
446     for (auto& el: base_arr2) el += bsp_pid();
447     std::vector<int> control_array4{nextnode + 1, nextnode + 2,
448                                   nextnode + 3, bsp_pid() + 3,
449                                   bsp_pid() + 4};
450     bsp_array<int> a18 = get_array(base_arr2);
451     bsp_sync();
452
453     a18.bsp_get(nextnode, 1, 0, 3);
454     bsp_sync();
455
456     assertvec(a18.BSPunsafe_access(), control_array4, bsp_pid(), __LINE__);
457
458     if (bsp_pid() == 0) {
459         endfun();
460         std::cout
461             << "PART 4: ARRAYS (PAR)"
462             << std::endl;
463         std::cout
464             << "Test 23&24: bsp_put(T elem, int dest, int pos)"
465             << std::endl;
466         std::cout << "Different positions" << std::endl;
467     }
468
469     bsp_array<int> a19 = get_array(base_array);
470
471     a19.bsp_put(bsp_pid(), 0, bsp_pid());
472     bsp_sync();
473

```

```

474 auto compare_array5 = std::vector<int>{0, 1, 2, 6, 5, 4, 3, 2, 1, 0};
475 if (bsp_pid() == 0) {
476     assertvec(a19.BSPunsafe_access(), compare_array5, bsp_pid(),
477             __LINE__);
478     endfun();
479     std::cout << "Same position" << std::endl;
480 }
481
482 bsp_array<int> a20 = get_array(base_array);
483
484 if (bsp_pid() != 0) a20.bsp_put(bsp_pid(), 0, 0);
485 bsp_sync();
486
487 if (bsp_pid() == 0) {
488     assertint(a20.BSPunsafe_access().at(0), 1, 2, bsp_pid(), __LINE__);
489     endfun();
490     std::cout
491         << "Test 25&26 : bsp_put(bsp_variable<T>, dest, pos)"
492         << std::endl;
493     std::cout << "Different positions" << std::endl;
494 }
495
496 bsp_array<int> a21 = get_array(base_array);
497 bsp_variable<int> v11 = get_variable(bsp_pid());
498
499 a21.bsp_put(v11, 0, bsp_pid());
500 bsp_sync();
501
502 if (bsp_pid() == 0) {
503     assertvec(a21.BSPunsafe_access(), compare_array5, bsp_pid(),
504             __LINE__);
505     endfun();
506     std::cout << "Same position" << std::endl;
507 }
508
509 bsp_array<int> a22 = get_array(base_array);
510 bsp_variable<int> v12 = get_variable(bsp_pid());
511
512 if (bsp_pid() != 0) a22.bsp_put(v12, 0, 0);
513 bsp_sync();
514
515 if (bsp_pid() == 0) {
516     assertint(a22.BSPunsafe_access().at(0), 1, 2, bsp_pid(), __LINE__);
517     endfun();
518     std::cout
519         << "Test 27 : bsp_put(std::vector<T> other, int dest)"
520         << std::endl;
521 }
522
523 bsp_array<int> a23 = get_empty_array<int>(5);
524
525 if (bsp_pid() != 0) a23.bsp_put(std::vector<int>(5, bsp_pid()), 0);
526 bsp_sync();
527
528 if (bsp_pid() == 0) {
529     assertvec(a23.BSPunsafe_access(), std::vector<int>(5, 1),
530             std::vector<int>(5, 2), bsp_pid(), __LINE__);
531     endfun();

```



```

532         std::cout
533             << "Test 28 : bsp_put(bsp_array<T> other, int dest)"
534             << std::endl;
535     }
536
537     bsp_array<int> a24 = get_empty_array<int>(5);
538     bsp_array<int> a25 = get_array(std::vector<int>(5, bsp_pid()));
539
540     if (bsp_pid() != 0) a24.bsp_put(a25, 0);
541     bsp_sync();
542
543     if (bsp_pid() == 0) {
544         assertvec(a24.BSPunsafe_access(), std::vector<int>(5, 1),
545                 std::vector<int>(5, 2), bsp_pid(), __LINE__);
546         endfun();
547         std::cout
548             << "Test 29&30: bsp_put(vector<T>, dest, "
549             << "src_off, dst_off, len)"
550             << std::endl;
551         std::cout << "Separate positions" << std::endl;
552     }
553
554     const std::vector<int>& confront1{9, 1, 2, 3, 5, 4, 3, 98, 99, 0};
555     const std::vector<int>& subst1{0, 1, 2, 3, 4, 5};
556     const std::vector<int>& subst2{95, 96, 97, 98, 99};
557
558     bsp_array<int> a26 = get_array(base_array);
559
560     if (bsp_pid() == 1) a26.bsp_put(subst1, 0, 1, 1, 3);
561     if (bsp_pid() == 2) a26.bsp_put(subst2, 0, 3, 7, 2);
562     bsp_sync();
563
564     if (bsp_pid() == 0) {
565         assertvec(a26.BSPunsafe_access(), confront1, bsp_pid(), __LINE__);
566         endfun();
567         std::cout << "Overlapping positions" << std::endl;
568     }
569
570     const std::vector<int>& confront2{9, 8, 7, 1, 2, 3, 4, 98, 99, 0};
571     const std::vector<int>& confront3{9, 8, 7, 1, 2, 3, 97, 98, 99, 0};
572
573     bsp_array<int> a27 = get_array(base_array);
574
575     if (bsp_pid() == 1) a27.bsp_put(subst1, 0, 1, 3, 4);
576     if (bsp_pid() == 2) a27.bsp_put(subst2, 0, 2, 6, 3);
577     bsp_sync();
578
579     if (bsp_pid() == 0) {
580         assertvec(a27.BSPunsafe_access(), confront2, confront3, bsp_pid(),
581                 __LINE__);
582         endfun();
583         std::cout
584             << "Test 31&32 : bsp_put(bsp_array<T>, dest, "
585             << "src_off, dst_off, len)"
586             << std::endl;
587         std::cout << "Separate positions" << std::endl;
588     }
589

```

```

590     bsp_array<int> a28 = get_array(base_array);
591     bsp_array<int> a29 = get_array(bsp_pid() > 1 ? subst2 : subst1);
592
593     if (bsp_pid() == 1) a28.bsp_put(a29, 0, 1, 1, 3);
594     if (bsp_pid() == 2) a28.bsp_put(a29, 0, 3, 7, 2);
595     bsp_sync();
596
597     if (bsp_pid() == 0) {
598         assertvec(a28.BSPunsafe_access(), confront1, bsp_pid(), __LINE__);
599         endfun();
600         std::cout << "Overlapping positions" << std::endl;
601     }
602
603     bsp_array<int> a30 = get_array(base_array);
604     bsp_array<int> a31 = get_array(bsp_pid() > 1 ? subst2 : subst1);
605
606     if (bsp_pid() == 1) a30.bsp_put(a31, 0, 1, 3, 4);
607     if (bsp_pid() == 2) a30.bsp_put(a31, 0, 2, 6, 3);
608     bsp_sync();
609
610     if (bsp_pid() == 0) {
611         assertvec(a30.BSPunsafe_access(), confront2, confront3, bsp_pid(),
612                 __LINE__);
613         endfun();
614         std::cout
615             << "PART 5 : DIRECT GETS"
616             << std::endl;
617         std::cout
618             << "Test 33 (var) : T bsp_direct_get(int source)"
619             << std::endl;
620     }
621
622     bsp_variable<int> v13 = get_variable(bsp_pid());
623     bsp_sync();
624
625     int dg1 = v13.bsp_direct_get(0);
626     v13.bsp_put(bsp_pid(), 0);
627     bsp_sync();
628
629     assertint(dg1, 0, bsp_pid(), __LINE__);
630
631     if (bsp_pid() == 0) {
632         endfun();
633         std::cout
634             << "Test 34 (arr): T bsp_direct_get(int source, int pos)"
635             << std::endl;
636     }
637
638     bsp_array<int> a32 = get_array(base_array);
639     bsp_sync();
640
641     int dg2 = a32.bsp_direct_get(0, 3);
642     a32.bsp_put(bsp_pid(), 0, 3);
643     bsp_sync();
644
645     assertint(dg2, 6, bsp_pid(), __LINE__);
646
647     if (bsp_pid() == 0) {

```

```

648         endfun();
649         std::cout
650             << "Test 35 (arr) : vector<T> bsp_direct_get(int source)"
651             << std::endl;
652     }
653
654     bsp_array<int> a33 = get_array(fillpids);
655     bsp_sync();
656
657     auto dg3 = a33.bsp_direct_get(0);
658     a33.bsp_put(std::vector<int>(5, bsp_pid()));
659     bsp_sync();
660
661     assertvec(dg3, std::vector<int>(5, 0), bsp_pid(), __LINE__);
662     if (bsp_pid() == 0) {
663         endfun();
664         std::cout
665             << "TEST ENDED"
666             << std::endl;
667         endfun(true);
668     }
669 }
670
671 };
672
673 int main() {
674     std::vector<std::unique_ptr<bsp_node>> nodes;
675     for (size_t i{0}; i < 3; ++i) {
676         nodes.push_back(std::make_unique<unit_tests>());
677     }
678     bsp_program mbsp(std::move(nodes));
679     mbsp.start();
680 }

```

## B.5 Java programs

### BSPpsrs.java

```

1  import com.multicorebsp.core.*;
2  import org.apache.commons.math3.random.MersenneTwister;
3  import org.apache.commons.math3.util.MathArrays;
4
5  import java.util.*;
6  import java.util.stream.IntStream;
7
8  public class BSPpsrs extends BSP_PROGRAM {
9
10     private int n;
11     private int n_procs;
12     private int seed;
13
14     @Override
15     protected void main_part() throws InterruptedException {
16         try {
17             bsp_begin(n_procs);
18         } catch (IllegalAccessException | InstantiationException |
19             EmptyQueueException | InterruptedException e) {

```

```

20         e.printStackTrace();
21     }
22 }
23
24 @Override
25 protected void parallel_part() throws InterruptedException,
26     IllegalAccessException, EmptyQueueException {
27     BSP_INTEGER problem_size = new BSP_INTEGER(this, 0);
28     if (bsp_pid() == 0) {
29         for (int i = 0; i < bsp_nprocs(); ++i) {
30             problem_size.bsp_put(n, i);
31         }
32     }
33     bsp_sync();
34     n = problem_size.read();
35
36     int p = bsp_nprocs();
37     int s = bsp_pid();
38     long t = System.currentTimeMillis();
39     int numels = n / p;
40     BSP_INT_ARRAY to_sort = new BSP_INT_ARRAY(this, numels);
41
42     if (s == 0) {
43         MersenneTwister mtw;
44         if (seed == -1) {
45             mtw = new MersenneTwister();
46         } else {
47             mtw = new MersenneTwister(seed);
48         }
49         int[] data = IntStream.range(0, n).toArray();
50         MathArrays.shuffle(data, mtw);
51         long t2 = System.currentTimeMillis();
52         System.out.println("Ended vector generation and shuffling (spent "
53             + (t2 - t) + " ms)");
54         System.out.println("Starting parallel part");
55
56         t = System.currentTimeMillis();
57         int count = 0;
58         for (int i = 0; i < n; i += numels) {
59             int last = Math.min(n, i + numels);
60             to_sort.bsp_put(data, i, count++, 0, last - i);
61         }
62     }
63     bsp_sync();
64
65     int[] vec = to_sort.getData();
66
67     Arrays.sort(vec);
68
69     ArrayList<Integer> primary_samples = new ArrayList<>();
70     int samplesize = vec.length / p;
71     int i;
72     for (i = 0; i < vec.length; i += samplesize) {
73         primary_samples.add(vec[i]);
74     }
75
76     if (i != vec.length - 1) {
77         primary_samples.add(vec[i - 1]);

```

```

78     }
79
80     IntArrayList psam = new IntArrayList();
81     psam.array.addAll(primary_samples);
82
83     BSP_ARRAY<IntArrayList> ps_array = new BSP_ARRAY<>(this, p, psam);
84     for (i = 0; i < p; ++i) {
85         ps_array.bsp_put(psam, i, s);
86     }
87     bsp_sync();
88
89     ArrayList<Integer> ps_all = new ArrayList<>();
90     ArrayList<Integer> secondary_samples = new ArrayList<>();
91     try {
92         for (i = 0; i < p; ++i) {
93             ps_all.addAll(ps_array.read(i).array);
94         }
95         Collections.sort(ps_all);
96         samplesize = ps_all.size() / p;
97     } catch (NullPointerException e) {
98         e.printStackTrace();
99     }
100
101     for (i = 0; i < ps_all.size(); i += samplesize) {
102         secondary_samples.add(ps_all.get(i));
103     }
104
105     if (i == ps_all.size())
106         secondary_samples.add(ps_all.get(i - 1));
107
108     int upperbound;
109     int count = 1;
110
111     do {
112         upperbound = secondary_samples.get(count++);
113     } while (vec[0] > upperbound);
114
115     count--;
116     ArrayList<Integer> temp = new ArrayList<>();
117
118     BSP_ARRAY<IntArrayList> portion =
119         new BSP_ARRAY<>(this, p, new IntArrayList());
120
121     for (i = 0; i < vec.length; ++i) {
122         if (vec[i] > upperbound) {
123             portion.bsp_put(new IntArrayList(temp), count - 1, s);
124             temp.clear();
125             upperbound = secondary_samples.get(++count);
126         }
127         temp.add(vec[i]);
128     }
129
130     portion.bsp_put(new IntArrayList(temp), count - 1, s);
131     bsp_sync();
132
133     ArrayList<Integer> secondary_block = new ArrayList<>();
134     BSP_ARRAY<IntArrayList> final_arr = new BSP_ARRAY<>(this, p,
135         new IntArrayList());

```

```

136     for (IntArrayList pbl : portion) {
137         secondary_block.addAll(pbl.array);
138     }
139
140     Collections.sort(secondary_block);
141     final_arr.bsp_put(new IntArrayList(secondary_block), 0, s);
142     bsp_sync();
143
144     if (s == 0) {
145         ArrayList<Integer> conclusion = new ArrayList<>();
146         for (IntArrayList el : final_arr)
147             conclusion.addAll(el.array);
148
149         boolean passed = true;
150         for (int j = 0; j < n; ++j) {
151             if (conclusion.get(j) != j) {
152                 passed = false;
153                 break;
154             }
155         }
156
157         System.out.println("Check " + (passed ? "passed" : "failed"));
158         long t1 = System.currentTimeMillis();
159         System.out.println("Parallel part took " + (t1 - t) + " ms.");
160     }
161
162     bsp_sync();
163 }
164
165 public static void main(String[] args) {
166     if (args.length < 2) {
167         System.err.println("Usage: BSPpsrs <N> <P> (seed)");
168         System.err.println(" where N is the problem size (power of 2)");
169         System.err.println("       P is the number of threads (power of 2)");
170         System.err.println("       N must be >= P^3");
171         System.err.println("       seed is an optional seed " +
172             "for permutations (leave blank to randomize it)");
173         System.exit(1);
174     }
175     int n, p, s;
176     n = Integer.parseInt(args[0]);
177     p = Integer.parseInt(args[1]);
178     if (n < p * p * p) {
179         System.err.println("N must be >= P^3");
180         System.exit(1);
181     }
182     s = (args.length >= 3 ? Integer.parseInt(args[2]) : -1);
183     BSPpsrs my_computation = new BSPpsrs();
184     my_computation.n = n;
185     my_computation.n_procs = p;
186     my_computation.seed = s;
187     my_computation.start();
188 }
189 }

```

## IntArrayList.java

```
1 import com.multicorebsp.core.CompulsaryCloneable;
2
3 import java.util.ArrayList;
4
5 public class IntArrayList implements CompulsaryCloneable<IntArrayList> {
6
7     public ArrayList<Integer> array;
8
9     public IntArrayList() {
10         array = new ArrayList<>();
11     }
12     public IntArrayList(ArrayList<Integer> other) {
13         array = new ArrayList<>(other);
14     }
15
16     @Override
17     public Object clone() throws CloneNotSupportedException {
18         super.clone();
19         return safeClone();
20     }
21
22     @Override
23     public IntArrayList safeClone() {
24         return new IntArrayList(array);
25     }
26 }
```

## SequentialInprod.java

```
1 public class SequentialInprod {
2
3     public static void main(String[] args) {
4         if (args.length < 1) System.exit(1);
5
6         int n = 1;
7
8         try {
9             n = Integer.parseInt(args[0]);
10        } catch (NumberFormatException e) {
11            System.err.println("Argument must be an integer");
12            System.exit(-1);
13        }
14
15        double[] x = new double[n];
16
17        for (int i = 0; i < n; ++i) {
18            x[i] = i + 1;
19        }
20
21        double inprod = 0.0;
22        long t1 = System.currentTimeMillis();
23
24        for (int i = 0; i < n; ++i) {
25            inprod += x[i] * x[i];
26        }
27    }
28 }
```

```

27     long t2 = System.currentTimeMillis();
28
29     System.out.println("Processor 0: sum of squares up to " + n + "*"
30         + n + " is " + inprod);
31     System.out.println("Processor 0: local time taken is " + (t2 - t1));
32
33     double dn = (double) n;
34     dn *= n + 1.0;
35     dn *= 2.0 * n + 1.0;
36     dn /= 6.0;
37     System.out.println("Checksum: " + dn);
38
39     System.exit(0);
40 }
41
42 }

```

## SequentialSort.java

```

1  import java.util.*;
2  import java.util.stream.Collectors;
3
4  public class SequentialSort {
5
6      private static void arrshuffle(int[] array, Random rd) {
7          for (int i = array.length - 1; i > 0; i--) {
8              int index = rd.nextInt(i + 1);
9              // Simple swap
10             int a = array[index];
11             array[index] = array[i];
12             array[i] = a;
13         }
14     }
15
16     public static void main(String[] args) {
17         if (args.length < 1) {
18             System.err.println("Usage: seqsort <N> (seed)");
19             System.err.println(" where N is the problem size (power of 2)");
20             System.err.println("      seed is an optional seed for " +
21                 "permutations (leave blank to randomize it)");
22             System.exit(1);
23         }
24         int n = 1, s = -1;
25         try {
26             n = Integer.parseInt(args[0]);
27         } catch (NumberFormatException e) {
28             System.err.println("First argument must be an integer");
29             System.exit(-1);
30         }
31         if (args.length >= 2) {
32             try {
33                 s = Integer.parseInt(args[1]);
34             } catch (NumberFormatException e) {
35                 System.err.println("Second argument must be an integer");
36                 System.exit(-1);
37             }
38         }

```



```

39     long t = System.currentTimeMillis();
40     Random rd = new Random();
41     if (s == -1) {
42         s = rd.nextInt();
43     }
44     rd.setSeed(s);
45     int[] data = new int[n];
46     for (int i = 0; i < n; ++i) {
47         data[i] = i;
48     }
49
50     arrshuffle(data, rd);
51
52     List<Integer> list = Arrays.stream(data).boxed()
53         .collect(Collectors.toList());
54
55     long t1 = System.currentTimeMillis();
56     System.out.println("Ended vector generation and shuffling (spent " +
57         (t1 - t) + " ms)");
58     System.out.println("starting sequential part");
59
60     long t2 = System.currentTimeMillis();
61
62     Arrays.sort(data);
63
64     long t3 = System.currentTimeMillis();
65
66     Collections.sort(list);
67
68     long t4 = System.currentTimeMillis();
69
70     boolean passed = true;
71     for (int j = 0; j < n; j++) {
72         if (data[j] != j) {
73             passed = false;
74             break;
75         }
76     }
77
78     System.out.println("Check " + (passed ? "passed" : "failed"));
79     System.out.println("Parallel part took " + (t3 - t2) + " ms.");
80     System.out.println("Collections part took " + (t4 - t3) + " ms.");
81 }
82 }

```

## SequentialFFT.java

```

1 import java.util.Arrays;
2
3 public class SequentialFFT {
4
5     private final static double PI = 3.141592653589793;
6
7     private static void ufft(double[] x, int offset, int n, boolean sign,
8         double[] w) {
9         for (int k = 2; k <= n; k *= 2) {
10             int nk = n / k;

```

```

11     for (int r = 0; r < nk; ++r) {
12         int rk = 2 * r * k;
13         for (int j = 0; j < k; j += 2) {
14             double wr = w[j * nk];
15             double wi;
16             if (sign) {
17                 wi = w[j * nk + 1];
18             } else {
19                 wi = -w[j * nk + 1];
20             }
21
22             int j0 = rk + j + offset;
23             int j1 = j0 + 1;
24             int j2 = j0 + k;
25             int j3 = j2 + 1;
26
27             double taur = wr * x[j2] - wi * x[j3];
28             double taui = wi * x[j2] + wr * x[j3];
29
30             x[j2] = x[j0] - taur;
31             x[j3] = x[j1] - taui;
32             x[j0] += taur;
33             x[j1] += taui;
34         }
35     }
36 }
37
38
39 private static void ufft_init(int n, double[] w) {
40     assert (w.length == n);
41
42     if (n == 1) return;
43
44     w[0] = 1.0;
45     w[1] = 0.0;
46
47     if (n == 4) {
48         w[2] = 0.0;
49         w[3] = -1.0;
50     } else if (n >= 8) {
51         double theta = -2.0 * PI / (double) (n);
52         for (int j = 1; j <= n / 8; j++) {
53             w[2 * j] = Math.cos(j * theta);
54             w[2 * j + 1] = Math.sin(j * theta);
55         }
56         for (int j = 0; j < n / 8; j++) {
57             int n4j = n / 4 - j;
58             w[2 * n4j] = -w[2 * j + 1];
59             w[2 * n4j + 1] = -w[2 * j];
60         }
61         for (int j = 1; j < n / 4; j++) {
62             int n2j = n / 2 - j;
63             w[2 * n2j] = -w[2 * j];
64             w[2 * n2j + 1] = w[2 * j + 1];
65         }
66     }
67 }
68

```

```

69 private static void twiddle_init(int n, double alpha, int[] rho,
70                               double[] w, int offset) {
71     double theta = -2.0 * PI * alpha / (double) (n);
72     for (int j = 0; j < n; ++j) {
73         double rt = (double) (rho[j]) * theta;
74         w[offset + 2 * j] = Math.cos(rt);
75         w[offset + 2 * j + 1] = Math.sin(rt);
76     }
77 }
78
79 private static void permute(double[] x, int n, int[] sigma) {
80     assert (x.length / 2 == sigma.length);
81
82     for (int j = 0; j < n; ++j) {
83         if (j < sigma[j]) {
84             int j0 = 2 * j;
85             int j1 = j0 + 1;
86             int j2 = 2 * sigma[j];
87             int j3 = j2 + 1;
88             double tmpr = x[j0];
89             double tmpi = x[j1];
90             x[j0] = x[j2];
91             x[j1] = x[j3];
92             x[j2] = tmpr;
93             x[j3] = tmpi;
94         }
95     }
96 }
97
98 private static void bitrev_init(int[] rho) {
99     int n = rho.length;
100
101     int binary_len = (int) (Math.ceil(Math.log((double)(n))/Math.log(2.0)));
102     boolean[] bits = new boolean[binary_len];
103     int[] pwrs = new int[binary_len];
104     pwrs[0] = 1;
105     for (int j = 1; j < binary_len; ++j) {
106         pwrs[j] = pwrs[j - 1] * 2;
107     }
108     int j = 0;
109     while (j < n - 1) {
110         j++;
111         int lastbit = 0;
112         while (bits[lastbit]) {
113             bits[lastbit] = false;
114             lastbit++;
115         }
116         bits[lastbit] = true;
117         int val = 0;
118         for (int k = 0; k < binary_len; ++k) {
119             if (bits[k]) {
120                 val += pwrs[binary_len - k - 1];
121             }
122         }
123         rho[j] = val;
124     }
125 }
126

```

```

127 private static void fft_init(int n, double[] w,
128                             double[] tw, int[] rho_np) {
129     bitrev_init(rho_np);
130     ufft_init(n, w);
131
132     twiddle_init(n, 0, rho_np, tw, 0);
133 }
134
135 private static void calcError(double[] xlocal, double[] xarr) {
136     double error = 0.0;
137     for (int c = 0; c < xlocal.length; c++) {
138         double lerror = Math.abs(xlocal[c] - xarr[c]);
139         error += lerror;
140     }
141     System.out.println("local error is " + (error/(double)xlocal.length));
142 }
143
144 public static void main(String[] args) {
145     if (args.length < 1) {
146         System.err.println("Usage: seqfft <N>");
147         System.err.println("where N is the problem size (power of 2)");
148         System.exit(-1);
149     }
150     int n = 1;
151     try {
152         n = Integer.parseInt(args[0]);
153         if (!(n > 0 && (n & n - 1) == 0))
154             throw new NumberFormatException("not power of 2");
155     } catch (NumberFormatException e) {
156         System.err.println("Argument must be an integer (power of 2)");
157         System.exit(-1);
158     }
159     double[] data = new double[n];
160
161     n /= 2;
162     for (int i = 0; i < n; i += 2) {
163         data[i] = (double) (i) / 2.0;
164     }
165
166     double[] old = Arrays.copyOf(data, data.length);
167
168     double[] w = new double[n];
169     double[] tw = new double[(2 * n)];
170     int[] rho_np = new int[(n)];
171
172     long t = System.currentTimeMillis();
173     fft_init(n, w, tw, rho_np);
174     // forward fft
175     permute(data, n, rho_np);
176     ufft(data, 0, n, true, w);
177
178     // inverse fft
179     permute(data, n, rho_np);
180     ufft(data, 0, n, false, w);
181
182     double ninv = 1.0 / (double) (n);
183     for (int j = 0; j < 2 * n; ++j) {
184         data[j] *= ninv;

```

```

185     }
186
187     calcError(data, old);
188     long t1 = System.currentTimeMillis();
189
190
191     System.out.println("Sequential part took " + (t1 - t) + " ms.");
192
193     calcError(data, old);
194
195
196     System.exit(0);
197 }
198 }

```

### SequentialLU\_commons.java

```

1  import org.apache.commons.math3.linear.LUdecomposition;
2  import org.apache.commons.math3.linear.MatrixUtils;
3  import org.apache.commons.math3.linear.RealMatrix;
4
5  import java.util.Random;
6
7  public class SequentialLU_commons {
8
9      public static void main(String[] args) {
10         if (args.length < 1) {
11             System.err.println("Usage: seqlu <n>");
12             System.err.println("\twill stat LU decomposition on a " +
13                 "n by n matrix.");
14             System.exit(1);
15         }
16
17         int n = 1;
18         try {
19             n = Integer.parseInt(args[0]);
20         } catch (NumberFormatException e) {
21             System.err.println("Argument must be an integer");
22             System.exit(-1);
23         }
24
25         Random rd = new Random();
26
27         double[][] a = new double[n][n];
28         double[][] l = new double[n][n];
29         double[][] u = new double[n][n];
30
31         for (int i = 0; i < n; ++i) {
32             for (int j = 0; j < n; ++j) {
33                 double val = Math.random();
34                 a[i][j] = val;
35                 l[i][j] = u[i][j] = 0.0;
36             }
37         }
38
39         RealMatrix A = MatrixUtils.createRealMatrix(a);
40

```

```
41     long t1 = System.currentTimeMillis();
42     LUdecomposition luCalc = new LUdecomposition(A);
43     l = luCalc.getL().getData();
44     u = luCalc.getU().getData();
45     long t2 = System.currentTimeMillis();
46
47     System.out.println("Sequential computation took " + (t2 - t1) + " ms.");
48     System.exit(0);
49 }
50 }
```

# Bibliography

- [1] S.G. Akl and W. Rheinboldt. *Parallel Sorting Algorithms*. Notes and reports in computer science and applied mathematics. Elsevier Science, 2014. ISBN: 9781483268088. URL: <https://books.google.it/books?id=jhHjBQAAQBAJ>.
- [2] Marco Aldinucci, Salvatore Ruggieri, and Massimo Torquati. “Porting decision tree algorithms to multicore using FastFlow”. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2010, pp. 7–23.
- [3] Marco Aldinucci et al. “Fastflow: high-level and efficient streaming on multi-core”. In: *Programming multi-core and many-core computing systems, parallel and distributed computing* (2017).
- [4] Marco Aldinucci et al. “Targeting distributed systems in fastflow”. In: *European Conference on Parallel Processing*. Springer. 2012, pp. 47–56.
- [5] B. Andrist and V. Sehr. *C++ High Performance: Boost and optimize the performance of your C++17 code*. Packt Publishing, 2018. ISBN: 9781787124776. URL: <https://books.google.it/books?id=EyBKDwAAQBAJ>.
- [6] *Apache Giraph*. URL: <https://giraph.apache.org/> (visited on 04/20/2020).
- [7] *Apache Hama*. URL: <https://hama.apache.org/> (visited on 04/20/2020).
- [8] R.H. Bisseling. *Author’s page on the University of Utrecht website*. URL: <http://www.staff.science.uu.nl/~bisse101/Software/software.html> (visited on 04/21/2020).
- [9] R.H. Bisseling. *Parallel Scientific Computation: A Structured Approach Using BSP and MPI*. Oxford University Press, 2004. ISBN: 9780198529392. URL: <https://books.google.it/books?id=iCsTDAAAQBAJ>.
- [10] Dina Bitton et al. “A taxonomy of parallel sorting”. In: *ACM Computing Surveys (CSUR)* 16.3 (1984), pp. 287–318.
- [11] Å. Björck. *Numerical Methods in Matrix Computations*. Texts in Applied Mathematics. Springer International Publishing, 2014. ISBN: 9783319050898. URL: <https://books.google.it/books?id=jo08BAAAQBAJ>.
- [12] Olaf Bonorden et al. “The paderborn university BSP (PUB) library”. In: *Parallel Computing* 29.2 (2003), pp. 187–207.
- [13] Jeff Bonwick. “The Slab Allocator: An Object-Caching Kernel Memory Allocator”. In: *Proceedings of the USENIX Summer 1994 Technical Conference on USENIX Summer 1994 Technical Conference - Volume 1*. USTC’94. Boston, Massachusetts: USENIX Association, 1994, p. 6.

- [14] *BSPLib Home Page*. URL: <http://www.bsp-worldwide.org/> (visited on 04/23/2020).
- [15] Rainer Buchty et al. “A survey on hardware-aware and heterogeneous computing on multicore processors and accelerators”. In: *Concurrency and Computation: Practice and Experience* 24.7 (2012), pp. 663–675. DOI: 10.1002/cpe.1904. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.1904>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.1904>.
- [16] Jan-Willem Burchage, Tom Bannink, and Rob H. Bisseling. “Bulk: A Modern C++ Interface for Bulk-Synchronous Parallel Programs”. In: *EuroPar 2018: Parallel Processing*. Ed. by Marco Aldinucci, Luca Padovani, and Massimo Torquati. Cham: Springer International Publishing, 2018, pp. 519–532. ISBN: 978-3-319-96983-1.
- [17] Avery Ching et al. “One trillion edges: Graph processing at facebook-scale”. In: *Proceedings of the VLDB Endowment* 8.12 (2015), pp. 1804–1815.
- [18] James W Cooley and John W Tukey. “An algorithm for the machine calculation of complex Fourier series”. In: *Mathematics of computation* 19.90 (1965), pp. 297–301.
- [19] *FastFlow commit 8b9105d*. URL: <https://github.com/fastflow/fastflow/commit/8b9105de8784335fc2fcf28c3ba955cd68dfed5b> (visited on 04/22/2020).
- [20] *FastFlow official repository on GitHub*. URL: <https://github.com/fastflow/fastflow> (visited on 04/22/2020).
- [21] *FastFlow online tutorial*. URL: <http://calvados.di.unipi.it/dokuwiki/doku.php/ffnamespace:tutorial> (visited on 04/23/2020).
- [22] Alexandros V Gerbessiotis. “Extending the BSP model for multi-core and out-of-core computing: MBSP”. In: *Parallel Computing* 41 (2015), pp. 90–102.
- [23] Alexandros V. Gerbessiotis and Constantinos J. Siniolakis. *Efficient Deterministic Sorting on the BSP Model*. Tech. rep. PARALLEL PROCESSING LETTERS, 1996.
- [24] Michael T Goodrich. “Communication-efficient parallel sorting”. In: *SIAM Journal on Computing* 29.2 (1999), pp. 416–432.
- [25] Khaled Hamidouche, Joel Falcou, and Daniel Etiemble. “Hybrid Bulk Synchronous Parallelism Library for Clustered Smp Architectures”. In: *Proceedings of the Fourth International Workshop on High-Level Parallel Programming and Applications*. HLPP '10. Baltimore, Maryland, USA: Association for Computing Machinery, 2010, pp. 55–62. ISBN: 9781450302548. DOI: 10.1145/1863482.1863494. URL: <https://doi.org/10.1145/1863482.1863494>.
- [26] Thomas Heller et al. “Hpx—an open source c++ standard library for parallelism and concurrency”. In: *Proceedings of OpenSuCo* (2017), p. 5.
- [27] Jonathan MD Hill et al. “BSPLib: The BSP programming library”. In: *Parallel Computing* 24.14 (1998), pp. 1947–1980.
- [28] *Intel Xeon Phi x200 Family specifications*. URL: <https://ark.intel.com/content/www/us/en/ark/products/series/92650/intel-xeon-phi-x200-product-family.html> (visited on 04/22/2020).



- [29] Peter Kankowski. *Hash Functions: An Empirical Comparison*. 2009. URL: <https://www.codeproject.com/Articles/32829/Hash-Functions-An-Empirical-Comparison>.
- [30] Richard M. Karp. *A Survey of Parallel Algorithms for Shared-Memory Machines*. Tech. rep. USA, 1988.
- [31] David Lecomber. *An Object-Oriented Programming Model for BSP Computations*. Tech. rep. In Proc. PPECC Workshop on Parallel and Distributed Computing, 1994.
- [32] Xiaobo Li et al. "On the versatility of parallel sorting by regular sampling". In: *Parallel Computing* 19.10 (1993), pp. 1079–1103.
- [33] Frédéric Loulergue, Frédéric Gava, and David Billiet. "Bulk synchronous parallel ML: modular implementation and performance prediction". In: *International Conference on Computational Science*. Springer. 2005, pp. 1046–1054.
- [34] *LU decomposition documentation for the ALGLIB library*. URL: <https://www.alglib.net/matrixops/lu.php> (visited on 04/22/2020).
- [35] *LU decomposition documentation for the Apache Commons Java library*. URL: <http://commons.apache.org/proper/commons-math/javadocs/api-3.6/org/apache/commons/math3/linear/LUdecomposition.html> (visited on 04/22/2020).
- [36] Grzegorz Malewicz et al. "Pregel: a system for large-scale graph processing". In: *Proceedings of the 2010 international conference on Management of data*. New York, NY, USA, 2010, pp. 135–146. URL: <http://doi.acm.org/10.1145/1807167.1807184>.
- [37] Michael McCool, James Reinders, and Arch Robison. *Structured parallel programming: patterns for efficient computation*. Elsevier, 2012.
- [38] *MulticoreBSP for C*. URL: <http://www.multicorebsp.com/download/c/> (visited on 04/20/2020).
- [39] *MulticoreBSP for Java*. URL: <http://www.multicorebsp.com/download/java/> (visited on 04/20/2020).
- [40] *MulticoreBSP for Java documentation*. URL: <http://www.multicorebsp.com/javadoc/> (visited on 04/16/2020).
- [41] *Repository for the library implemented as part of this thesis*. URL: <https://gitlab.com/leombro/master-thesis> (visited on 04/16/2020).
- [42] Luis FG Sarmenta. "An adaptive, fault-tolerant implementation of BSP for Java-based volunteer computing systems". In: *International Parallel Processing Symposium*. Springer. 1999, pp. 763–780.
- [43] Hideyuki Shamoto et al. "Large-scale distributed sorting for GPU-based heterogeneous supercomputers". In: *2014 IEEE International Conference on Big Data (Big Data)*. IEEE. 2014, pp. 510–518.
- [44] R. Shams et al. "A Survey of Medical Image Registration on Multicore and the GPU". In: *IEEE Signal Processing Magazine* 27.2 (2010), pp. 50–60.
- [45] R. Shaposhnik, C. Martella, and D. Logothetis. *Practical Graph Analytics with Apache Giraph*. Apress, 2015. ISBN: 9781484212516. URL: [https://books.google.it/books?id=Fwb%5C\\_CgAAQBAJ](https://books.google.it/books?id=Fwb%5C_CgAAQBAJ).

- [46] Hanmao Shi and Jonathan Schaeffer. “Parallel sorting by regular sampling”. In: *Journal of parallel and distributed computing* 14.4 (1992), pp. 361–372.
- [47] Kamran Siddique, Zahid Akhtar, and Yangwoo Kim. “Researching Apache Hama: A Pure BSP Computing Framework”. In: *Advanced Multimedia and Ubiquitous Engineering*. Ed. by James J. (Jong Hyuk) Park et al. Springer Singapore, 2016, pp. 215–221. ISBN: 978-981-10-1536-6.
- [48] Dharendra Pratap Singh, Ishan Joshi, and Jaytrilok Choudhary. “Survey of GPU Based Sorting Algorithms”. In: *Int. J. Parallel Program.* 46.6 (Dec. 2018), pp. 1017–1034. ISSN: 0885-7458. DOI: 10.1007/s10766-017-0502-5. URL: <https://doi.org/10.1007/s10766-017-0502-5>.
- [49] *std::allocate\_shared documentation*. URL: [https://en.cppreference.com/w/cpp/memory/shared\\_ptr/allocate\\_shared](https://en.cppreference.com/w/cpp/memory/shared_ptr/allocate_shared) (visited on 04/16/2020).
- [50] Maarten van Steen and Andrew S. Tanenbaum. *Distributed Systems*. 3rd ed. 2017. URL: <https://www.distributed-systems.net/>.
- [51] Bjarne Stroustrup. “What is object-oriented programming?” In: *IEEE software* 5.3 (1988), pp. 10–20.
- [52] *Ten and Now: Almost 10 Years of Intel CPUs compared*. URL: <https://www.techspot.com/article/1039-ten-years-intel-cpu-compared/> (visited on 04/22/2020).
- [53] *The djb2 hash function*. 1991. URL: <https://groups.google.com/forum/#!msg/comp.lang.c/LSKWxiuNOAk/zstZ3SRhCjgJ> (visited on 04/15/2020).
- [54] *The Doolittle Algorithm*. URL: [https://vismor.com/documents/network\\_analysis/matrix\\_algorithms/S4.SS2.php](https://vismor.com/documents/network_analysis/matrix_algorithms/S4.SS2.php) (visited on 04/22/2020).
- [55] Alexandre Tiskin. “The design and analysis of bulk-synchronous parallel algorithms”. PhD thesis. 1998.
- [56] Massimo Torquati. *Introduction to FastFlow programming: allocators*. URL: <http://didawiki.di.unipi.it/lib/exe/fetch.php/magistraleinformaticanetworking/spm/spm1617dic15-2.pdf> (visited on 04/16/2020).
- [57] Massimo Torquati. “Single-Producer/Single-Consumer Queues on Shared Cache Multi-Core Systems”. In: *CoRR* abs/1012.1824 (2010). arXiv: 1012.1824. URL: <http://arxiv.org/abs/1012.1824>.
- [58] Matthew Travers. “Cpu power consumption experiments and results analysis of intel i7-4820k”. In: (2015).
- [59] L. G. Valiant. “General Purpose Parallel Architectures”. In: *Handbook of Theoretical Computer Science (Vol. A): Algorithms and Complexity*. Cambridge, MA, USA: MIT Press, 1991, pp. 943–973. ISBN: 0444880712.
- [60] Leslie G Valiant. “A bridging model for multi-core computing”. In: *Journal of Computer and System Sciences* 77.1 (2011), pp. 154–166.
- [61] Leslie G. Valiant. “A Bridging Model for Parallel Computation”. In: *Commun. ACM* 33.8 (Aug. 1990), pp. 103–111. ISSN: 0001-0782. DOI: 10.1145/79173.79181. URL: <https://doi.org/10.1145/79173.79181>.
- [62] *Wikipedia article on Daniel J. Bernstein*. URL: [https://en.wikipedia.org/wiki/Daniel\\_J.\\_Bernstein](https://en.wikipedia.org/wiki/Daniel_J._Bernstein) (visited on 04/15/2020).

- [63] A. Williams. *C++ Concurrency in Action: Practical Multithreading*. Manning Pubs Co Series. Manning, 2012. ISBN: 9781933988771. URL: <https://books.google.it/books?id=EttPPgAACAAJ>.
- [64] J. C. Wyllie. "The Complexity of Parallel Computations". PhD thesis. Computer Science Department, Cornell University, Ithaca, NY, 1981.
- [65] Yongsuk Kee and Soonhoi Ha. "xBSP: an efficient BSP implementation for cLAN". In: *Proceedings First IEEE/ACM International Symposium on Cluster Computing and the Grid*. 2001, pp. 237–244.
- [66] A.N. Yzelman and Rob H. Bisseling. "An Object-Oriented Bulk Synchronous Parallel Library for Multicore Programming". In: *Concurr. Comput. : Pract. Exper.* 24.5 (Apr. 2012), pp. 533–553. ISSN: 1532-0626. DOI: 10.1002/cpe.1843. URL: <https://doi.org/10.1002/cpe.1843>.
- [67] Albert-Jan Yzelman et al. "MulticoreBSP for C: A High-Performance Library for Shared-Memory Parallel Programming". In: *International Journal of Parallel Programming* 42 (Aug. 2014). DOI: 10.1007/s10766-013-0262-9.