

ARCHITECTURE INDEPENDENT PARALLEL MODELING  
THE BULK-SYNCHRONOUS PARALLEL (BSP) MODEL, AND  
THE LOGP MODEL

**Disclaimer:** THESE NOTES DO NOT SUBSTITUTE THE TEXTBOOK FOR THIS CLASS. THE NOTES SHOULD BE USED IN CONJUNCTION WITH THE TEXTBOOK AND THE MATERIAL PRESENTED IN CLASS. IF A STATEMENT IN THESE NOTES SEEMS TO BE INCORRECT, REPORT IT TO THE INSTRUCTOR SO THAT IT BE FIXED IMMEDIATELY. THESE NOTES ARE ONLY DISTRIBUTED TO THE STUDENTS TAKING THIS CLASS WITH A. GERBESSIOTIS IN FALL 2004; DISTRIBUTION OUTSIDE THIS GROUP OF STUDENTS IS NOT ALLOWED.

**Required Reading.** L.G.Valiant. A bridging model for parallel computation. Communications of the ACM, 33(8):103-111, August 1990, accessible through the Course Web-page.

**Recommended Reading.** Culler, Karp, Patterson et al. LogP: Towards a Realistic model of parallel computation. Proceedings of the 4th ACM SIGPLAN Symposium on Principles and practice of parallel programming, San Diego, CA, USA. Pages 1-12. 1993.

## *BSP Model* Motivation

---

Since the introduction and adaptation of the von-Neumann model for sequential computing, the effects of computer revolution on society have been pretty significant. A general purpose computer performs well on computer programs written on a variety of standardized programming languages like C, Fortran, Cobol or Lisp and the same computer program can be easily ported on other platforms.

It has always been realized that parallel computers will eventually supersede sequential machines. This has yet to happen despite advances in computer technology and the fact that chip technology seems to have reached physical limitations; nowadays, fast machines are not much faster than the slowest ones which are as fast (or perhaps faster) as a supercomputer of twenty years ago. Small incremental improvements that may lead to stagnation of sequential computing seem inevitable. Despite these shortcomings of sequential computing, there has been no significant spread of use of parallel computers and few companies have realized that their future may rely on parallel platforms. The main reason for this has been that parallel computer platforms are built in such a way that are too hardware specific, programs written for them exhibit poor performance unless the programmer fine-tunes its code to take into consideration features of the particular architecture. Not only the code is non-portable but scalability comes at a high cost as well. On the other hand parallel algorithms designed and analyzed by the theorists work on parallel models that usually ignore communication and/or synchronization issues, like the PRAM and its variants, and work only under unlimited parallelism assumptions.

## *BSP Model*

# Parallel Computing Failures

---

One of the earliest attempts to model a parallel computer has been the Parallel Random Access Machine (PRAM) which is one of the most widely studied abstract parallel models. A PRAM consists of a collection of processors which work synchronously and which communicate with a global shared random access memory which can access in unit time. There are many different types of PRAMs which are distinguished from the way they access the shared memory (eg CRCW, EREW PRAMs). Numerous parallel algorithms have been developed for this parallel computer model.

More realistic models of parallel computing view a parallel computer as a collection of sequential processors, each one having its own local memory (*distributed-memory model*). The processors are interconnected by a network which allows them to communicate by sending and receiving messages. Constraints such as the maximum number of pins on a chip, or the maximum width of a data bus, limit the capacity of a processor to communicate with any other processor. It is only possible for a single processor to communicate directly with few others, in most cases those physically close to it. If a message needs to be sent to a distant processor it is relayed through a number of intermediate processors.

As it has already been mentioned, the parallel machines built in the 1980s and early 90s failed to garner general acceptance mainly because of the lack of a stable, unified and bridging parallel programming model. These deficiencies made programming of such machines difficult (cf. assembly vs higher level programming languages), time consuming, non-portable and architecture-specific. Recently, the introduction of realistic parallel computer models such as the Bulk-Synchronous Parallel (BSP) model of computation by L.G. Valiant comes to address these limitations of parallel computing. Our hope is that further architectural convergences will occur with the goal of writing software that will be portable and run with high performance on a variety of architectures from networks/clusters of workstations (NOW/COW) to parallel supercomputers.

## *BSP Model* The Bulk-Synchronous Parallel model

---

The *Bulk-Synchronous Parallel* (BSP) model of computation has been proposed by L.G. Valiant, as a **unified** framework for the design, analysis and programming of general purpose parallel computing systems. It allows the design of algorithms that are both **scalable and portable**.

In a BSP program, processors jointly advance through its various phases called **supersteps** with the required computations and remote communication occurring between them; at the end of a superstep processors check themselves in order to proceed to the following superstep.

The BSP model consists of three parts:

- (1) a collection of *processor-memory components*,
- (2) a *communication network* that can deliver messages point-to-point among the components, and
- (3) a *facility for global synchronization*, in barrier style, of all or a subset of the components.

A time step (as opposed to a CPU instruction or cycle) would refer to the time needed to perform a local computation (such as a fetch from memory and a floating-point operation followed by a store operation).

It should be noted that, although the model stresses global barrier-style synchronization, pairs of processing units may synchronize pairwise by sending messages to and from an agreed memory location. However, such message exchanges should respect the superstep rules.

As mentioned, computation on the BSP model proceeds in a succession of *supersteps*. A superstep may be thought of as a segment of computation during which each processor performs a given task using data already available there locally before the start of the superstep. Such a task may include (i) local computations, (ii) message transmissions, and (iii) message receipts.

## Hardware Abstraction: Parameters ( $p, L, g$ )

---

The tuple  $(p, L, g)$  characterizes the behavior and performance of a BSP computer.

- $p$  is the number of components available.
- $L$  is the minimum time between successive synchronization operations,
- $g$  is the ratio of the total throughput of the whole system (in a steady state, i.e. in a state of continuous message usage) in terms of basic computational operations, to the throughput of the communication network in terms of words of information delivered.
  
- A lower bound on the value of  $L$  is the time for a remote memory operation/message dispatch to become effective and is thus dependent on say, the diameter of the interconnection network.
- The time for barrier synchronization also poses a lower bound on the effective value of  $L$ .
- An upper bound on  $L$  is application specific and expressed in terms of problem size  $n$  as well.
- The value of  $g$  is measured while the network is in a steady state, i.e. latency issues become insignificant in the measurement of communication time; parameter  $L$  is large enough for the theoretical bound on  $g$  to be realizable.

The theoretical definition of  $g$  relates to the routing of  $h$ -relations; when each processor sends or receives at most  $h$  messages (of the same size) an  $h$ -relation is realized and the cost assigned to this communication is  $gh$  provided that  $h \geq h_0$ , where  $h_0$  is a machine dependent parameter.

Otherwise the cost of communication is  $L$ .

This way latency issues associated with small message size/messages are taken into consideration.

In practice, and in various libraries implementing the BSP model message requests issued/accepted by processors are of variable size. The  $h$  of an  $h$ -relation relates then to the total size of communicated data and  $g$  is expressed in terms of basic computational operations (sometimes, floating-point operations) or absolute time (seconds) per data-unit (a byte or word of information).

In practice, the parameter  $L$  of the BSP model not only hides the cost of communication when each processor sends or receives a small number of messages but also the cost of communication where each processor may send or receive a large number of messages but each one is of small size.

For any BSP computer, the values of  $L$  and  $g$  are likely to be a non-decreasing functions of  $p$ .

The use of  $L$  and  $g$  to characterize the communication and synchronization performance of a BSP computer is important because such issues are abstracted in only two parameters thus allowing considerations to be moved from a local level to a global one.

For the sake of an example, the values for  $L$  and  $g$  for some abstract network topologies are as follows. A ring has  $L = O(p)$ ,  $g = O(p)$ , a 2d-array (mesh)  $L = O(\sqrt{p})$ ,  $g = O(\sqrt{p})$ , a butterfly  $L = O(\log p)$ ,  $g = O(\log p)$ , and a hypercube  $L = O(\log p)$ ,  $g = O(1)$ . For the case of a hypercube, as  $g = O(1)$  the cost of routing a permutation, i.e. an one-relation is not  $1 \cdot g$  but  $L$ . Thus  $h_0$  for the hypercube is such that  $h_0 = \Theta(\log p)$ .

## BSP Model

### Cost Modeling under the BSP model

---

If in a superstep

- (1) an  $h$ -relation is realized, and
- (2)  $x$  computational operations are performed,

then, the cost of the superstep is given by the following formula.

- (1)  $\max \{L, x + gh\}$ .

Alternative costs are  $\max \{L, x, gh\}$  and  $L + x + gh$ .

The maximum size of a superstep depends on the problem in hand. Under the BSP programming paradigm the objective is to maximize the size of supersteps, decrease their number and increase processor utilization.

The description of BSP programs can be simplified by separating computation and communication and assuming that each superstep contains either local computations or communication.

Machine	Mflop ( $\geq$ )	$p$	$L$ ( $\geq$ ) in flops/ $\mu$ sec	$g$ in flops/word/ $\mu$ sec/word
SGI Power Challenge	80	4	2000/26	9.3/0.13
MultiProc Sun	10	4	120/12	4.1/0.41
IBM SP2	30	4	3600/140	8.0/0.30
Digital Alpha Farm	10	4	47000/4664	81/8.10
PCs in Cluster	200	4	180000/900	65/0.325

## Optimality of Algorithms under the BSP model

---

Two modes of programming on the BSP model were envisaged: **automatic mode** where programs are written, say PRAM style, in a high level language that hides memory distribution from the user and **direct mode** where the programmer retains control of memory allocation. In the **direct mode** of programming small multiplicative constant factors in runtime are important. It can be shown that the automatic mode achieves optimality within constant factors by simulating say PRAM algorithms on the BSP.

The term **slack** in the context of algorithm design refers to the ratio of the problem size over the processor number of the BSP machine. The question is whether the direct mode can be beneficial in circumstances where:

- small multiplicative constant factors in runtime are important,
- where smaller problem instances can be run more efficiently in direct mode (less slack is required) than in automatic mode,
- where the available machine has high  $g$  (automatic mode requires  $g$  to be constant), and
- $L$  is high for direct but not for automatic mode for the problem instance in hand.



## One-Optimality under the BSP model

---

Although it is difficult to measure performance to high accuracy, since the operations that are counted have to be carefully defined, we can make such measures meaningful by measuring ratios between runtimes on pairs of models that have the same set of local instructions.

The performance of a BSP algorithm  $P$  is thus described in three parts.

- A sequential algorithm  $S$  with which we are comparing  $P$  is first specified.
- The model of computation used for both algorithms is then defined and the basic computational operations that will be counted in both  $P$  and  $S$  are also described and the charging policy is made explicit.
- Second, two ratios  $\pi$  and  $\mu$  are specified.
  - $\pi$ , is the ratio between the computation time  $C_P$ , of the BSP algorithm, over the time  $C_S$  of the comparing sequential algorithm divided by  $p$ , i.e.,  $\pi = pC_P/C_S$ .
  - $\mu$ , is the ratio between the communication time  $M_P$  required by the communication supersteps of the BSP algorithm and the computation time of  $S$  divided by  $p$ , i.e.,  $\mu = pM_P/C_S$ .
  - When communication time is described, it is necessary that the amount of information that can be conveyed in a single message be made explicit.
  - Finally, conditions on  $n$ ,  $p$ ,  $L$  and  $g$  are specified that are sufficient for the algorithm to be plausible and the claimed bounds on  $\pi$  and  $\mu$  to be valid.

Corollaries describe sufficient conditions for the most interesting optimality criteria, such as  $c$ -optimality, i.e.,  $\pi = c + o(1)$  and  $\mu = o(1)$ . All asymptotic bounds refer to the problem size as  $n \rightarrow \infty$ .

## Software Support under the BSP model

---

The BSP model, unlike other models of parallel computation, is not just an architectural-oriented theoretical model; it can also serve as a paradigm for programming parallel computers.

- (1) The fundamental concept introduced by the BSP model is the notion of the superstep, and that all remote memory accesses occur between supersteps as part of a global operation among the processors the results of these accesses become effective at the end of the current superstep.

Although it may have been apparent why some consider the BSP as a satisfactory unifying and bridging model for parallel computation, one may ask the question of how successful it has been as a practical model and what level of software support there is for BSP. The BSP model has been realized as a library of functions for process creation and destruction, remote memory access and message passing, and global, barrier-style, synchronization. The abstraction offered by the BSP model is such that any library offering such facilities can be used for programming according to the BSP programming paradigm. The Oxford BSP Library that supports Direct Remote Memory Access (DRMA) for parallel programs, the Green BSP library that supports message passing and the Oxford BSP Toolset that supports both DRMA and message passing are some of the libraries that specifically allow programming according to the BSP paradigm. Just as the von-Neumann model encompasses various programming language paradigms eg. functional, logic programming, the BSP does not dictate a particular mode of programming as well. All three libraries present a particular set of choices to the user. Some of the elements of this set are:

- **Data Parallel Program Structure.** It allows large-scale parallelism by splitting data and concurrently working on the individual pieces.
- **SPMD programs.** A Single Program Multiple Data programming style is used as perhaps implied by the structure of supersteps.
- **Direct mode of global memory management.** The programmer has direct access to memory allocation and determines how data are partitioned.
- **Static processor allocation.** The number of participating processors is determined in the beginning of the execution and cannot vary during the computation dynamically.

## *BSP Model*

### Performance vs Running Time Prediction

---

Whereas performance of a particular BSP algorithm can be reliably predicted, one *should not* expect the BSP cost model to accurately predict the running time and behavior of a particular implementation. Accurate prediction is difficult even for sequential algorithms due to the existence of varying memory hierarchies on real machines; adding parallelism and the side-effects of communication introduces two more difficulties.

There have been various attempts to more accurately predict parallel performance by extending the BSP model. The E-BSP model is one such approach and is more explicit and specific about the communication network of a particular hardware platform, and the patterns of communication involved in routing. The attractiveness of the BSP cost model is its simplicity and generality; introducing more parameters to describe the performance of a communication network under various patterns of communication increases the complexity of describing the performance of all but the simplest algorithms with perhaps only small gains in prediction accuracy. Such an approach may also be problematic as it is the main reason parallel computing failed in the past: the attempt to realize for a particular algorithm those patterns of communication that are optimal for a given platform (and communication network), whereas they may lead in significant degradation in performance, if utilized in other platforms (portability vs. efficiency). Another variant of the BSP model introduces one more parameter,  $B$ , related to message size, and associates  $g$  with that message size, enforcing this way coarse-grained communication of messages of size equal to  $B$ . The original BSP model does not elaborate in detail between fine-grained and coarse-grained communication. If small  $h$ -relations are communicated (where “small” is to mean less than some parameter  $h_0$ , usually assumed to be equal to  $L/g$ ) a cost  $L$  is assigned to such a communication; no mention of message size is inferred. Presumably, for BSP to be an abstract and general-purpose model, details of how communication is performed efficiently are left to the BSP library implementor. In practice, the generality of the BSP model works well if one interprets the cost model so as to absorb in  $L$  not only “small”  $h$ -relations but also “small” messages; therefore, the value of  $B$  is reflected in the choice of  $L$  and  $g$  (as is,  $h_0$  as well) without the need of introducing an extra parameter.

## *The LogP Model*

### Introduction

---

The LogP model (Culler, Karp, Patterson, Sahay, Schauser, Santos, Subramonian and von Eicken, 1993) has also been suggested as a realistic model for the design of parallel algorithms that work predictably well on a wide range of parallel machines. It models a parallel machine as a distributed memory multiprocessor in which processors communicate by point-to-point messages. It is an *asynchronous* model that does not enforce synchronization of the processors, as the BSP model does. If processor synchronization is required in a program, the programmer must provide it. A parallel machine under the LogP model can be characterized by the following tuple  $(p, L, g, o)$  of parameters. Each parameter is explained in more detail below (note that LogP may use the same notation for some of its parameters as the BSP model but these may have different meaning).

## The LogP Model Parameters

---

$p$ : The number of processor/memory components (as in the BSP model).

$L$ : an upper bound on the *latency* or delay of the machine, incurred while communicating a message of very small size (one or a small number of words) from a source to a destination component. In other words, such messages are delivered by the router within time  $L$ .

$o$ : the *overhead*, defined as the length of time that a processor is engaged in the transmission or receipt of each message (i.e the time required for the submission of a message to the router or acquisition of a received message from it).

$g$ : the *gap*, defined as the minimum time interval between consecutive message transmissions or consecutive message receptions at a processor.  $1/g$  gives the bandwidth of the system per processor.

In addition, the communication capacity of the system is limited. This means that at most  $L/g$  messages can originate or be destined to any processor at any time. If a processor needs to transmit a message that would violate this condition, it stalls until this transmission is possible (i.e. there is some available capacity). Writing under the LogP model into a remote memory location takes time  $L + 2o$ .

## *The LogP Model*

### BSP vs LogP

---

Although BSP and LogP are similarly powered models (one can simulate a BSP computation on the LogP and the other way around), the BSP model seems to be more usable as a theoretical model for the design and analysis of parallel algorithms and as a programming paradigm for writing parallel software that is scalable and transportable (portable and efficient among a variety of hardware platforms).

The importance of BSP is that it introduces a new abstraction for communication in terms of the BSP parameters  $L$  and  $g$  so that considerations in programming parallel machines move from a local (detailed) level to a global (abstract) one. The two parameters abstract all communication and synchronization issues related to parallel computing and allow the design of software (a unifying property) that work on any machine independent of the underlying architecture (eg. shared memory vs distributed memory).

Under the BSP model, an algorithm designer describes the performance of a parallel algorithm in terms of  $p$ ,  $L$ ,  $g$  and problem size  $n$ . A collection of algorithms that solves a given problem can then be formed with varying performance characteristics (for example, one algorithm may be suitable for machines with small  $L$ , another for machines with small  $g$  or  $n$ , and so on). For a given machine whose BSP parameters are known or are measurable, that algorithm is chosen from the collection whose performance (computation and communication) on the particular machine is the best available.