

Towards a highly scalable Multi-BSP fast Fourier transform Albert-Jan N. Yzelman^{1,2}, Rob H. Bisseling³, Dirk Roose², and Karl Meerbergen²

Corresponding author: Albert-Jan.Yzelman@cs.kuleuven.be

Bulk Synchronous Parallel (BSP)

Valiant introduced the BSP model in 1990. The model consists of (1) an abstract **parallel computer**, (2) an imposed structure on **parallel algorithms**, and (3) the resulting **cost function**.

A BSP computer executes *p* processes concurrently, and does so at a homogeneous speed r. Each process has access to its own local memory. Inter-process communication happens through a full-duplex blackbox network.

Preparing this network for all-to-all communication and synchronisation costs *l* units, while a simultaneous send and receive of one data word costs g units. In illustration, an abstract BSP computer (*p*,*r*,*g*,*l*):



The network is coloured green, memory is coloured blue, and units that execute the BSP processes are in black.

A BSP algorithm separates local computation phases from communication phases. There is no computation while communicating, and vice versa. In illustration:







The cost of the *i*th computation phase depends on the BSP process *s* which was assigned the most work *w*:

$$T_{\rm comp}^i = \max_{0 \le s < p} w_i^{(s)}$$

If in the *i*th communication phase process *s* transmits *t* words and receives *r* words, the *h*-relation gives us the linear part of the communication cost in g. Adding the constant latency cost *l* yields:

$$T_{\text{comm}}^{i} = \left(\max_{0 \le s < p} \max\{t_{i}^{(s)}, r_{i}^{(s)}\}g\right) + l = h_{i}g + l$$

¹Flanders ExaScience Lab (Intel Labs Europe), Belgium ² Department of Computer Science, KU Leuven, Belgium ³ Department of Mathematics, Utrecht University, the Netherlands

Fast Fourier transform (FFT)

Consider the FFT on complex vectors *x*, *y* of power-of-two length *n*, which computes $y = F_n x$,

$$y_i = \sum_{j=0}^{n-1} x_j e^{-2\pi i i j/n}$$

Let the matrix *B* be the butterfly matrix containing the weights *W*, and let *V* be the bit-reversal matrix. Then the FFT can succinctly be written as follows (see also the illustration on the right).

$$F_n x = \left[\prod_{i=0}^{m-1} \left(I_{2^i} \otimes B_{2^{m-i}} \right) \right] V_n x$$



 $W_{n/2}$

$$u_2 = diag(1, e^{-2\pi i 1/n}, \dots, e^{-2\pi i (n/2-1)/n})$$

Butterfly matrix weights



Butterfly matrix (left), even-odd sort matrix (right)



Structure of the butterfly matrices involved for n=8



For parallelisation, the computation is split in two. For integer m, q we assume $n=2^{m}$ and $p=2^{q}$. The default value for t is m-q.

The Multi-BSP FFT applies the flat BSP FFT algorithm recursively by splitting the computations *Ln* and *Rn* as well. Note only the leaves of the Multi-BSP tree can compute. An illustration for L=2, $p_0=2$, $p_1=2$, and n=8:



Implementation of the (Multi-BSP) FFT uses **MulticoreBSP for C**, and is freely available at: www.multicorebsp.com

Ref.: A. N. Yzelman, R.H. Bisseling, D. Roose, K. Meerbergen, "MulticoreBSP for C: a high-performance library for shared-memory parallel programming", International Journal of Parallel Programming, in press (2013).

The Multi-BSP model

Modern and future supercomputers consist of nodes with shared-memory parallel capabilities, interconnected by increasingly hierarchical networks. This gives rise to Non-Uniform Memory Access (NUMA) effects on various levels of the compute hierarchy. Valiant updated BSP to account for this situation, resulting in the Multi-BSP model (2008):



An abstract Multi-BSP computer. Each node in the tree is a BSP computer with local memory. Only the leaf nodes contain compute elements.



A Multi-BSP algorithm operating on the various levels of a Multi-BSP computer hierarchy. Successive communication phases are useful when each op*erates on a different level of the communication network.*

Writing *L* for the number of levels in the Multi-BSP computer hierarchy, *N* for the number of computation phases, and *Me* for the number of communication phases at the *e*th level of the Multi-BSP tree. Then:

This directly leads to the Multi-BSP cost model: $T = \sum_{i=0}^{N-1} T_{\rm comp}^{i} + \sum_{e=0}^{L-1} \sum_{i=0}^{M_e-1} T_{\rm comm,e}^{i}$



$$T_{\mathrm{comm},e}^{\iota} = h_i g_e + l_e$$

If *L*=1 and *M*₀=*N*, this reduces to the flat BSP model: $T = \sum_{i=0}^{N-1} \left(T_{\rm comp}^i + T_{\rm comm}^i \right)$