

## Overview

We present high-performance techniques for FFT-based dense polynomial multiplication on multi-cores, with a focus on **unbalanced input data**. We show that *balanced data* can **maximize parallel speed-up** and **minimize cache complexity** for bivariate multiplication. Then, we show how multivariate (and univariate) multiplication can be efficiently reduced to *balanced bivariate multiplication*. This approach brings in practice dramatic improvements for unbalanced input data, even if the implementation relies on serial one-dimensional FFTs. This permits the use of memory efficient FFT techniques, such as TFFT, which is hard to parallelize.

## FFT-based Multivariate Multiplication

Let  $\mathbb{K}$  be a field and  $f, g \in \mathbb{K}[x_1 < \dots < x_n]$  be polynomials. Define  $d_i = \deg(f, x_i)$  and  $d'_i = \deg(g, x_i)$ , for all  $i$ . Assume there exists a primitive  $s_i$ -th root  $\omega_i \in \mathbb{K}$ , for all  $i$ , where  $s_i$  is a power of 2 satisfying  $s_i \geq d_i + d'_i + 1$ . Then  $fg$  can be computed as follows.

**Step 1.** Evaluate  $f$  and  $g$  at each point of the  $n$ -dimensional grid  $((\omega_1^{e_1}, \dots, \omega_n^{e_n}), 0 \leq e_1 < s_1, \dots, 0 \leq e_n < s_n)$  via  $n$ -D FFT.

**Step 2.** Evaluate  $fg$  at each point  $P$  of the grid, simply by computing  $f(P)g(P)$ ,

**Step 3.** Interpolate  $fg$  (from its values on the grid) via  $n$ -D FFT.

## Complexity Estimates

• Let  $s = s_1 \dots s_n$ . The number of operations in  $\mathbb{K}$  for computing  $fg$  based on FFTs is

$$\frac{9}{2} \sum_{i=1}^n (\prod_{j \neq i} s_j) s_i \lg(s_i) + (n+1)s = \frac{9}{2} s \lg(s) + (n+1)s.$$

• Assuming serial 1-D FFTs,  $\frac{9}{2}(s_1 \lg(s_1) + \dots + s_n \lg(s_n))$  is the span of **Step 1** and the **parallelism** is lower bounded by

$$s / \max(s_1, \dots, s_n). \quad (1)$$

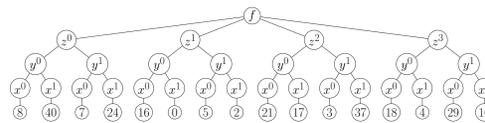
• Let  $L$  be the size of a cache line. For some constant  $c > 0$ , the number of cache misses of **Step 1** is upper bounded by

$$n \frac{cs}{L} + cs \left( \frac{1}{s_1} + \dots + \frac{1}{s_n} \right). \quad (2)$$

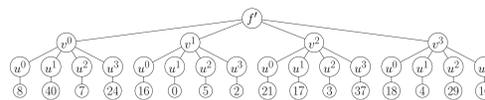
• **Remark:** For  $n \geq 2$ , Expr. (2) is minimized at  $n = 2$  and  $s_1 = s_2 = \sqrt{s}$ . Moreover, when  $n = 2$ , under a fixed  $s = s_1 s_2$ , Expr. (1) is maximized at  $s_1 = s_2 = \sqrt{s}$ .

## Contraction to Bivariate

• **Example.** Let  $f \in \mathbb{K}[x, y, z]$  where  $\mathbb{K} = \mathbb{Z}/41\mathbb{Z}$ , with  $\deg(f, x) = \deg(f, y) = 1$ ,  $\deg(f, z) = 3$  and recursive dense representation:

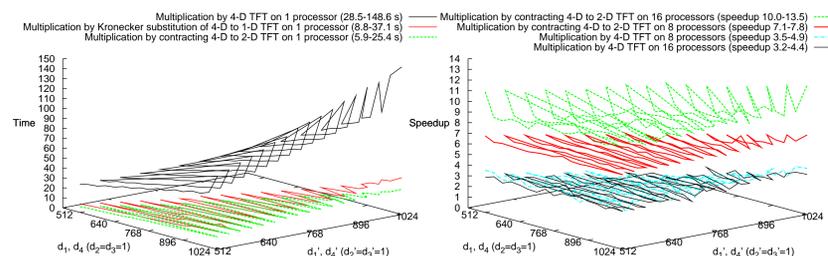


Contracting  $f(x, y, z)$  to  $f'(u, v)$  by  $x^{e_1}y^{e_2}z^{e_3} \mapsto u^{e_1+2e_2}, v^{e_3}$ :



• **Remark.** The data is “essentially” unchanged by contraction, which is a property of recursive dense representation.

• Below, the **left** figure displays the timing of 4-variate multiplication via 4-D TFFT, 1-D TFFT by Kronecker substitution and contraction to balanced 2-D TFFT on 1 processor; The **right** figure shows the speedups of 4-variate multiplication using 4-D TFFT and contraction to balanced 2-D TFFT on 8 and 16 processors.

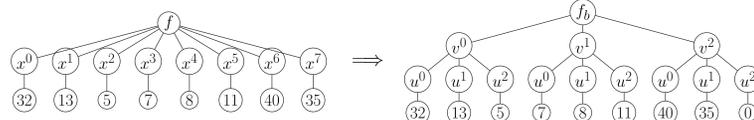


## Extension from Univariate to Bivariate

• **Example:** Consider  $f, g \in \mathbb{K}[x]$  univariate, with  $\deg(f) = 7$  and  $\deg(g) = 8$ ;  $fg$  has “dense size” 16. We obtain an integer  $b$ , such that  $fg$  can be performed via  $f_b g_b$  using “nearly square” 2-D FFTs, where  $f_b := \Phi_b(f)$ ,  $g_b := \Phi_b(g)$  and

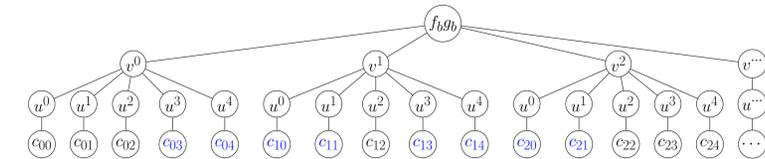
$$\Phi_b : x^e \mapsto u^{e \bmod b} v^{e \text{ quo } b}.$$

Here  $b = 3$  works since  $\deg(f_b g_b, u) = \deg(f_b g_b, v) = 4$ ; moreover the dense size of  $f_b g_b$  is 25. Extending  $f(x)$  to  $f_b(u, v)$  gives

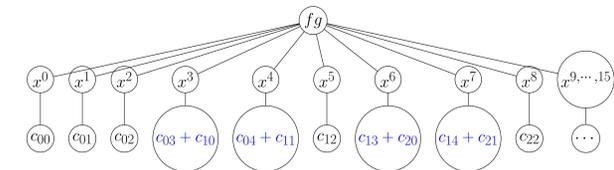


• **Proposition:** For any non-constant  $f, g \in \mathbb{K}[x]$ , one can always compute  $b$  such that  $|\deg(f_b g_b, u) - \deg(f_b g_b, v)| \leq 2$  and the dense size of  $f_b g_b$  is **at most twice** that of  $fg$ .

• **Example (ctnd):** Computing the bivariate product  $f_b g_b$ :



Converting back to  $fg$  from  $f_b g_b$  requires only to traverse the coefficient array once and perform **at most  $\deg(fg, x)$  additions**.



## Balanced Multiplication

• **Definition.** A pair of bivariate polynomials  $p, q \in \mathbb{K}[u, v]$  is *balanced* if  $\deg(p, u) + \deg(q, u) = \deg(p, v) + \deg(q, v)$ .

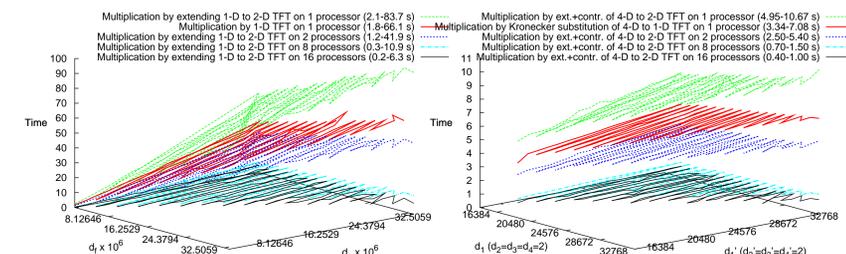
• **Algorithm.** Let  $f, g \in \mathbb{K}[x_1 < \dots < x_n]$ . W.l.o.g. one can assume  $d_1 \gg d_i$  and  $d'_1 \gg d'_i$  for  $2 \leq i \leq n$  (up to variable re-ordering and contraction). We obtain  $fg$  by

**Step 1.** Extending  $x_1$  to  $\{u, v\}$ .

**Step 2.** Contracting  $\{v, x_2, \dots, x_n\}$  to  $v$ .

Determine the above extension  $\Phi_b$  such that  $f_b, g_b$  is (nearly) a **balanced pair** and  $f_b g_b$  has dense size at most twice that of  $fg$ .

• The **left** figure shows the timing of univariate multiplication via 1-D TFFT and extension to balanced 2-D TFFT on 1, 2, 8, 16 processors; The **right** one shows timing of our balanced multiplication for an unbalanced 4-variate case on 1, 2, 8, 16 processors vs the method based on 1-D TFFT via Kronecker substitution.



**Acknowledgements.** This work was supported by NSERC and MITACS NCE of Canada, and NSF Grants 0540248, 0615215, 0541209, and 0621511. We are very grateful for the help of Professor Charles E. Leiserson at MIT, Dr. Matteo Frigo and all other members of Cilk Arts.