

Balanced Dense Polynomial Multiplication on Multi-cores

Marc Moreno Maza

Ontario Research Centre for Computer Algebra
University of Western Ontario, London, Canada
moreno@csd.uwo.ca

Yuzhen Xie

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, Cambridge, USA
yxie@csail.mit.edu

Abstract— In symbolic computation, polynomial multiplication is a fundamental operation akin to matrix multiplication in numerical computation. We present efficient implementation strategies for FFT-based dense polynomial multiplication targeting multi-cores. We show that *balanced input data* can maximize parallel speedup and minimize cache complexity for bivariate multiplication. However, *unbalanced input data*, which are common in symbolic computation, are challenging. We provide efficient techniques, what we call *contraction* and *extension*, to reduce multivariate (and univariate) multiplication to *balanced bivariate multiplication*. Our implementation in Cilk++ demonstrates good speedup on multi-cores.

Keywords- parallel symbolic computation; parallel polynomial multiplication; parallel multi-dimensional FFT/TFT; Cilk++; multi-core;

I. INTRODUCTION

Polynomials and matrices are the fundamental objects on which most computer algebra algorithms operate. In the last decade, significant efforts have been deployed by different groups of researchers for delivering highly efficient software packages for computing symbolically with polynomials and matrices. Among them: LinBox [12], MAGMA [13], NTL [16]. However, most of these works are dedicated to serial implementation, in particular in the case of polynomials. None of the computer algebra software packages available today offers parallel implementation of asymptotically fast algorithms for polynomial arithmetic. The work reported hereafter aims at filling this gap.

We present high-performance techniques for the implementation of multivariate polynomial multiplication targeting multi-cores. Symbolic computations with polynomials rely, indeed, directly or indirectly on multiplication. We commit ourselves to polynomials over finite fields since the so-called *modular techniques* reduce all computations to such fields of coefficients. In addition, we focus on dense polynomial arithmetic because most computer algebra algorithms, such as the Euclidean Algorithm and its variants, tend to densify intermediate data, even when the input and output polynomials are sparse. See Chapter 5 in [8] for an extensive presentation of these ideas.

Dense representations permit the use of multiplication algorithms based on *Fast Fourier Transform* (FFT) which run in quasi-linear sequential time w.r.t. output size, when counting the number of operations on coefficients. This result holds for univariate as well as for multivariate polynomials. We observe that reducing multivariate multiplication to univariate one through Kronecker's substitution is not

an option in our context. Indeed, this would lead us to manipulate univariate polynomials of very large degrees, say in the order of a machine word. Meanwhile, we aim at computing over field Z/pZ where p is a machine word prime number, for efficiency reasons. Therefore, we would not always be able to find in Z/pZ the appropriate primitive roots of unity for performing a Cooley-Tukey 1-D FFT.

In the multivariate case, the row-column algorithm for multi-dimensional FFT, reviewed in Section II, proceeds one dimension after another and performs several one-dimensional FFTs along one dimension at a time. This yields concurrent execution without even requiring that each one-dimensional FFT is computed in a parallel fashion. We take advantage of this flexibility to make use of non-standard and memory-efficient one-dimensional FFT techniques, such as Truncated Fourier Transform (TFT), for which no efficient parallel algorithm is known. More importantly, we do not seek a very fine grain of parallelism in our multiplication code since it will itself be a low-level routine in higher-level codes for computing polynomial GCDs and solving polynomial systems, for which parallel algorithms are available and distributed computing is desired.

Efficient implementation of algorithms on multi-cores makes necessary to consider complexity measures such as parallel speedup and cache complexity. We analyze the performances of dense multiplication based on row-column multi-dimensional FFT for these complexity measures in Section III. On bivariate input and when the partial degrees of the product are equal, the performances are nearly optimal; we call *balanced* this degree configuration. When the ratio between these two partial degrees is large, our experimentation confirms poor performances.

Motivated by these theoretical and experimental results, we show how multivariate multiplication can be efficiently reduced to *balanced bivariate multiplication*, based on 2-D FFT. With respect to a multiplication based on n -dimensional FFT, our approach may increase the input data size by at most of factor of 2. However, it provides much larger parallel speedup as reported in our experimentation.

Our approach combines two fundamental techniques that we call *contraction* and *extension*, presented in Sections IV and V. The first one reduces multivariate multiplication to bivariate one, without ensuring that dimension sizes are equal; however, the work remains unchanged and in many practical cases the parallelism and cache complexity are improved substantially.

The technique of extension turns univariate multiplication to bivariate one. This has several applications. First, it permits to overcome the difficult cases where primitive roots of unity of “large” orders cannot be found in the field of coefficients. Secondly, combined with the technique of contraction, this leads in Section VI to balanced bivariate multiplication.

The techniques proposed in this paper are implemented in the Cilk++ language [3], which extends C++ to the realm of multi-core programming based on the multi-threaded model realized in [7]. The Cilk++ language is also equipped with a provably efficient parallel scheduler by work-stealing [2]. We use the serial C routines for 1-D FFT and 1-D TFT from the `modpn` library [10]. Our integer arithmetic modulo a prime number relies also on the efficient functions from `modpn`, in particular the improved Montgomery trick [14], presented in [11]. This trick is another important specificity of 1-D FFTs over finite fields which makes their parallelization even more difficult. All our benchmarks are carried out on a 16-core machine with 16 GB memory and 4096 KB L2 cache. All the processors are Intel Xeon E7340 @ 2.40GHz.

II. BACKGROUND

Throughout this paper \mathbb{K} designates the finite field Z/pZ with p elements, where $p > 2$ is a prime number. In this section, we review algorithms and complexity results for multiplying multivariate polynomials over \mathbb{K} by means of FFT techniques. We start by stressing the specificities of performing FFTs over finite fields.

A. FFTs over Finite Fields

Using the Cooley-Tukey algorithm [4] (and its extensions such as Bluestein’s algorithm) one can compute the *Discrete Fourier Transform* (DFT) of a vector of s complex numbers within $O(s \lg(s))$ scalar operations. For vectors with coordinates in the prime field \mathbb{K} two difficulties appear with respect to the complex case.

First, in the context of symbolic computation, it is desirable to restrict ourselves to radix 2 FFTs since the radix must be invertible in \mathbb{K} and one may want to keep the ability of computing modulo small primes p , even $p = 3, 5, 7, \dots$ for certain types modular methods, such as those for polynomial factorization; see [8, Chapter 14] for details. As a consequence the FFT of a vector of size s over \mathbb{K} has the same running time for all s in a range of the form $[2^\ell, 2^{\ell+1})$. This *staircase* phenomenon can be smoothed by the so-called *Truncated Fourier Transform* (TFT) [9]. In most practical cases, the TFT performs better in terms of running time and memory consumption than the radix-2 Cooley-Tukey Algorithm; see the experimentation reported in [11]. However, the TFT has its own practical limitations. In particular, no efficient parallel algorithm is known for it.

Another difficulty with FFTs over finite fields comes from the following fact: a primitive s -th root of unity exists in \mathbb{K}

if and only if s divides $p - 1$. Therefore, the product of two univariate polynomials f, g over \mathbb{K} can be computed by evaluation and interpolation based on the radix 2 Cooley-Tukey Algorithm (see the algorithm of Section II-B with $n = 1$) if and only if the degree d of the product fg is less than the largest power of 2 dividing $p - 1$. When this holds, computing fg amounts to $\frac{9}{2} \lg(s)s + 3s$ operations in \mathbb{K} using the Cooley-Tukey Algorithm (and $\frac{9}{2}(\lg(s) + 1)(d + 1) + 3s$ operations in \mathbb{K} using TFT) where s is the smallest power of 2 greater than d . When this does not hold, one can use other techniques, such as the Schönage-Strassen Algorithm [8, Chapter 8], which introduces “virtual primitive roots of unity”. However, this increases the running time to $O(s \lg(s) \lg(\lg(s)))$ scalar operations.

B. Multivariate Multiplication

Let $f, g \in \mathbb{K}[x_1, \dots, x_n]$ be two multivariate polynomials with coefficients in \mathbb{K} and with n ordered variables $x_1 < \dots < x_n$. For each i , let d_i and d'_i be the degree in x_i of f and g respectively. For instance, if $f = x_1^3 x_2 + x_3 x_2^2 + x_3^2 x_1^2 + 1$ we have $d_1 = 3$ and $d_2 = d_3 = 2$. We assume the existence of primitive s_i -th roots of unity ω_i , for all i , where s_i is a power of 2 satisfying $s_i \geq d_i + d'_i + 1$. Then, the product fg is 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 multi-dimensional 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 multi-dimensional FFT.

The above procedure amounts to

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

operations in \mathbb{K} , where $s = s_1 \dots s_n$. In practice the benefit of using 1-D TFT instead of 1-D FFT increases with the number of variables and the number of cores in use. The cut-off criteria and a detail performance evaluation are reported in [15].

Consider now the following map from the monomials of $\mathbb{K}[x_1, \dots, x_n]$ to those of $\mathbb{K}[x_1]$:

$$x_1^{e_1} x_2^{e_2} x_3^{e_3} \dots x_n^{e_n} \mapsto x_1^{e_1 + \alpha_2 e_2 + \alpha_3 e_3 + \dots + \alpha_n e_n}$$

where $\alpha_2 = d_1 + d'_1 + 1$, $\alpha_3 = \alpha_2(d_2 + d'_2 + 1)$, \dots , $\alpha_n = \alpha_{n-1}(d_{n-1} + d'_{n-1} + 1)$. This induces a polynomial ring homomorphism Ψ (called the *Kronecker substitution*) from $\mathbb{K}[x_1, \dots, x_n]$ to $\mathbb{K}[x_1]$, hence a map satisfying $\Psi(a + b) = \Psi(a) + \Psi(b)$ and $\Psi(ab) = \Psi(a)\Psi(b)$, for all polynomials a, b . Moreover, one can check that fg is the only pre-image of $\Psi(f)\Psi(g)$. This latter polynomial has degree

$$\delta_n := (d_1 + d'_1 + 1) \dots (d_n + d'_n + 1) - 1.$$

It follows from this construction that one can reduce multivariate multiplication to univariate one. If \mathbb{K} admits primitive s -th roots of unity for $\delta_n < s = 2^\ell$ for some ℓ , then using the FFT-based multiplication, one can compute fg at at most $\frac{9}{2} \lg(s)s + 3s$ operations in \mathbb{K} . Using the TFT approach, this upper bound becomes $\frac{9}{2}(\lg(s) + 1)(\delta_n + 1) + 3s$.

Multivariate multiplications based on multi-dimensional FFT and Kronecker substitution have similar serial running time. However, the latter approach has at least two drawbacks. First, the field \mathbb{K} may not admit primitive s -th roots of unity. Recall that primitive s -th roots of unity exist in \mathbb{K} if and only if s divides $p - 1$, see [8, Chapter 8]. Secondly, as mentioned above, it is desirable to achieve efficient parallel multiplication without assuming that 1-D FFTs are performed in a parallel fashion.

III. MAIN RESULTS

The specificities of 1-D FFTs over finite fields, see Section II-A, lead us to the following hypothesis. We assume throughout this paper that we have at our disposal a **black box** computing the DFT at a primitive 2^ℓ -th root of unity (when \mathbb{K} admits such value) of any vector of size s in the range $(2^{\ell-1}, 2^\ell]$ in time $O(s \lg(s))$. However, we do not make any assumptions about the algorithm and its implementation. In particular, we do not assume that this implementation is a parallel one. As mentioned in the introduction, we do not seek a very fine-grained parallelism for our polynomial multiplication since it is meant to be a core-routine in higher-level parallel code. Therefore, we rely on the row-column multi-dimensional FFT to create concurrent execution in the algorithm presented in Section II-B.

This strategy has its own challenges. Suppose that one dimension x_1 has a small size s_1 , say in the order of units, whereas another dimension x_2 has size s_2 in the thousands. Then a lot of small FFTs have to be performed along x_1 while only a few large FFTs can be run simultaneously along x_2 . In the former case, the parallel overhead may dominate, reducing severely the benefits of concurrent runs. In the latter case, the measured speedup factor may simply be too small by lack of parallelism.

We formalize this remark in Section III-A where we give a lower bound for the parallel running time of the algorithm of Section II-B. Then, in Section III-B we give an upper bound for the number of cache misses of the same algorithm. We observe that these lower and upper bounds reach a “local” maximum and minimum respectively, when the number n of variables equals 2 and the dimension sizes of the 2-D FFT are equal. Therefore, the algorithm of Section II-B performs very well in terms of parallelism and cache complexity on bivariate polynomial input when the partial degrees of the product are equal. For this reason, we introduce in Section III-C the concept of *balanced bivariate multiplication*.

In Section III-D, we claim that dense multivariate multiplication can be efficiently reduced to balanced bivariate multiplication. Efficiently means here that the overheads of the reduction are in general much less than the performance gains. Sections IV to VI formally establish this reduction and prove its performances, including experimental results.

A. Parallel Running Time Estimates

Let us consider the parallel running time of the algorithm of Section II-B with the multi-threaded programming model of [7]. Under the assumption that 1-D FFT may be run sequentially, the following estimate holds for the span of **Step 1**:

$$\frac{9}{2} (s_1 \lg(s_1) + \dots + s_n \lg(s_n)).$$

Therefore, the parallelism (i.e. theoretical speedup) of **Step 1** is lower bounded by

$$\frac{s_1 \dots s_n \lg(s_1 \dots s_n)}{s_1 \lg(s_1) + \dots + s_n \lg(s_n)}$$

and thus by

$$s / \max(s_1, \dots, s_n). \quad (2)$$

Similar estimates can be given for **Step 3** while the costs of **Step 2** can be neglected comparing to the others.

Observe that $\max(s_1, \dots, s_n)$ is lower bounded by $s^{1/n}$. Hence, for a fixed s , one could conclude that the larger is n , the better. In practice, this would be a mistake. Suppose for instance that $s = 2^\ell$ for some ℓ . Assume first that each s_i can be set to 2, implying $n = \ell$. In this case the serial and parallel running time for **Step 1** are respectively given by $\frac{9}{2} \ell 2^\ell$ and $\frac{9}{2} \ell 2$; hence the theoretical speedup is $2^{\ell-1}$. Alternatively, we can set each s_i to be $2^{\ell/2}$, implying $n = 2$. The parallel running time now becomes $\frac{9}{2} \ell 2^{\ell/2}$; hence the theoretical speedup is $2^{\ell/2}$. Apparently the parallelism for the case $n = \ell$ is more attractive than that of $n = 2$. But this is neglecting parallel overhead! In our analysis of the case $n = \ell$, we are implicitly assuming that one can run concurrently $2^{\ell-1}$ threads with each of them doing very little work (actually computing a 1-D FFT of size 2). Not only is this not realistic as ℓ becomes large, but also this makes simply no sense since the overhead of executing a thread is certainly greater than the cost of executing an FFT of size 2. In the case $n = 2$, a theoretical speedup of $2^{\ell/2}$ is already good, say for $\ell \geq 20$. In addition, the work performed by each thread is an FFT of size $2^{\ell/2}$ which is larger (for the Cilk++ concurrency platform) than the overhead of executing that thread.

Our experimentation hereafter confirms that the case $n = 2$ performs better than $n = \ell$ for a fixed $s = 2^\ell$. Finally, we observe that for $n = 2$ and for a fixed s the lower bound $s / \max(s_1, \dots, s_n)$ is maximum at $s_1 = s_2 = \sqrt{s}$.

B. Cache Complexity Estimates

We now turn to cache complexity estimates, using the theoretical model introduced in [6]. We focus on **Step 1** again. Let L be the size of a cache line. We assume that the cache size is large enough such that all data involved by P concurrent runs of 1-D FFT (where P is the number of processors) fit in cache. This is justified in the experimentation with our *balanced bivariate multiplication* (see Section VI) where our 16-core machine has 4MB of L2 unified cache and each FFT vector has at most size 128KB. We also assume that our n -D FFT is performed without data transposition by loading directly from main memory to cache the vectors on which 1-D FFT is run. This technique is used in the implementation of the FFTW [5]. Therefore, cache misses arise essentially when reading data before performing a 1-D FFT. For a vector of size s_i the number of cache misses is at most $s_i/L + 1$. Thus the number of cache misses at **Step 1** and **Step 3** fits in

$$O(\sum_{i=1 \dots n} (\prod_{j \neq i} s_j)(s_i/L + 1)).$$

At **Step 2**, this number is within $O(\frac{s}{L} + 1)$. Hence, if $Q(s_1, \dots, s_n)$ denotes the total number of cache misses for the whole algorithm, we obtain

$$Q(s_1, \dots, s_n) \leq cs \frac{n+1}{L} + cs(\frac{1}{s_1} + \dots + \frac{1}{s_n}) \quad (3)$$

for some constant c . As in Section III-A let us consider $s = s_1 \dots s_n$ to be fixed. The following is easy to prove:

$$\frac{n}{s^{1/n}} \leq \frac{1}{s_1} + \dots + \frac{1}{s_n}.$$

Moreover this latter inequality is an equality when each s_i equals $s^{1/n}$. Noting $\frac{n+1}{n} \leq 2$ for $n \geq 1$ we deduce:

$$Q(s_1, \dots, s_n) \leq ncs(\frac{2}{L} + \frac{1}{s^{1/n}}) \quad (4)$$

when $s_i = s^{1/n}$ holds for all i . This suggests to minimize n , thus setting $n = 2$. Therefore, for fixed s , the upper bound of (3) reaches a local minimum at $n = 2$ and $s_1 = s_2 = \sqrt{s}$.

C. Balanced Bivariate Multiplication

The analysis of Sections III-A and III-B suggests that, for bivariate input, the algorithm of Section II-B is nearly optimum in terms of parallelism and cache complexity when $s_1 = s_2$, that is, when the partial degrees of the product are equal. This brings the definition and proposition below.

Definition 1: The pair of polynomials $f, g \in \mathbb{K}[x_1, \dots, x_n]$ is *balanced* if all the partial degrees of their product are equal, that is, if $d_1 + d'_1 = d_i + d'_i$ holds for all $2 \leq i \leq n$.

Proposition 1: Under our assumption of 1-D FFT black box, for two multivariate polynomials f, g the theoretical speedup of the algorithm in Section II-B is lower bounded

by $s/\max(s_1, \dots, s_n)$ and its cache complexity is within $O(s \frac{n+1}{L} + s(\frac{1}{s_1} + \dots + \frac{1}{s_n}))$. For fixed s and n , these lower and upper bounds are respectively maximized and minimized when the pair f, g is balanced. The second bound reaches a local minimum at $n = 2$ and $s_1 = s_2 = \sqrt{s}$.




Figure 1. Speedup of bivariate multiplication on balanced input.




Figure 2. Speedup of multiplication for non-bivariate or non-balanced input.

We present here experimental results which confirm the above analysis. Figure 1 provides speedup factors of different balanced pairs of bivariate polynomials. The number associated with each curve is the common partial degree of the input. This illustrates the good performances of the algorithm of Section II-B for such input. For the partial degree 8191, our implementation reaches a speedup factor of 14 on 16 cores. Figure 2 provides speedup factors of different pairs which are either non-bivariate or non-balanced. The performances reported there are clearly much less satisfactory than what could be observed on Figure 1. Note that these poor results are obtained on both non-balanced bivariate and balanced non-bivariate input.

D. Reduction to Balanced Bivariate Multiplication

The results of Sections III-A to III-C indicate that a reduction to bivariate multiplication with balanced input could improve the performance of multivariate multiplication based on FFT techniques. This is, indeed, possible and we achieve this reduction through the rest of the paper.

In Section IV we describe a first fundamental technique, that we call *contraction*. This generalization of Kronecker’s substitution allows us to turn a n -variate multiplication (for $n > 2$) into a bivariate multiplication without any overheads in terms of serial running time. This technique provides performance improvements on many practical cases.

In Section V, we study how univariate polynomial multiplication can be performed efficiently via bivariate multiplication based on 2-D TFT. This technique, that we call *extension*, has several motivations. First, under our assumption of 1-D FFT black-box (which may be a serial program) this trick creates concurrent execution for FFT-based univariate multiplication. Secondly, when the base field \mathbb{K} does not possess primitive roots of unity of sufficiently large orders for performing a Cooley-Tukey radix-2 FFT, this trick can reduce the computations to a case where this latter algorithm can be applied. Finally, this technique of extension, together with that of contraction studied in Section IV, is the basis of dense multivariate multiplication via *balanced bivariate multiplication*, presented in Section VI.




Figure 3. Timing of 4-variate multiplication with unbalanced input via 4-D TFT vs balanced 2-D TFT methods.

Figure 3 gives the timing of a 4-variate multiplication with unbalanced input via 4-D TFT method vs balanced 2-D TFT method. Even on 1 core, the balanced 2-D TFT method is 2.5 times faster. The work by the two methods is essentially the same. However, our balanced 2-D TFT method has better cache efficiency. Moreover, it scales well on 16 cores. As a result, the total “net speedup” using balanced 2-D TFT method instead of the direct 4-D TFT method reaches 31 on 16 cores.

IV. CONTRACTION

Before introducing the concept of contraction in Definition 3, we specify in Definition 2 how polynomials are represented in our implementation. Proposition 2 states that contraction can be performed essentially “at no cost” with this representation. The experimental results reported at the end of this section illustrate the benefits of contraction.

As in Section II, let $f, g \in \mathbb{K}[x_1, \dots, x_n]$ be multivariate polynomials with coefficients in the prime field $\mathbb{K} = \mathbb{Z}/p\mathbb{Z}$ and with ordered variables $x_1 < \dots < x_n$. For each i , let d_i and d'_i be the degree in x_i of f and g respectively.

Definition 2: Let ℓ_1, \dots, ℓ_n be positive integers such that $\ell_i > d_i$ holds for all $1 \leq i \leq n$. Define $\underline{\ell} := (\ell_1, \dots, \ell_n)$. We call $\underline{\ell}$ -recursive dense representation (RDR, for short) of f any one-dimensional array F of size $\ell := \ell_1 \cdots \ell_n$ and with integer indices in the range $0 \cdots (\ell - 1)$ such that the following two conditions hold.

- (i) the coefficient in f of the monomial $x_1^{e_1} x_2^{e_2} \cdots x_n^{e_n}$, for $0 \leq e_i \leq d_i$, is in the slot $F[j]$ of F with index $j = e_1 + \ell_1 e_2 + \ell_1 \ell_2 e_3 + \cdots + (\ell_1 \cdots \ell_{n-1}) e_n$,
- (ii) the coefficient $F[j]$ is 0 whenever the index j equals $e_1 + \ell_1 e_2 + \ell_1 \ell_2 e_3 + \cdots + (\ell_1 \cdots \ell_{n-1}) e_n$ where $0 \leq e_i < d_i$ for $1 \leq i \leq n$ and $d_i < e_i$ holds for some i ; such a coefficient $F[j]$ is called a *padding zero*.

Remark 1: When $n = 1$, an $\underline{\ell}$ -RDR of f is given by any vector F of size at least $d_1 + 1$ where $F[i]$ is the coefficient of x_1^i in f when $0 \leq i \leq d_1$ and 0 otherwise. Consider now $n > 1$ and let ℓ_1, \dots, ℓ_n be positive integers such that $\ell_i > d_i$ holds for all $1 \leq i \leq n$. For all $0 \leq i \leq d_n$, let $c_i \in \mathbb{K}[x_1, \dots, x_{n-1}]$ be the coefficient of f regarded as a univariate polynomial in x_n and C_i be an $(\ell_1, \dots, \ell_{n-1})$ -RDR of c_i . Let $Z_{d_n+1}, \dots, Z_{\ell_n-1}$ be zero-vectors, all of size $\ell_1 \cdots \ell_{n-1}$. Then an $\underline{\ell}$ -RDR of f is obtained by concatenating $C_0, C_1, \dots, C_{d_n}, Z_{d_n+1}, \dots, Z_{\ell_n-1}$ in this order. This fact justifies the term *recursive dense representation*.

Recall how the product fg can be computed in parallel via n -dimensional FFT in the context of our implementation. During *Step 1* and *Step 3* of the algorithm of Section II-B, n -dimensional FFT’s are performed by computing in parallel one-dimensional FFT’s along x_i , for $i = 1, \dots, n$ successively. As pointed out in Section III, this approach suffers from the following bottleneck. In practice (and in particular when solving systems of polynomial equations) the partial degree d_1 is likely to be large whereas d_2, \dots, d_n are likely to be as small as 1 or 2. This implies that the n -dimensional FFT approach will compute a lot of “small 1-D FFTs” (whereas such 1-D FFTs of short vectors are not worth the game) and only a few “large 1-D FFTs” concurrently (leading to poor parallelism). To deal with this “unbalanced work” the techniques developed in this paper aim at transforming the polynomials f and g into bivariate ones in a way that they can be efficiently multiplied by a

2-D FFT approach. In this section, we accomplish a first step toward this goal using the notion of *contraction*.

Definition 3: Let ℓ_1, \dots, ℓ_n be positive integers such that $\ell_i > d_i$ holds for all $1 \leq i \leq n$. Let m be an integer satisfying $1 \leq m < n$. Define $\alpha_1 = 1, \alpha_2 = \ell_1, \alpha_3 = \ell_1 \ell_2, \dots, \alpha_m = \ell_1 \ell_2 \dots \ell_{m-1}, \alpha_{m+1} = 1, \alpha_{m+2} = \ell_{m+1}, \dots, \alpha_n = \ell_{m+1} \ell_{m+2} \dots \ell_{n-1}$. Then, we set $\underline{\alpha} := (\alpha_1, \dots, \alpha_n)$. Consider the following map from the monomials of $\mathbb{K}[x_1, \dots, x_n]$ to those of $\mathbb{K}[x_1, x_{m+1}]$:

$$x_1^{e_1} x_2^{e_2} \dots x_n^{e_n} \mapsto x_1^{c_1} x_{m+1}^{c_2}$$

where $c_1 = \alpha_1 e_1 + \alpha_2 e_2 + \dots + \alpha_m e_m$ and $c_2 = \alpha_{m+1} e_{m+1} + \alpha_{m+2} e_{m+2} + \dots + \alpha_n e_n$. This induces a polynomial ring homomorphism $\Psi_{\underline{\alpha}}$, from $\mathbb{K}[x_1, \dots, x_n]$ to $\mathbb{K}[x_1, x_{m+1}]$, that we call $\underline{\alpha}$ -contraction. Hence, this map satisfies $\Psi_{\underline{\alpha}}(ab) = \Psi_{\underline{\alpha}}(a)\Psi_{\underline{\alpha}}(b)$, for all polynomials a, b .

Proposition 2: With the notations of Definition 3, define $t_1 = \ell_1 \ell_2 \dots \ell_m, t_2 = \ell_{m+1} \dots \ell_n$ and $\underline{t} = (t_1, t_2)$. Let F be an one-dimensional array of size $t_1 t_2 = \ell_1 \dots \ell_n$. If F is an ℓ -RDR of f , then F is also a \underline{t} -RDR of $\Psi_{\underline{\alpha}}(f)$. Conversely, if F is a \underline{t} -RDR of $\Psi_{\underline{\alpha}}(f)$, then it is an ℓ -RDR of f .

Proposition 2 follows easily from the *recursive structure* of RDR's, as pointed out in Remark 1. We explain now how we make use of contraction for computing the product fg . We define $\ell_i := d_i + d'_i + 1$, for all i . Then, we set $\underline{\ell} := (\ell_1, \dots, \ell_n)$. Let F and G be $\underline{\ell}$ -RDR of f and g respectively. We choose an integer m such that the “distance” given by $|\ell_1 \dots \ell_m - \ell_{m+1} \dots \ell_n|$ is minimum. With these values for $\underline{\ell}$ and m , consider the $\underline{\alpha}$ -contraction of Definition 3. Then $\Psi_{\underline{\alpha}}(f)$ and $\Psi_{\underline{\alpha}}(g)$ are two bivariate polynomials in $\mathbb{K}[x_1, x_m]$. By the choice of ℓ_i 's, the polynomial fg is the only pre-image of $\Psi_{\underline{\alpha}}(f) \times \Psi_{\underline{\alpha}}(g)$ under $\Psi_{\underline{\alpha}}$. Therefore, this map can be used to compute fg via a 2-D FFT approach. Moreover, it follows from Proposition 2 that this change of representation is made at no cost! In addition, by the choice of m , the degrees of $\Psi_{\underline{\alpha}}(fg)$ w.r.t. x_1 and x_{m+1} are as close to each other as possible.

Let us compare the work, the parallelism and the cache complexity of the multiplication based on n -D FFT with the multiplication based on contraction and 2-D FFT approach. To keep the discussion simple, let us assume that we can choose $s_i = \ell_i$ for all i . (Recall that s_i is the size of our 1-D FFT input vectors along x_i .) This is actually realistic if all 1-D FFTs are computed by TFT, which is the case in our implementation. It follows from Proposition 2 and Expression (1) that the work is unchanged. The inequality

$$\frac{3}{L} + \left(\frac{1}{s_1 \dots s_{m-1}} + \frac{1}{s_m \dots s_n} \right) \leq \frac{n+1}{L} + \left(\frac{1}{s_1} + \dots + \frac{1}{s_n} \right)$$

combined with Expression (3) suggests that contraction is likely to reduce cache misses. As discussed in Section III-A, “contracting dimensions” will keep enough theoretical speedup while reducing parallel overhead.

Experimental results. In our experimentation illustrated in Figure 4, we study the case of multiplying two 4-variate polynomials f and g . Their partial degrees d_2, d_3, d'_2, d'_3 are all equal to 1 while $d_1 = d'_1$ and $d_4 = d'_4$ vary in the range $1024 \dots 2047$. These degree patterns are typical in computing normal forms based on the algorithm in [11].




Figure 4. Timing (s) for 4-variate multiplication by direct 4-D TFT on 1 core vs Kronecker’s substitution on 1 core vs contraction from 4-D to 2-D TFT on 1 core and 16 cores.

On 1 core, we compare three methods for computing the product fg : direct 4-D TFT, 1-D TFT via Kronecker substitution (see Section II-B) and our contraction to 2-D method. We should observe first that the Kronecker substitution method fails on most input due to the fact that \mathbb{K} does not have primitive roots of unity of sufficiently large order; in those cases our contraction method clearly outperforms the direct 4-D TFT method, which was expected, based on our complexity estimates. When the Kronecker substitution method does not fail, our contraction method is clearly the most efficient technique. The fact that the contraction outperforms Kronecker’s substitution in this case can probably be explained by cache complexity arguments, see the discussion on FFT computation in [6]. The results also show that the contraction method scales very well on 16 cores, reaching a speedup factor between 8.2 to 13.2 for this large range of problems. The net speedup on 16 cores w.r.t. the direct 4-D method on 1 core is between 16 and 30.

V. EXTENSION

Let $f, g \in \mathbb{K}[x_1]$ be two non-constant univariate polynomials with coefficients in the prime field $\mathbb{K} = \mathbb{Z}/p\mathbb{Z}$ and with respective degrees d_f, d_g . Let $b \geq 2$ be an integer. We consider the map Φ_b from $\mathbb{K}[x_1]$ to $\mathbb{K}[x_1, x_2]$ that replaces any monomial x_1^a by $x_1^u x_2^v$, where u, v are the remainder and quotient in the Euclidean division of a by b , and that leaves all coefficients unchanged. More formally Φ_b is the canonical ring homomorphism from $\mathbb{K}[x_1]$ to the residue class ring $\mathbb{K}[x_1, x_2]/\langle x_1^b - x_2 \rangle$.

We determine a value for b such that the product fg can be obtained by computing $\Phi_b(f)\Phi_b(g)$ using a number

of operations in \mathbb{K} which is at most twice the number of operations for a direct computation in $\mathbb{K}[x_1]$. Moreover, we impose that the pair $\Phi_b(f), \Phi_b(g)$ is balanced, or nearly balanced. Indeed, the cache complexity upper bound given by Expression (3) is minimized in this case.

To this end, we need some notation. Let s_1, s_2 be positive integers and F, G, H be (s_1, s_2) -RDR's of $\Phi_b(f)$, $\Phi_b(g)$ and $\Phi_b(f)\Phi_b(g)$. Since the product $\Phi_b(f)\Phi_b(g)$ will be computed by means of 2-D TFTs applied to F and G , we shall determine b, s_1, s_2 such that both $s := s_1 s_2$ and $|s_1 - s_2|$ are as small as possible in order to reduce work and cache complexity, and improve parallelism as well. Let q_f, r_f (resp. q_g, r_g) be the quotient and remainder in the Euclidean division of d_f (resp. d_g) by b . Since $\Phi_b(f)\Phi_b(g)$ will be calculated by 2-D TFTs over \mathbb{K} , this polynomial product will be obtained as an element of $\mathbb{K}[x_1, x_2]$ (not one of $\mathbb{K}[x_1, x_2]/\langle x_1^b - x_2 \rangle$). Hence the degrees of $\Phi_b(f)\Phi_b(g)$ w.r.t. x_1 and x_2 will be at most $2b - 2$ and $q_f + q_g$, respectively. Thus we should set $s_1 = 2b - 1$ and $s_2 = q_f + q_g + 1$. Roughly speaking, the RDR's F, G of f and g contain at least 50% of padding zeros. More precisely, the size (i.e. number of slots) of each of the arrays F, G, H is

$$s = 2b(q_f + q_g) + (s_1 - s_2) + 1.$$

It follows from Lemma 1 hereafter that it is always possible to choose b such that the absolute value $|s_1 - s_2|$ is at most equal to 2. This implies the following: $s \leq 2(d_f + d_g) + 3$. Since the size of the univariate polynomial fg is $d_f + d_g + 1$, this b "essentially" realizes our objectives of increasing the data size at most by a factor of 2, while ensuring that our 2-D TFT's will operate on (nearly) square 2-D arrays.

Lemma 1: With $d_f, d_g, b, q_f, q_g, r_f, r_g$ as above, given a positive integer σ , define $t_1 := (2b - 1)$ and $t_2 := (q_f + q_g + 1)\sigma$. There exists at least one integer b such that we have $-1 \leq t_1 - t_2 \leq 2\sigma$. In particular, for $\sigma = 1$, the inequality $|s_1 - s_2| \leq 2$ can be achieved. If, in addition, $d_f = d_g$ is satisfied, there exists b such that $|s_1 - s_2| \leq 1$ holds.

Proof: We solve for b (as positive integer) the quadratic equation $2b^2 - \sigma b - (d_f + d_g)\sigma = 0$, which means $s_1 = s_2$. Its discriminant is $\Delta := (\sigma + 1)^2 + 8(d_f + d_g)\sigma$. Let k be the positive integer satisfying $k^2 \leq \Delta < (k + 1)^2$. We define $b_i := \frac{\sigma + k + i}{4}$. For $i = -1, 0, 1, 2$, elementary calculations bring the inequalities: $-2 \leq t_1 - t_2 \leq 2\sigma$, $-1 \leq t_1 - t_2 \leq 2\sigma$, $-1 \leq t_1 - t_2 \leq 2\sigma$ and $1 \leq t_1 - t_2 \leq 2\sigma$. For the case where $\sigma = 1$ and $d_f = d_g$, we have $|s_1 - s_2| \leq 1$ for either $b = k'$ or $b = k' + 1$ with $k'^2 \leq d_f < (k' + 1)^2$. ■

Once the bivariate product $\Phi_b(f)\Phi_b(g)$ is computed, one task remains: converting this polynomial to the univariate polynomial fg . This operation is non-trivial since x_2 stands for x_1^b meanwhile the degree of $\Phi_b(f)\Phi_b(g)$ w.r.t. x_1 can be larger than b , but at most equal to $2b - 2$. Elementary algebraic manipulations lead to the pseudo-code below which

constructs an $(d_f + d_g + 1)$ -RDR of fg from H , the (s_1, s_2) -RDR of $\Phi_b(f)\Phi_b(g)$. Recall that s_1 and s_2 have been set to $2b - 1$ and $q_f + q_g + 1$ respectively. Define $d := d_f + d_g$ and $q := q_f + q_g$. This procedure clearly runs in $O(d)$ operations in \mathbb{K} . Finally, we obtain Proposition 3.

```

for  $u := 0 \dots (b - 1)$  do  $U[u] := H[u]$ ; end do;
for  $w := 1 \dots q$  do
   $X := wb$ ;  $Z := w(2b - 1)$ ;
   $Y := (w - 1)(2b - 1) + b$ ;
  for  $u := 0 \dots (b - 2)$  do
     $U[X + u] := H[Y + u] + H[Z + u]$ ; end do;
   $U[X + (b - 1)] := H[Z + (b - 1)]$ ;
end do;
 $X := (q + 1)b$ ;  $Z := d - X$ ;
 $Y := q(2b - 1) + b$ ;
for  $u := 0 \dots Z$  do  $U[X + u] := H[Y + u]$ ; end do;

```

Proposition 3: Let $f, g \in \mathbb{K}[x_1]$ have respective positive degrees d_f, d_g . Then, one can compute from f, g a pair of bivariate polynomials $h, k \in \mathbb{K}[x_1, x_2]$ within $O(d)$ bit operations, such that the product fg can be recovered from hk within $O(d)$ operations in \mathbb{K} , and such that $s_1 s_2 \leq 2(d_f + d_g) + 3$ and $|s_1 - s_2| \leq 2$ hold, where $\deg(h, x_1) + \deg(k, k_1) < s_1$, $\deg(h, x_2) + \deg(k, k_2) < s_2$ and $d = d_f + d_g$.




Figure 5. Univariate multiplication timing (s) via extension to 2-D TFT on 1, 8, 16 cores vs direct 1-D TFT.

We stress the fact that, the construction that has led to Proposition 3 permits to efficiently multiply univariate polynomials via 2-D TFT without requiring that 1-D TFTs are computed in a parallel fashion. Moreover this strategy allows us to take advantage of FFT techniques even if \mathbb{K} does not admit primitive roots of unity of sufficiently large order for using the radix 2 Cooley Tukey Algorithm. Corollary 1 states that over any field of characteristic different from 2 our techniques compute the product of two univariate polynomials of degree d within $O(d \lg^3(d))$ coefficient operations. This is slower than Schönage-Strassen Algorithm [8, Chapter 8] which runs in $O(d \lg(d) \lg(\lg(d)))$ but for which no parallelism or cache complexity results are known.

Corollary 1: Let $f, g \in \mathbb{K}[x_1]$ be non-zero polynomials of degree $d - 1$. Then, one can compute the product fg

within $O(d \lg^3(d))$ operations in \mathbb{K} .

Proof: The number of terms (null or not) in fg is $d_f + d_g + 1$, that is, $2d - 1$. The number of terms in hk is bounded by $2(d_f + d_g) + 2$, that is, $2(2d - 2) + 2 = 4d - 2$. (Here we have used the fact that the degrees of f and g are equal, see Lemma 1.) Therefore the overhead factor between a radix 2 Cooley Tukey Algorithm (if the appropriate primitive roots exist in \mathbb{K}) and our extension method is $(4d - 2) \lg(4d - 2)$ divided by $(2d - 1) \lg(2d - 1)$, which is at most 4 for $d \geq 2$. Observe that each partial degree of h, k is in $O(\sqrt{d})$. It could happen that one of these degrees is still too large for running the radix 2 Cooley Tukey Algorithm. One can, then, apply to h, k the construction of Proposition 3 in order to extend to 4 variables. Observe that the number of consecutive calls to this construction is at most $\lg(\lg(d))$; hence the cumulative overhead is within $4^{\lg(\lg(d))} = \lg^2(d)$ yielding the result. ■

Experimental results. We compare the timings of univariate polynomial multiplications based on direct 1-D TFT and our extension method to 2D, for input degree ranging between 8126460 and 32505900. The results are reported on Figure 5. On 1 core, our extension method is slower than the direct 1-D TFT method for about 30%. This is not a surprise since we know that the extension from 1-D to 2-D can increase the size s of the product by (at most) a factor of 2. On 2 cores, the extension method provides a speedup factor between 1.5 and 1.7 with respect to the direct 1-D TFT method. On 16 cores, this gain ranges between 6.5 and 11.5. These data also show that extending univariate polynomials to balanced bivariate pairs can create substantial parallelism even when 1-D FFTs are executed serially!

VI. BALANCED MULTIPLICATION

We turn now to the question of performing multivariate polynomial multiplication efficiently via bivariate multiplication, based on 2-D TFTs applied to (nearly) balanced pairs. We call *balanced multiplication* the resulting strategy. As in Sections II and IV, let $f, g \in \mathbb{K}[x_1, \dots, x_n]$ be two multivariate polynomials with coefficients in the prime field $\mathbb{K} = \mathbb{Z}/p\mathbb{Z}$ and with ordered variables $x_1 < \dots < x_n$. For each i , let d_i and d'_i be the degree with x_i of f and g respectively.

A first approach for computing the product fg via bivariate multiplication would be to convert the polynomials f and g to univariate polynomials via Kronecker's substitution and then, to apply the techniques of *extension* developed in Section V. This would have the following severe limitation. RDR's of the images of f and g by Kronecker's substitution must have padding zeros such that the product fg can be recovered. The extension technique of Section V requires also the introduction of padding zeros. A naive combination of these two transformations would introduce far too many padding zeros. We actually checked experimentally that this approach is unsuccessful.

In this section, we develop a “short cut” which combines extension and contraction in a single transformation. In order to focus on the main ideas, we shall assume first that the input polynomials are in *Shape Lemma position*, see Definition 4. This assumption is actually a practical observation (formalized by the so-called *Shape Lemma* [1]) for the polynomials describing the symbolic solutions of polynomial systems with finitely many solutions. In Remark 2 we describe how to relax this assumption.

Definition 4: The pair $f, g \in \mathbb{K}[x_1, \dots, x_n]$ is in *Shape Lemma position* if $d_1 + 1$ and $d'_1 + 1$ exceed the products $(d_2 + 1) \cdots (d_n + 1)$ and $(d'_2 + 1) \cdots (d'_n + 1)$ respectively.

This assumption suggests to extend the variable x_1 to two variables x_1, y such that f, g can be turned via a contraction Ψ_α from $\mathbb{K}[x_1, y, x_2, \dots, x_n]$ to $\mathbb{K}[x_1, y]$ into a balanced pair of bivariate polynomials.

For an integer $b \geq 2$, we consider the map Φ_b from $\mathbb{K}[x_1, x_2, \dots, x_n]$ to $\mathbb{K}[x_1, y, x_2, \dots, x_n]$ that replaces any monomial x_1^a by $x_1^u y^v$, where u, v are the remainder and quotient in the Euclidean division of a by b , and that leaves all coefficients and other monomials unchanged. More formally Φ_b is the canonical ring homomorphism from $\mathbb{K}[x_1, x_2, \dots, x_n]$ to the residue class ring $\mathbb{K}[x_1, y, \dots, x_n] / \langle x_1^b - y \rangle$.

We shall determine b such that after contracting the variables y, x_2, \dots, x_n onto y in the polynomials $\Phi_b(f)\Phi_b(g)$, the resulting bivariate polynomials h and k form a balanced pair. The construction is similar to that of Section V. Let s_1, s_2 be positive integers and H, K be (s_1, s_2) -RDR's of h and k . Define $\sigma := (d_2 + d'_2 + 1) \cdots (d_n + d'_n + 1)$. Let q_f and q_g be the quotients in the Euclidean division by b of d_1 and d'_1 respectively. Following the reasoning of Section V, we set $s_1 = 2b - 1$ and $s_2 = (q_f + q_g + 1)\sigma$ and we aim at determining b such that both $s := s_1 s_2$ and $|s_1 - s_2|$ are as small as possible, in order to reduce work and cache complexity, and improve speedup factors. Since σ is regarded as small (comparing to d_1 and d'_1), Lemma 1 provides us with a candidate b . Our experimental results confirm that this choice achieves our goals.

Remark 2: To transform a pair f, g into a pair in Shape Lemma position within $O(s)$ bit operations, we proceed as follows. We re-order the variables such that there exists an index j satisfying $1 \leq j < n$, $(d_1 + 1) \cdots (d_j + 1) \geq (d_{j+1} + 1) \cdots (d_n + 1)$ and $(d'_1 + 1) \cdots (d'_j + 1) \geq (d'_{j+1} + 1) \cdots (d'_n + 1)$. Then, contract x_1, \dots, x_j to x_1 . In rare cases, such a variable ordering may not exist and one can use Kronecker's substitution followed by extension.

Experimental results. We study the performance of our balanced multiplication method for 4-variate polynomial input. All partial degrees $d_2, d_3, d_4, d'_2, d'_3, d'_4$ are set to 2 while d_1 and d'_1 range between 32768 and 65536. Figure 6 illustrates our experimental results. On 1 core we compare




Figure 6. 4-variate multiplication timing (s) via balanced multiplication on 1, 2, 16 cores vs Kronecker substitution to 1-D TFT.

our balanced multiplication with the one through Kronecker's substitution. The latter approach performs slightly better than ours; indeed our method has a higher algebraic complexity, even though it improves on cache complexity. On 2 cores our method reaches a speedup gain of 1.75 w.r.t Kronecker's substitution and a speedup factor of 1.96 w.r.t. itself on 1 core. On 16 cores, these maxima become 10.3 and 11.3. For comparison with the direct 4-D TFT approach, see Figure 3 in Section III.

VII. CONCLUDING REMARKS

We have presented strategies for the implementation of dense polynomial multiplication on multi-core architectures. We have focused on polynomial multiplication over finite fields based on FFT techniques since this is a fundamental operation for symbolic computation. The techniques that we have developed for this operation are highly efficient in terms of parallelism and cache complexity. Our results are both theoretical and practical. We are not aware of similar work in the context of symbolic computation.

The design of our techniques has mainly two motivations. First, we aim at supporting higher-level parallel algorithms for solving systems of non-linear equations. Therefore, our multiplication must perform efficiently in terms of serial running time, parallelism and cache complexity, on any possible input degree patterns, insisting on those which put code efficiency to challenge.

Secondly, we have integrated the specificities of 1-D FFT computations over finite fields in the context of symbolic computation with polynomials. On one hand, these 1-D FFTs are applied to vectors which are large enough such that the base field may not contain the appropriate primitive roots of unity for a radix 2 Cooley Tukey Algorithm. On the other hand, the length of these vectors is not large enough for making efficient use of parallel code for 1-D FFT.

As a consequence of these constraints, we have assumed that 1-D FFTs in our implementation could be computed by a black box program, possibly a serial one. Therefore,

we had to take advantage of the row-column algorithm for multidimensional FFT computations. Our theoretical analysis has shown that balanced bivariate multiplication, as defined in Section III-C, is a good kernel.

Based on this observation, we have developed two fundamental techniques, contraction and extension, in order to efficiently reduce any dense multivariate polynomial multiplication to this kernel.

Our experimental results demonstrate that these techniques can substantially improve performances with respect to multiplication based on a direct (and potentially unbalanced) multidimensional FFT. Moreover, they can lead to efficient parallel code for univariate multiplication, despite of our 1-D FFT black box assumption. We believe that symbolic computation software packages such as MAPLE, MAGMA, NTL can greatly benefit from our work.

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

REFERENCES

- [1] E. Becker, T. Mora, M. G. Marinari, and C. Traverso. The shape of the shape lemma. *Proc. of ISSAC'94*, pages 129–133, 1994.
- [2] R. D. Blumofe and C. E. Leiserson. Scheduling multithreaded computations by work stealing. *IEEE FOCS94*, 1994.
- [3] Cilk Arts. Cilk++. <http://www.cilk.com/>.
- [4] J. Cooley and J. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comp.*, 19:297–301, 1965.
- [5] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proc. of the IEEE*, 93(2):216–231, 2005.
- [6] M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran. Cache-oblivious algorithms. *40th Annual Symposium on Foundations of Computer Science*, pages 285–297, 1999.
- [7] M. Frigo, C. E. Leiserson, and K. H. Randall. The implementation of the cilk-5 multithreaded language. *ACM SIGPLAN*, 1998.
- [8] J. Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 1999.
- [9] J. Hoeven. The Truncated Fourier Transform and applications. *Proc. of ISSAC'04*, pages 290–296, 2004.
- [10] X. Li, M. Moreno Maza, R. Rasheed, and É. Schost. The modpn library: Bringing fast polynomial arithmetic into maple. *Proc. of MICA'08*, 2008.
- [11] X. Li, M. Moreno Maza, and É. Schost. Fast arithmetic for triangular sets: From theory to practice. *Proc. of ISSAC'07*, pages 269–276, 2007.
- [12] LinBox Project. <http://www.linalg.org/>.
- [13] MAGMA Project. <http://magma.maths.usyd.edu.au/magma/>.
- [14] P. L. Montgomery. Modular multiplication without trial division. *Math. of Comp.*, 44(170):519–521, 1985.
- [15] M. Moreno Maza and Y. Xie. FFT-based dense polynomial arithmetic on multi-cores. *Proc. of HPCS'09*, 2009.
- [16] NTL. *The Number Theory Library*. <http://www.shoup.net/ntl>.