Journal of Symbolic Computation 46 (2011) 1017–1029

Contents lists available at ScienceDirect

Journal of Symbolic Computation

journal homepage: www.elsevier.com/locate/jsc

The F5 criterion revised

Alberto Arri [a,1], John Perry [b,2]

[a] Scuola Normale Superiore di Pisa - Piazza dei Cavalieri, 7 - 56126 Pisa, Italy
[b] University of Southern Mississippi, Hattiesburg, MS, USA

ARTICLE INFO

Article history:
Received 7 February 2011
Accepted 2 May 2011
Available online 10 May 2011

Keywords:
F5
Gröbner bases
Syzygies

ABSTRACT

The purpose of this work is to generalize part of the theory behind Faugère's "F5" algorithm. This is one of the fastest known algorithms to compute a Gröbner basis of a polynomial ideal I generated by polynomials $f_1, \ldots, f_m$. A major reason for this is what Faugère called the algorithm's "new" criterion, and we call "the F5 criterion"; it provides a sufficient condition for a set of polynomials G to be a Gröbner basis. However, the F5 algorithm is difficult to grasp, and there are unresolved questions regarding its termination.

This paper introduces some new concepts that place the criterion in a more general setting: $\mathcal{S}$-Gröbner bases and primitive $\mathcal{S}$-irreducible polynomials. We use these to propose a new, simple algorithm based on a revised F5 criterion. The new concepts also enable us to remove various restrictions, such as proving termination without the requirement that $f_1, \ldots, f_m$ be a regular sequence.

1. Introduction

Since their introduction by Buchberger (1965), Gröbner bases and their computation have attracted significant attention in the computer algebra community. The best-known algorithm used to compute a Gröbner basis is the original algorithm due to Buchberger, and named after him. Its efficiency has been constantly enhanced through the years, but there remains room for improvement. Various criteria have since been introduced to detect useless computations – for example, (Buchberger, 1965,

E-mail addresses: arri@sns.it (A. Arri), john.perry@usm.edu (J. Perry).
1 Present address: Microsoft Corporation.
2 Fax: +1 601 266 5818.

0747-7171/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jsc.2011.05.004

1979; Gebauer and Möller, 1988) – but even so, the algorithm spends most of its time reducing polynomials to zero ("zero reductions").

Lazard (1983) pointed out that one can view the computation of a Gröbner basis as the reduction to row-echelon form of the Macaulay matrix of the ideal. This led to the Staggered Linear Basis algorithm of Gebauer and Möller (1986), as well as the "F4" algorithm of Faugère (1999). Möller, Mora, and Traverso exploited the relationship between zero reductions and syzygies (Möller et al., 1992), but although the algorithm they presented successfully detected many zero reductions, in practice it took too much memory and time (see Section 8 of Möller et al., 1992). Faugère (2002) combined aspects of these approaches into algorithm "F5", which for a certain class of polynomial system eliminates *all* zero reductions. This algorithm exhibits impressive performance.

By Faugère's admission, the theory behind the algorithm's new criterion, which we call *the F5 criterion*, is merely sketched, so as to leave more room for examples and an accurate description of the algorithm. The proof of the algorithm's termination and correctness were likewise only outlined. Additionally, some arguments were made under strong assumptions, such as that the input sequence $f_1, \ldots, f_m$ had only principal syzygies (such a sequence is called a *regular* sequence).

We pause a moment to consider some variants of F5. Bardet described an implementation of F5 in matrix form, where termination is ensured by manually supplying a maximal degree (Bardet, 2006). Stegers filled in some details of Faugère's proof in Stegers (2006), but stopped at two conjectures, one of which Gash later showed to be false (Gash, 2008).

The purpose of this paper is to present a simpler algorithm that illustrates the fundamental principles of F5 without sacrificing termination. We begin by defining a function $\mathcal{S}$ which is equivalent to that of Faugère, then develop a structured theory, introducing new concepts such as *primitive $\mathcal{S}$-irreducible polynomials* and *$\mathcal{S}$-Gröbner bases*. These make the study of the problem more accessible, and suggest a new version of the F5 criterion which depends neither on the regularity of the input, nor on a particular ordering on the module of syzygies.

From this theory, we develop a new, simpler algorithm. We must emphasize that the algorithm is a simple demonstration of the criterion, and not a deep treatment of how to implement a highly efficient algorithm; nevertheless, the new concepts allow us to prove correctness and termination *for any input*. Note that although some F5-style algorithms provide explicit termination mechanisms (Bardet, 2006; Gash, 2008), these mechanisms rely on previously-developed, non-F5 criteria to compute a maximal degree explicitly; by contrast, the termination criterion used here is precisely the generalized F5 criterion used to detect useless computations. Later, we show that if we know that the input is a regular sequence and we use a specific ordering on $\operatorname{Syz} \mathcal{F}$, we can avoid all the reductions to zero. We compare the results to both F5 and the Staggered Linear Basis algorithm, showing how this new algorithm differs from each.

The paper's structure is as follows. Sections 2–4 cover background material; although most of this is relatively straightforward, an important and novel contribution of the paper appears at the end of Section 4 with Proposition 14. The proof of that theorem leads to the concept of *primitive $\mathcal{S}$*-irreducible polynomials, from which we obtain in Section 5 a new characterization theorem for a Gröbner basis (Theorem 18). In Section 6, we use this characterization to formulate the new algorithm, and we prove that it terminates correctly. Section 7 compares this algorithm to the Staggered Linear Basis algorithm and F5, illustrating the differences concretely. Section 8 describes some conclusions and possible future directions.

2. Preliminaries

Let $P = k[x_1, \ldots, x_n]$ be the polynomial ring over the field k with n indeterminates, let μ be any admissible ordering on $\mathbb{T}^n$, the monoid of power products over $x_1, \ldots, x_n$: $\mathbb{T}^n = \left\{\prod_{i=1}^n x_i^{\alpha_i} \mid \alpha_i \in \mathbb{N}\right\}$.

Let P^m be the free P-module generated by $\{e_1, \ldots, e_m\}$ and let μ' be any admissible ordering on $\mathbb{T}^n_m$, the set of module terms of P^m: $\mathbb{T}^n_m = \{te_l \mid t \in \mathbb{T}, l \in \{1, \ldots, m\}\}$.

Fix $\mathcal{F} = (f_1, \ldots, f_m) \in P^m$ and let $I \subseteq P$ be the ideal generated by $\mathcal{F}$, and define $v : P^m \to I$ as the P-module homomorphism such that $v(e_i) = f_i$, and let $\operatorname{Syz} \mathcal{F} = \ker v$, so that $\operatorname{Syz} \mathcal{F}$ is the module of syzygies of $\mathcal{F}$, $\operatorname{LT}(\operatorname{Syz} \mathcal{F}) \subseteq \mathbb{T}^n_m$ is set of leading module terms of $\operatorname{Syz} \mathcal{F}$, and $\operatorname{NS}(\operatorname{Syz} \mathcal{F}) = \mathbb{T}^n_m \setminus \operatorname{LT}(\operatorname{Syz} \mathcal{F})$ is the *normal set* of the syzygies of $\mathcal{F}$.