IMPERIAL

Department of Mathematics
Imperial College London

Professor David M. Evans
Chair in Pure Mathematics
Room 661, Huxley Building
180 Queen's Gate
LONDON, SW7 2AZ
UK
T +44 (0) 20 7594 9257
E david.evans@imperial.ac.uk
W https://profiles.imperial.ac.uk/david.evans

14 July 2025.

Report on PhD Thesis of Arka Ghosh, University of Warsaw: Linear Algebra in Orbit-finite Dimension.

The above-mentioned thesis is a very substantial piece of work and I am pleased to report that it is **sufficient to grant a PhD.** The thesis is very well written and I found very few typographical errors or questionable statements. A small number of comments for the author are appended at the end of this report, but these do not need to be addressed before the thesis can be passed. The more substantive ones could be asked during the public defence. Many of the results (with the exception of those in Chapter 7) have already been presented at leading conferences in the area and have been published. However, the presentation in the thesis is often different, including more examples and sometimes with different proofs. Nevertheless, it would have been good to have had a statement in the thesis about the precise contribution of the candidate to these collaborations.

My background is in model theory, homogeneous structures and automorphism groups. I have no personal connection with the candidate or supervisors. I was previously unaware of much of the work referred to in the thesis (though I had looked at [6] in detail, due to connections with some of my own work) and very much enjoyed reading about it.

I will now give my comments on individual Chapters (references are as in the thesis; I will not generally repeat statements of results or definitions):

Chapters 1 and 2 are introductory. Chapter 1 gives some background material on orbit-finite systems of linear equations and some computational motivation for studying these (including data Petri nets). This is nicely illustrated with a few examples. Section 1.3 outlines the main results proved in the thesis and explains the collaborations involved in the work and its relationship to work already presented and published by the Author and collaborators (including the Supervisors). Section 2 provides a clear and detailed introduction to the framework of the space $Lin_R(B)$ of orbit-finite vectors, where B is an orbit-finite set and R is a commutative ring. All of this is relative to a set $\mathbb{A}$ of atoms which could in principle carry some extra structure, but which is (almost always) taken to be countably infinite and have the full group of permutations of $\mathbb{A}$ acting on it as automorphisms (the 'equality atoms' case). Again, the material here is nicely illustrated with examples and the connection with orbit-finite systems of equations is explained well. The final section reviews in more detail what is known here and what is the contribution of the work in the thesis.

Chapter 3 is essentially devoted to a proof of a single result, Theorem 3.1, stating that $Lin_R(B)$ has an orbit-finite basis when B is an orbit-finite set. The proof is elegant: the basis is identified as the characteristic functions of *tight* orbits on B, so in particular does not depend on the ring R. The result is used significantly in subsequent chapters. The Chapter also discusses counterexamples to some potential extensions of the result. Most of the material here is taken from the paper [16] jointly authored with the student's PhD supervisors.

The results in Chapter 4 are also taken from [16], though a new proof of Theorem 4.4 is given. The results concern decidability of orbit-finite systems of linear equations, given basic effectivity assumptions on arithmetic in R. The first step (Theorem 4.3) is a reduction to decidability of the existence of a finite solution (this uses results in Section 3). The next step (Theorem 4.4) is to prove the decidability of this. Interestingly, the new proof of this goes via working with the case where $\mathbb{A}$ has the additional structure of a dense linear order, and follows arguments from [6]. The Chapter concludes with some complexity results which can be read off from the proofs of the decidability results.

In Chapters 5, 6 and 7, the ring R is a subring of the reals and the focus shifts to the more general case of (orbit-finite) systems of linear inequalities and decidability issues. The material in Chapters 5 and 6 is based on the joint work [16,17,18] with the student's supervisors. In Chapter 5, it is shown (Theorem 5.4) that the general problem of existence of a solution of an orbit-finite system of linear inequalities over the integers is undecidable. The very nice proof (after an initial reduction) is given in 5.2 and uses undecidability of the reachability problem for d-counter machines. In contrast to this, the corresponding problem over the reals (or perhaps more properly, over a recursive subfield of the reals) is shown to be decidable (Theorem 5.3). One feature of the proof is the reduction to systems of polynomially parametrised systems of linear inequalities. I was initially very surprised about this, but once on sees how this comes about, it is very natural, but certainly ingenious. Here again, complexity results are given for the decidable problems, but the work required to obtain these is somewhat greater than was the case in Chapter 4. Methods involving polynomially parametrised systems are also used in Chapter 6 to prove results about computability of orbit-finite linear programming problems.

Chapter 7 continues the theme of orbit-finite linear programs. Here, the main questions concern the extent to which duality results (familiar from finite-dimensional linear programming) hold in this context. This is joint work with the student's supervisors and other collaborators; it is currently unpublished (see 1.3.1), but it is surely publishable. The main results are that weak duality holds for orbit-finite linear programs (Theorem 7.1) and strong duality holds if the specifying matrix A is row and column finite (Theorem 7.17). A number of interesting (counter-)examples are given relating to possible extensions of these results (I particularly like Example 7.16). The proof of Theorem 7.1 is a fairly straightforward reduction to finite linear-programming. The proof of Theorem 7.17 is much more difficult and involves ideas developed in previous Chapters, together with some additional techniques.

The thesis concludes with a very good list of interesting open questions. Many of these concern extending the work of the thesis from the case of equality atoms to more general omega-categorical $\mathbb{A}$ (in some cases imposing some additional model-theoretic conditions on $\mathbb{A}$). Examples indicate that this is not going to be routine, although undoubtedly some of the methods from the thesis can be extended and some ideas from model theory, which were not so prominent in the case of equality atoms, will be useful. There are also questions about the equality atoms case which remain to be investigated.

In summary, I am very pleased to confirm that this is a strong, well written thesis which contains many original and interesting results, together with promising ideas for future investigation. It fully merits the award of a PhD. The UK does not categorise PhD's and in particular, there is no equivalent of the 'magna cum laude' designation. I therefore have no calibration by which to judge whether the thesis should be awarded such a distinction and therefore would defer to the other referees in this matter.

Yours,

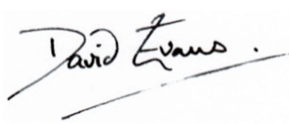

David M Evans, Chair in Pure Mathematics, Imperial College London, UK.

Some comments for the Author:

page 8, footnote 5: Do you mean this, or is there a typo?

page 84, for example Theorem 5.3: What does 'decidable' mean in the context of $\mathbb{R}$?

page 87, top of page, proof of (Ineq $\to$ Nonneg-eq): I do not understand what the system is here.

Corollary 5.27: in your proof of 5.25 you assume that T is the support of $\mathbf{x}$. So I think you need to explain in more detail how the Corollary follows.

Corollary 6.10: Is this true without restrictions on where the entries of A, b, c lie?

Remarks 7.10: do you have an opinion on this (especially in view of Example 7.16)?

In Definition 7.26, K is a finite superset of S. Immediately afterwards in Definition 7.28, it is an S-orbit. This is not a good choice of notation.