

CoCoALib - Slug #842

PPMonoidSparse: comparisons are VERY SLOW

15 Feb 2016 13:54 - John Abbott

| | | | |
|---|--------------|------------------------|--------------------|
| Status: | New | Start date: | 15 Feb 2016 |
| Priority: | Normal | Due date: | |
| Assignee: | | % Done: | 0% |
| Category: | Improving | Estimated time: | 0.00 hour |
| Target version: | CoCoALib-1.0 | Spent time: | 1.00 hour |
| Description | | | |
| I have just tried with Mario to use PPMonoidSparse rather than a normal dense repr (for a PPM with about 100 indets). The sparse impl was about 200 times slower than the dense one. | | | |
| Profiling showed that PPMonoidSparse::myCmpOrdvs was the culprit. | | | |
| Related issues: | | | |
| Related to CoCoALib - Feature #800: PPMonoidSparse: impl of sparse PPs | | Closed | 09 Nov 2015 |

History

#1 - 15 Feb 2016 13:54 - John Abbott

- Related to Feature #800: PPMonoidSparse: impl of sparse PPs added

#2 - 15 Feb 2016 13:57 - John Abbott

I have created a temporary hack to the code for the special case of lex ordering. Essentially I have written a new fn myCmpLex; this largely solved the problem of low speed (not as well as I had hoped, but still good enough for Mario to proceed).

PPMonoidSparse::myCmpOrdvs is very "clean" but also understandably very slow (e.g. it always performs a matrix-vector product using BigInt arithmetic).

We should try to make the code work better for the most common cases.