A Detailed Accounting of GENIE Cross-section Memory Usage in the NUMI Software Efforts

Robert Hatcher* [,1]

[1] *Fermi National Accelerator Laboratory*

September 27, 2010

1 Introduction

This note describes use of cross-section splines in NUMI experiments with an emphasis on the memory footprint they present while in use during `GENIE`[1][2][3] event generation, the disk space requirements and I/O demands.

The use of of `GENIE` for event generation requires the evaluation of numerous cross-sections. At the most basic level cross-sections are multi-dimensional (e.g. E_ν, $x_{\rm bj}$, y, ...) differential functions. The probability of a neutrino interaction of a particular neutrino involves only the cross-section integrated over all the variables except neutrino energy. Thus for efficiency reasons it is useful to precalculate the cross-sections for each physics process as a function of neutrino energy and only evaluate the differential cross-sections during the selection of kinematic variables of individual events. A representation of the 1-dimensional cross-sections is written to a file which is read at the beginning of event generation jobs. Separate functions are kept for each physics sub-process in order to efficiently select a process in the right proportions.

2 Numbers

The `R-2_6_0` release of `GENIE` with the selection of the default physics processes requires 577 individual sub-process cross-sections for every isotope assuming the full six neutrino flavors are under consideration. These are built up by applying nuclear effects to combinations of of 271/301 cross-sections of free protons and free neutrons.

In memory these cross-sections are represented by `genie::Spline` objects. Each of which is composed in part by a ROOT `TSpline3` object which holds the individual knot's information and makes the evaluation at arbitrary energies.

*rhatcher@fnal.gov

The energy of the NUMI beamline limits neutrinos to below 120 GeV. We have chosen to calculate them in the range 10 MeV to 200 GeV, with 500 knots spaced logarithmically to increase the sampling at low energies where the cross-sections vary more rapidly.

Initially cross-section splines were generated for 30 elements (90 isotopes) listed in Table 1. From the table one finds that Helium, Boron, Flourine, Cobalt, Zinc, and Bromine do not show up in any of the current NUMI geometry representations. It is not quite clear why these were included in the list to be generated; some might be due to anticipation of trace elements in material descriptions, others due to conflating the list of isotopes necessary in the PDG properties table that are needed because they are an end-product of an charged current interaction. It is also apparent that the geometry descriptions do not attempt (in general) to describe the isotopic content. This can lead to odd cases (cf. Copper) where the average atomic mass is not a isotope physically found in nature but must also be calculated.

3 Spline Files

The splines are generated with the `GENIE` tool `gmkspl` and written to XML files. The `GENIE` tool `gspladd` can be used to combine XML spline files to create a new file of the union of subprocesses. The size of the `R-2_6_0` file generated for NUMI experiments was 1.6 GB. With the XML parser built into `GENIE` they take roughly 166 seconds to read on a typical node. This all occurs at startup time for the event generation process.

Modificatons to a private instance of the `GENIE` code enabled the ability to store the cross-sections in ROOT files. This reduced the size of the file to 922 MB, and the read time to 90 seconds. This was an performance improvement at the cost of additional complexity and the inclusion into the `GENIE` codebase will need to be discussed with the lead developers.

4 Memory Footprint

The repository of cross-sections in memory is a `genie::XSecSplineList` singleton which is primarily a STL `map` between a STL `string` ("key") and a `genie::Spline`. Keys encode physics process names, parameter settings, and sub-process information. Typical names looks like:

```
genie::AhrensNCELPXSec/Default/nu:-12;tgt:1000000010;N:2112;proc:Weak[NC],QES;
genie::ReinSeghalRESPXSec/Default/nu:16;tgt:1000822080;N:2212;proc:Weak[NC],RES;res:8;
genie::QPMDISPXSec/Default/nu:16;tgt:1000822080;N:2112;q:1(v);proc:Weak[NC],DIS;
```

Isotopes are encoded as integers in the PDG2006 schema by the formula `10LZZZAAAI` where `ZZZ` is the atomic number (n_{proton}), `AAA` is the atomic mass ($n_{\text{proton}} + n_{\text{neutron}}$), `I` is the isomer number (ground state `I` $= 0$), and `L` is n_Λ (n_{strange}). Thus $Fe^5 6$ is `1000260560`.

The size of the `genie::Spline` is:

52	`sizeof(genie::Spline)`
14	"`genie::Spline`" (`genie::Spline.fName`)
132	`sizeof(TSpline3)`
8	"`Spline3`" (`TSpline3.fName`)
5	"`spl3`" (`TSpline3.fTitle`)
$n_{\text{knots}} \times 52$	`sizeof(TSplinePoly3)`

The generated splines have 500 knots. The average key length is 88 characters. So each spline takes approximately 26300 bytes in memory. Having 51925 subprocesses leads to a total of 1.3 GB in memory. This is a significant fraction of the memory allocated to processes on the nodes (generally the assumption is 2 GB per processor). Earlier estimates of the memory footprint grossly underestimated this size which lead to the author wrongly attributing the observed memory usage to memory leakage in the XML parsing stage. Apparently what wasn't considered was that the spline's knot information records three expansion coefficients as well as the abscissa and ordinate. The ROOT `TSplinePoly3` object also derives from `TObject` which introduces additional overhead.

To reduce the total usage one can:

- Reduce the number of knots by eliminating those above the NUMI energy. This extension of the energy range was done with non-beam sources in mind. Due to the use of logarithmic spacing it contributes little (only 26 (5%) of the knots are above 120 GeV).
- Reduce the total number of knots. 500 is probably overkill, but a study would need to be undertaken to investigate the effect.
- Remove unnecessary isotopes. This is the most straightforward approach. Bookkeeping is critical as `GENIE` event generation generates a fatal C++ `assert` if it encounters an isotope missing from the `genie::XSecSplineList`, so no isotope can be overlooked. Any gains made by this approach would be lost if `GENIE` were to be enhanced with ability to choose the isotopic content based on natural abundances rather than relying on them being encoded in the ROOT geometry.

Using the last approach the number of isotopes can currently be reduced from 90 to 27. A simple `awk` script can take the original XML file and weed out the undesired isotopes. It is probably best to continue to create the more complete list and then tailor it down in case geometries change and some of the currently unused isotopes are needed.

5 Conclusions

The cross-section XML spline files currently distributed to the NUMI experiments for the `R-2_6_0` release are overly large and costly in memory and startup time. By limiting the number of isotopes covered, the XML file reduces in size from 1.6 GB to 494 MB and the read time drops from 166 seconds to 50. The in-memory size is comparably abated.

References

[1] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.

[2] et. al. Costas Andreopoulos. The GENIE Neutrino Monte Carlo. `http://projects.hepforge.org/genie/`.

[3] et. al. Costas Andreopoulos. The GENIE Neutrino Monte Carlo Generator Physics & User Manual. `http://projects-docdb.fnal.gov:8080/cgi-bin/ShowDocument?docid=753`.

Table 1: Isotopes potentially used in NUMI experiments. Particular isotopes in found in ROOT geometries are: [a]MINOS, [b]ArgoNeuT, [c]MINERVA, [d]NOνA, [e]MicroBooNe, [f]LBNE. Entries with "–" for the abundance fraction are unphysical isotopes corresponding to the average atomic mass of the element or could arise from rounding. Entries enclosed with "[]" are isotopes that have non-zero abundances but were not calculated.

Element	Symbol	Z	Natural Abundance A (%)
Hydrogen	H	1	1(99.99%)^{a,b,c,d,e,f} 2(0.01%)
Helium	He	2	[3(0.000137%)] 4(100%)
Boron	B	5	10(19.90%) 11(80.10%)
Carbon	C	6	12(98.93%)^{a,b,c,d,e,f} 13(1.07%)
Nitrogen	N	7	14(99.63%)^{a,b,c,d,e} 15(0.37%)
Oxygen	O	8	16(99.76%)^{a,b,c,d,e,f} 17(0.04%) 18(0.20%)
Fluorine	F	9	19(100%)
Sodium	Na	11	23(100%)^{a,d,e,f}
Magnesium	Mg	12	24(78.99%)^{a} 25(10.00%) 26(11.01%)
Aluminum	Al	13	27(100%)^{a,c,d,e,f}
Silicon	Si	14	28(92.23%)^{a,b,d,e,f} 29(4.68%) 30(3.09%)
Phosphorus	P	15	31(100%)^{a}
Sulfer	S	16	32(94.93%)^{a,d} 33(0.76%) 34(4.29%) [36(0.02%)]
Chlorine	Cl	17	35(75.78%)^{d} 36(–) 37(24.22%)
Argon	Ar	18	36(0.34%) 38(0.06%) 39(–) 40(99.60%)^{b,d,e,f}
Potasssium	K	19	39(93.26%)^{a,d,e,f} 40(0.01%) 41(6.73%)
Calcium	Ca	20	40(96.94%)^{a,d,f} 41(–) 42(0.65%) 43(0.14%) 44(2.09%) [46(0.004%)] [48(0.19%)]
Titanium	Ti	22	46(8.25%) 47(7.44%) 48(73.72%)^{a,b,c,d,e,f} 49(5.41%) 50(5.18%)
Vanadium	V	23	50(0.26%) 51(99.75%)^{a}
Chromium	Cr	24	50(4.34%) 51(–) 52(83.79%)^{a,b,e} 53(9.50%) 54(2.37%)
Manganese	Mn	25	55(100%)^{a}
Iron	Fe	26	54(5.84%)^{a} 56(91.75%)^{a,b,c,d,e,f} 57(2.12%)^{a} 58(0.28%)^{a}
Cobalt	Co	27	59(100%)
Nickel	Ni	28	58(68.08%) 59(–)^{a,b,e} 60(26.22%) 61(1.14%) 62(3.63%) 64(0.93%)
Copper	Cu	29	63(69.17%) 64(–)^{a} 65(30.83%)
Zinc	Zn	30	64(48.63%) 65(–) 66(27.90%) 67(4.10%) 68(18.75%) 70(0.62%)
Bromine	Br	35	79(50.69%) 80(–) 81(49.31%)
Tin	Sn	50	[112(0.97%)] [114(0.66%)] [115(0.34%)] 116(14.54%) 117(7.68%) 118(24.22%) 119(8.59%)^{a} 120(32.58%) [122(4.63%)] [124(5.79%)]
Barium	Ba	56	[130(0.11%)] [132(0.10%)] 134(2.42%) 135(6.59%) 136(7.85%) 137(7.85%)^{d} 138(71.70%) $\Sigma \neq 100\%$
Lead	Pb	82	204(1.40%) 206(24.10%) 207(22.10%)^{a,c} 208(52.40%)