Interactive Programming for LLVM TableGen

David Spickett, Staff Software Engineer, Arm

The Problem

- TableGen powers large parts of LLVM.
- You'll need to learn it eventually.
- Existing TableGen is hard to learn from.
- Tutorials are too basic or too detailed.
  - https://llvm.org/docs/TableGen/
  - https://llvm.org/docs/TableGen/ProgRef.html

Existing Solutions

- REPLs - Python, cling, lisp, etc.
- Online compilers - https://godbolt.org/
- Explainer tools - https://explainshell.com/
- `clang -ast-dump` and similar.
- Jupyter Notebooks

Common themes:

- Learn what you want
- When you want

Jupyter Notebooks

- Text and code “cells”
- Edit and re-run cells
- Results shown inline

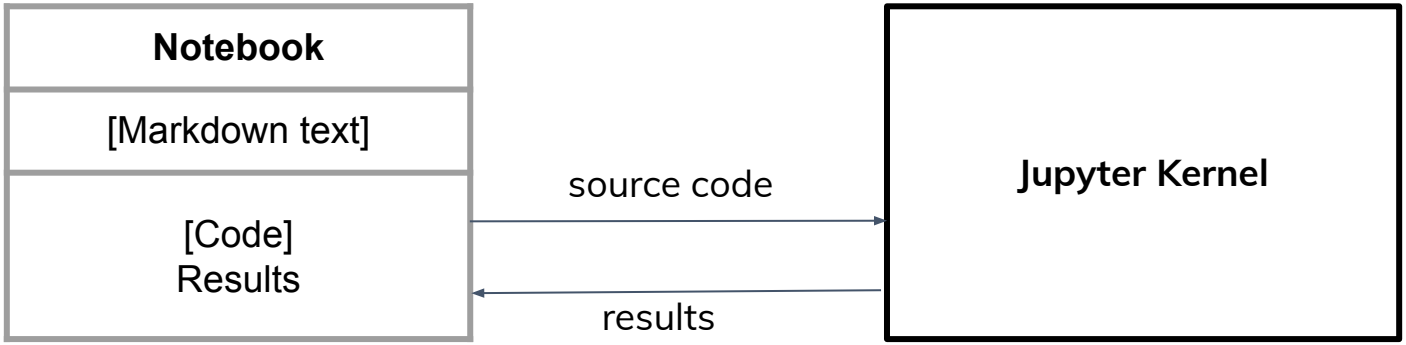

- Single .ipynb file
- Render to static formats like HTML

Jupyter Kernel for TableGen

This notebook is running `llvm-tblgen`.

In [1]:
```
%reset
// This is some tablegen
class Foo {}
```

```
------------ Classes -----------------
class Foo {
}
------------ Defs -----------------
```

- Based on the existing MLIR kernel
- Compiles with `llvm-tblgen`
- Cells linked by default (reset cache with `%reset`)
- Set compiler arguments with `%args`

The Goal

- An interactive, editable, TableGen tutorial.
- Read in Jupyter or as a static document.

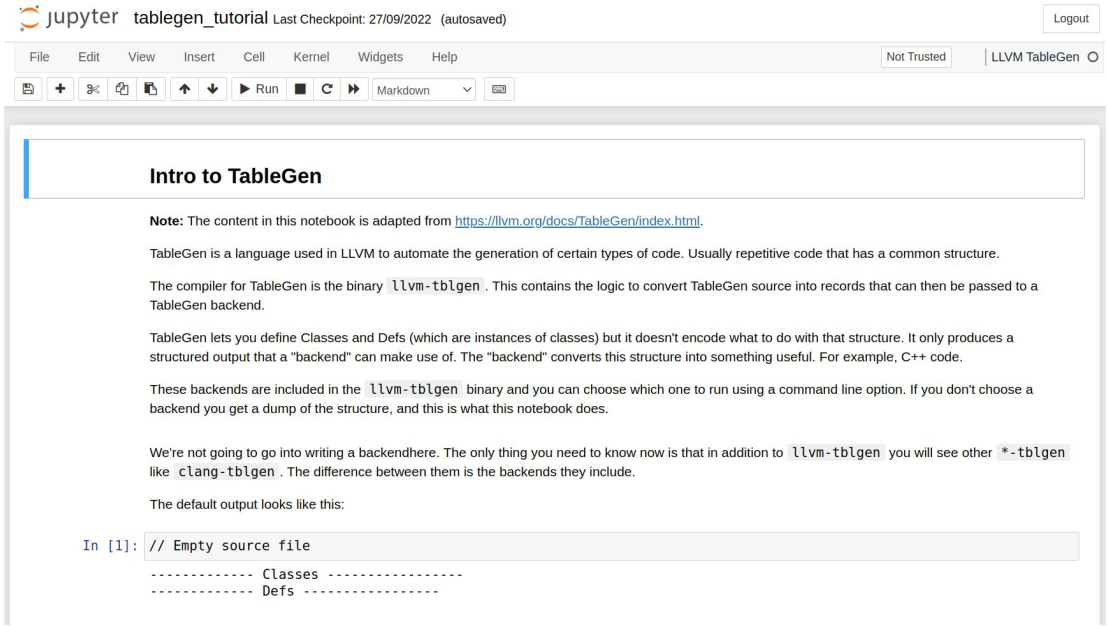

Status

- RFC: https://discourse.llvm.org/t/rfc-a-jupyter-kernel-for-tablegen/65003
- Patch series - https://reviews.llvm.org/D132378
- First tutorial notebook - https://reviews.llvm.org/D137085

Future work:

- Domain specific tutorials
- Visualise class structure
- Output filtering

Thank you

(thanks to Jacques Pienaar for the MLIR kernel)