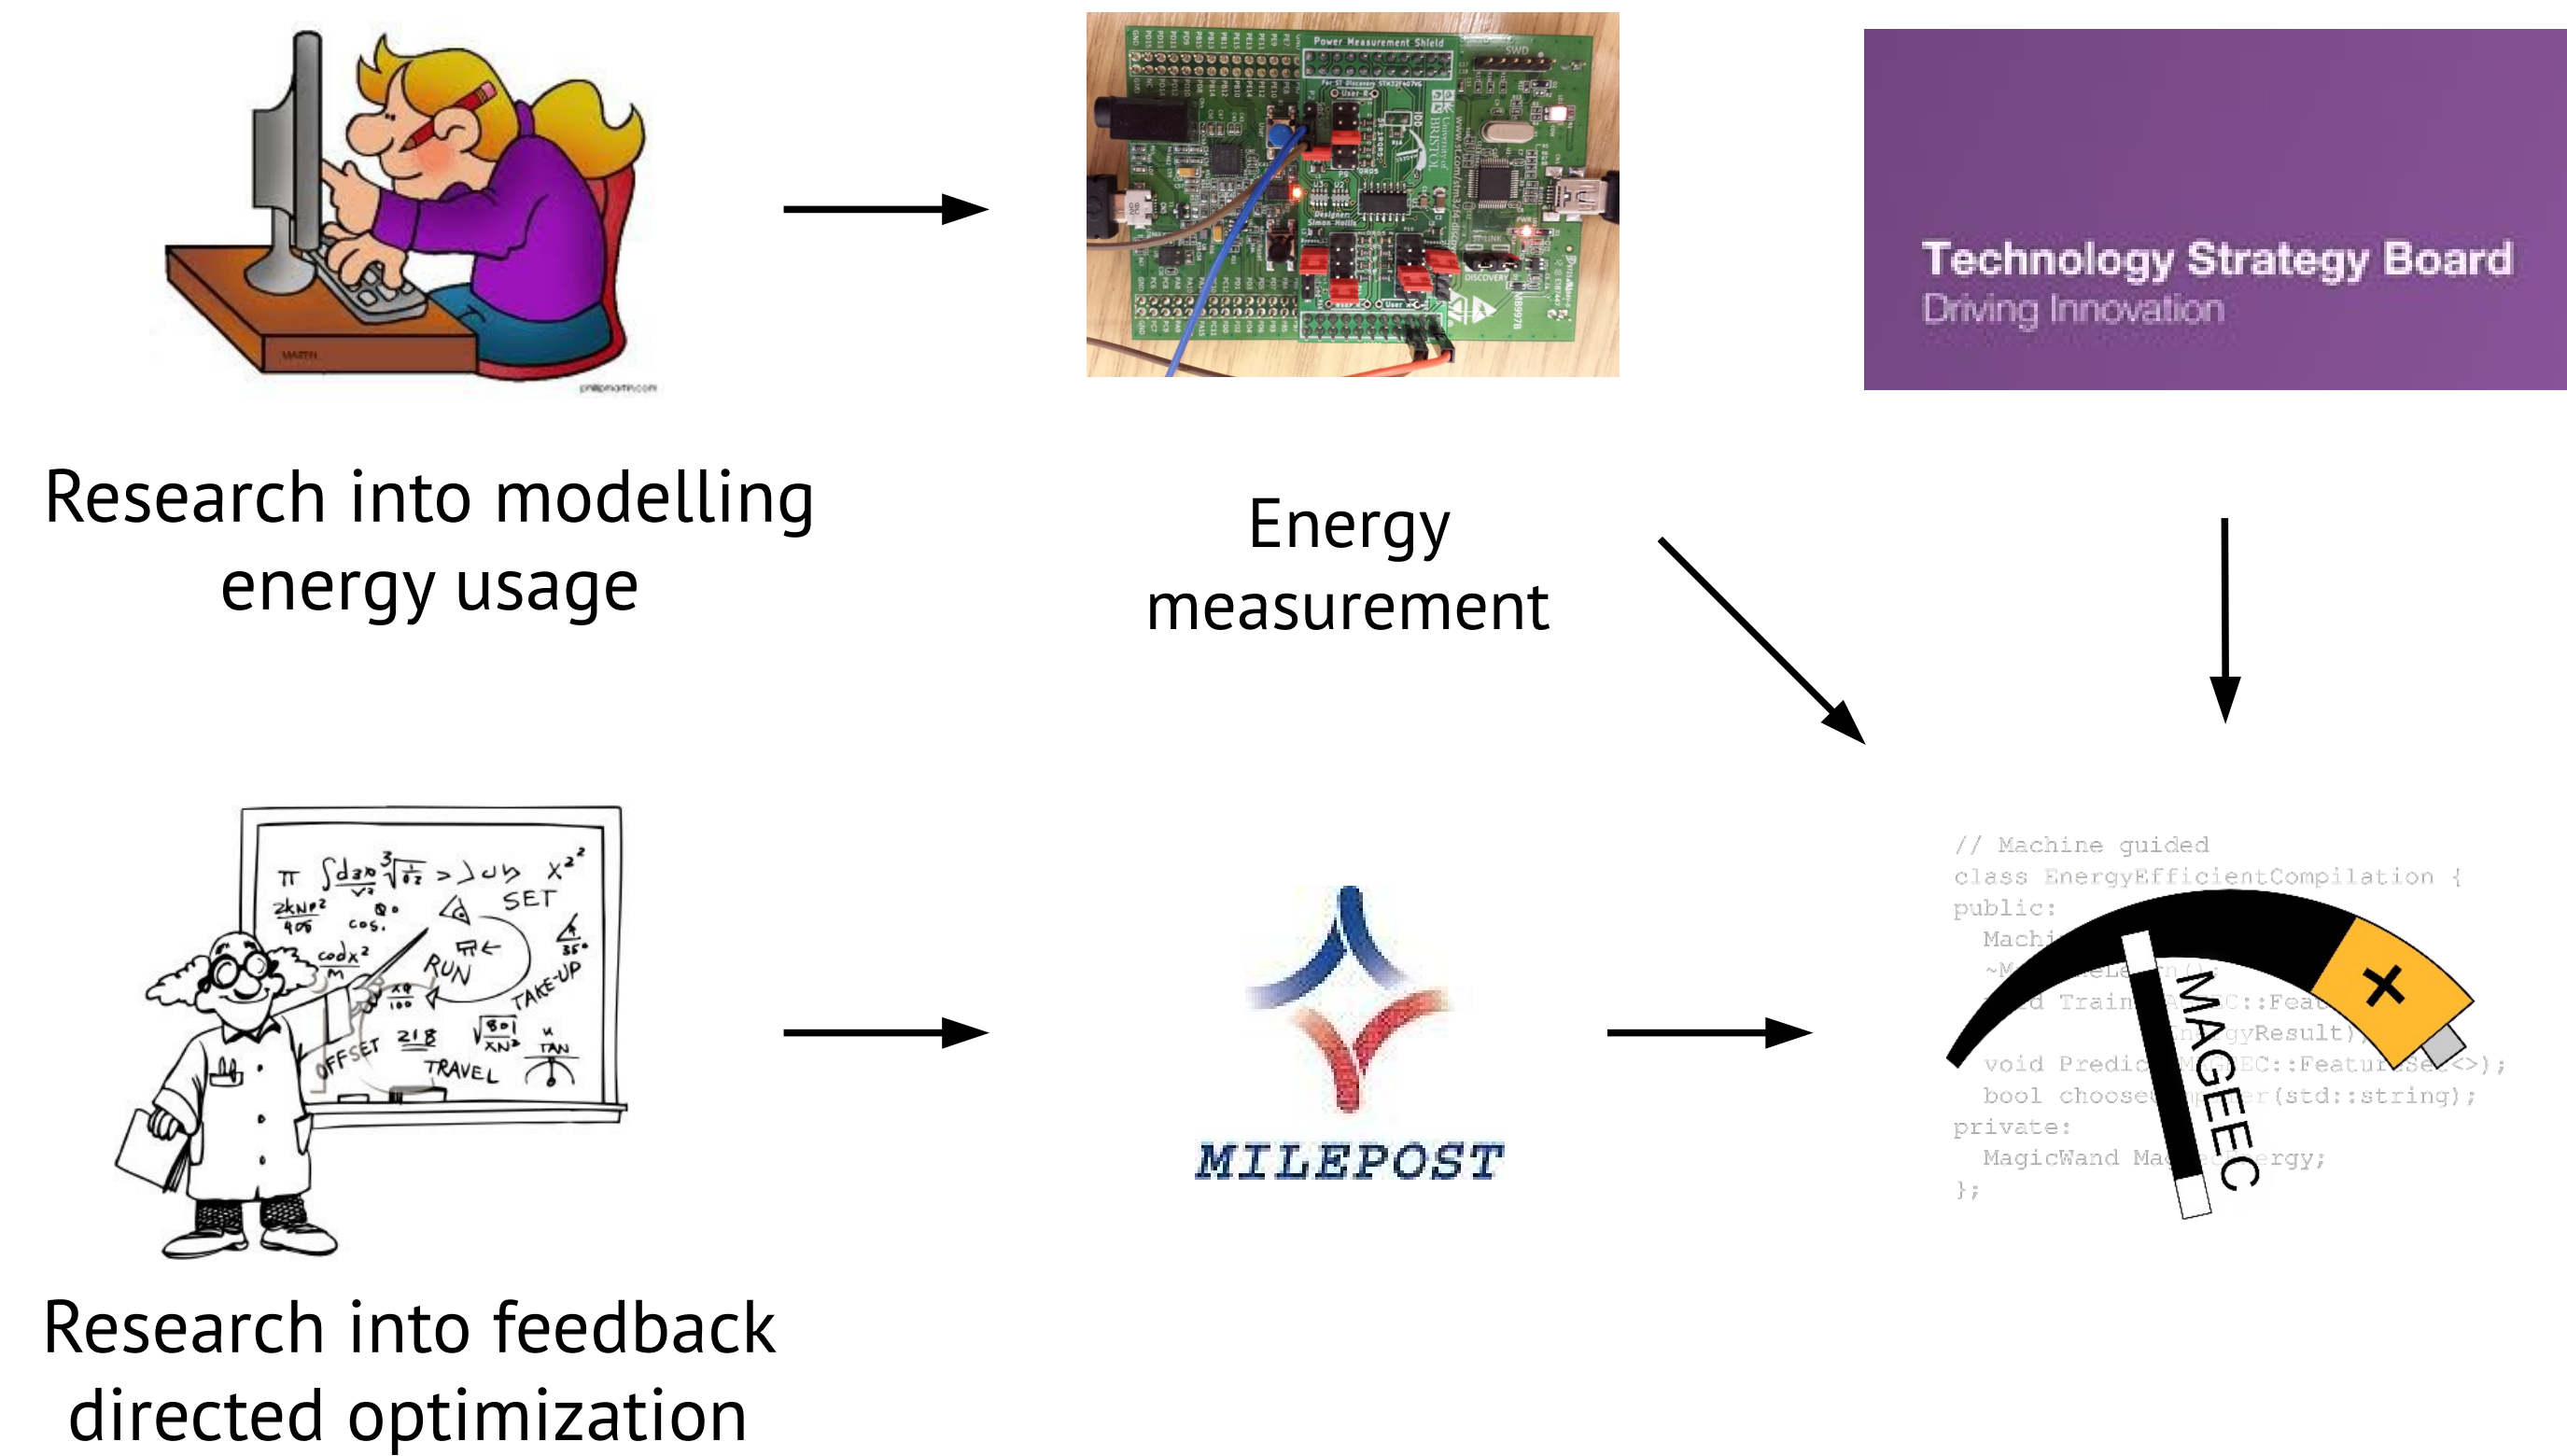


```
// Machine guided
class EnergyEfficientCompilation {
public:
    Mach
    ~Mach
    ~Mach Train
    ~Mach Train C::Feat
    ~Mach Train C::Feat EnergyResult)
    void Predict (EC::Feat use<<>);
    bool choose (er(std::string);
private:
    MagicWand Ma
    Energy;
};
```


# Machine Guided Energy Efficient Compilation

Using machine learning to select compiler optimizations that minimize energy consumption

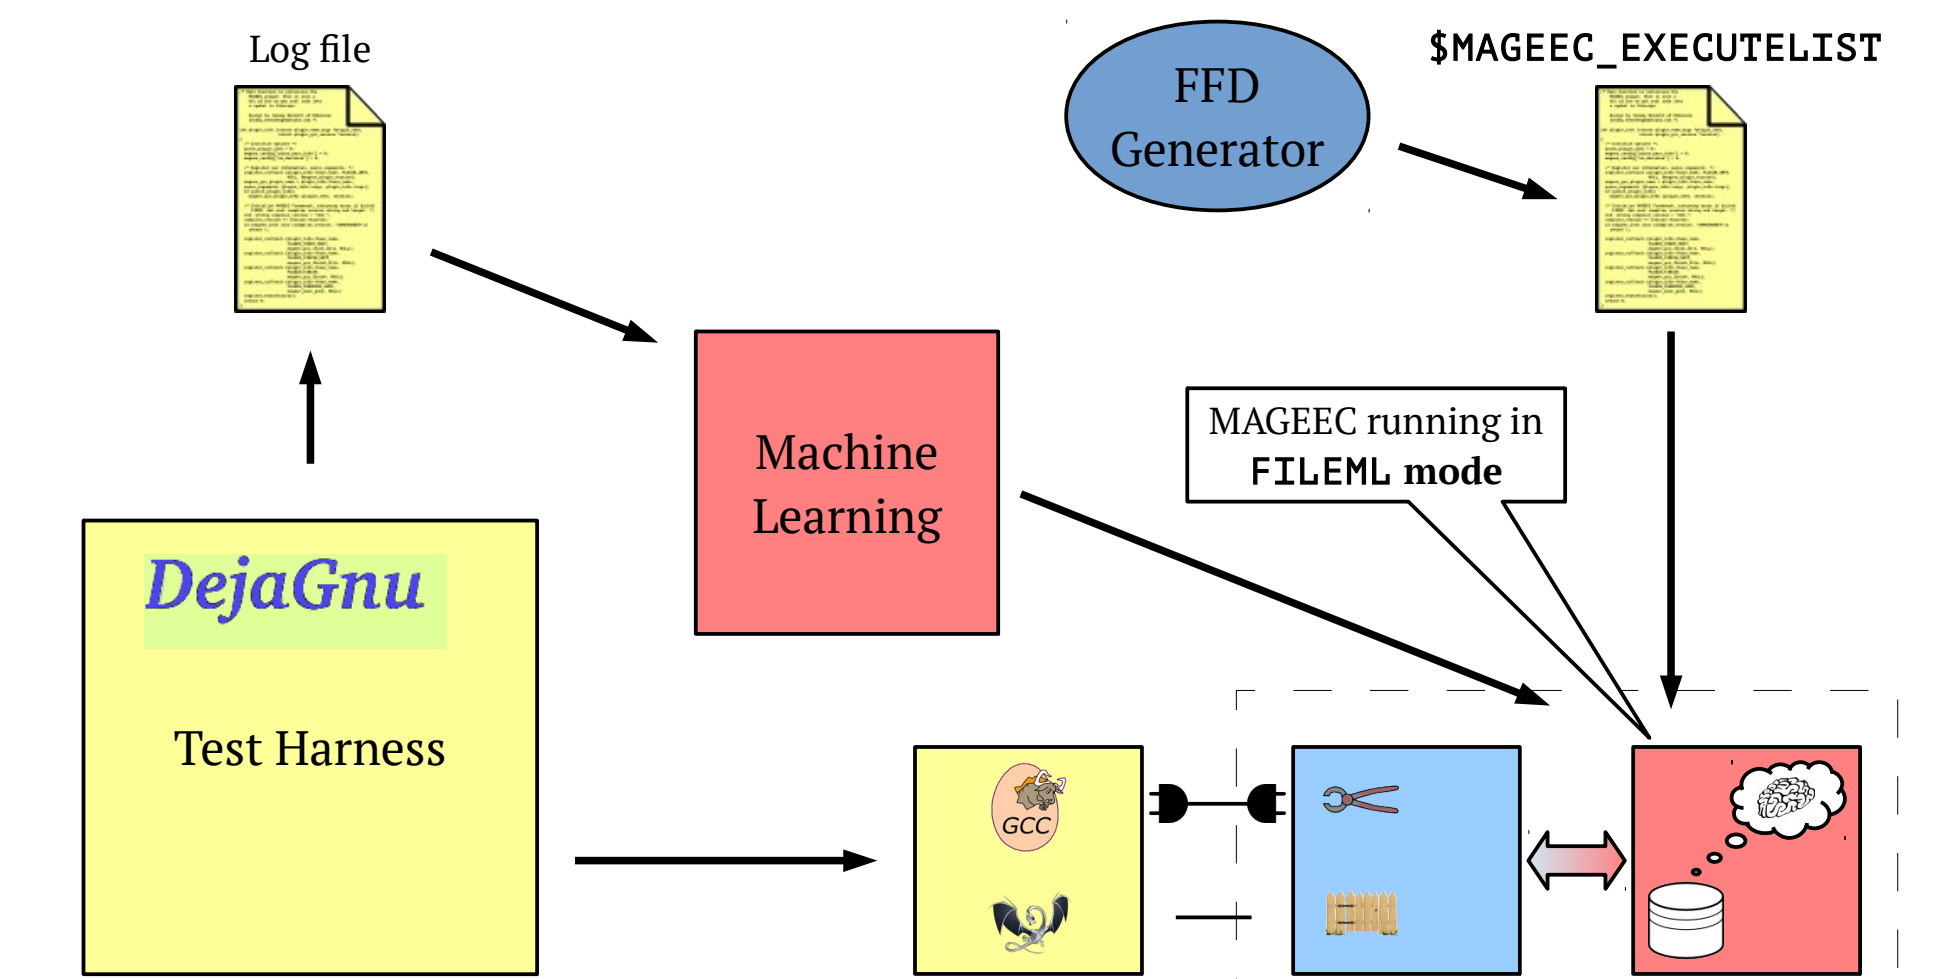
## How We Got Here



## Energy Measurement Hardware



## Training for LLVM



MAGEEC is trained using BEEBS ([www.beebs.eu](http://www.beebs.eu)), an open source benchmark suite optimized for deeply embedded devices.

To train we use a “clang-like” driver which uses *opt* for fine grained control over passes. We record file flags, passes and a feature vector which identify the program. After results have been gathered and normalized, we provide the machine learner with the “best” pass combination for each program.

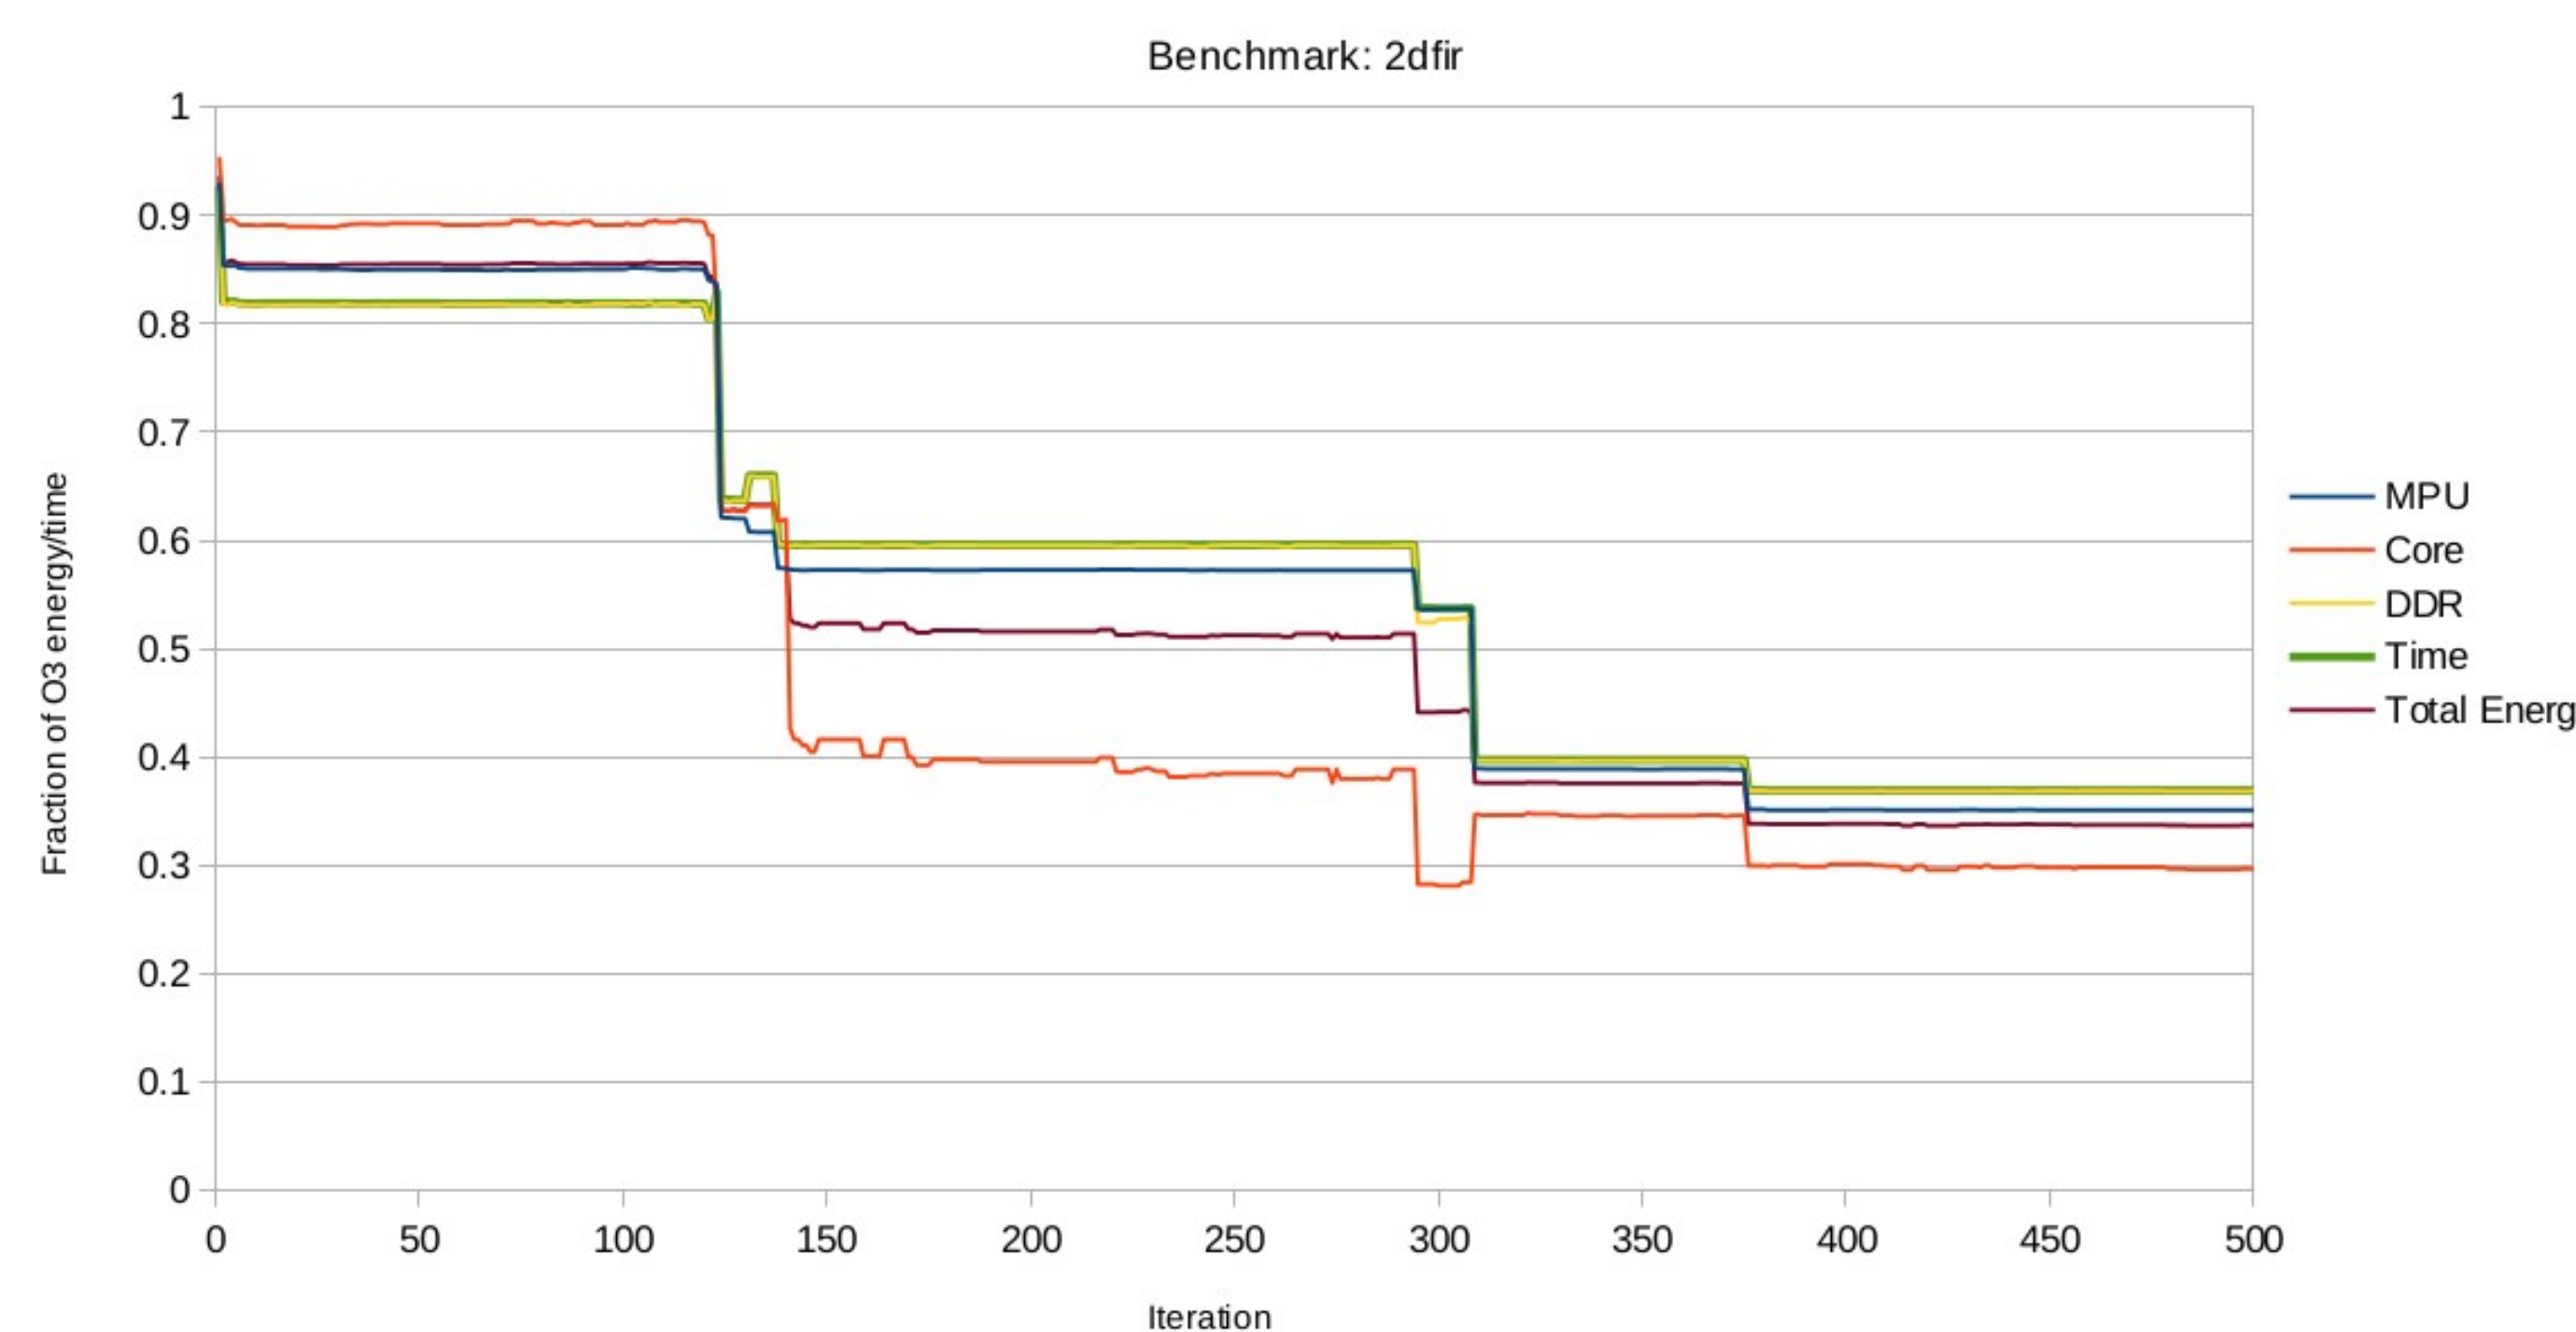
At run-time we switch to clang and use the machine learner to decide which passes should run to produce the best executable.

## Effects of Selecting Passes

Benchmark	% improvement	
	Iterative elimination	Modified version
2dfir	8.3	8.3
blowfish	1.5	1.5
crc32	0.7	1.0
fdct	11.0	15.3
float matmult	1.9	2.1
int matmult	4.6	4.7
sha	0.3	0.3

In this example, iterative compilation was used to select the optimal passes that should be used to compile a program. This shows that compared to -O3 improvements of up to 15% can be made.

## Effects of Reordering Passes



In this example, genetic algorithms were used to select optimizations to run, with the aim of minimizing energy. Over 500 generations energy was reduced by 65% with a time reduction of 60% compared to clang's -O3.

## Future Work

- Usability – having an “-Oe” flag.
- Improve MAGEEC's ability to live outside the compiler, which will require development of a plugin API for LLVM.
- Add knowledge of pass dependencies to training flow.
- In addition to disabling passes, allow pass reordering and re-execution.
- Experiment with other machine learners and data sets. It is not clear that the current choice (decision trees) is best for this problem.
- Explore other possible (multi-objective) optimization criteria. Candidates include code size, execution speed, run-time memory usage, build speed and build energy.



[mageec.org](http://mageec.org)

[github.com/mageec](https://github.com/mageec)

[beebs.eu](http://beebs.eu)

