

Benchmarking GNU Radio Kernels and Multi-Processor Scheduling

Nathan West, Doug Geiger, George Scheets

January 14, 2013

1 Introduction

The growth of Software Defined Radio (SDR) using general purpose processors (GPPs) brings a new engineering decision of processor selection to radio design. In this study we are concerned with comparing the performance of an SDR toolkit called GNU Radio on different processors. Properly selecting a processor for a radio application will depend on the application; however, a generic list of benchmarks would include

- Floating Point Operations per second (FLOPs) for common routines on multiple-processors
- time to complete common math/type conversions
- system latency

Knowledge of an application could then be paired with benchmark comparisons of potential processors to make an informed decision on the best processor for size, weight, power or cost constrained applications.

This study focuses on the first two items and leaves latency measurements as another project. [7] worked on latency measurements between an earlier version of GNU Radio and an Ettus USRP with a USB connection. Tallying FLOPs for FIR and FFT blocks in series and parallel covers two commonly used processes and tests the ability of a processor to multitask parallel operations and the ability of a processor to use multiple threads [3]. The second benchmark is common math and type conversions. GNU Radio provides the Vector Optimized Library of Kernels (VOLK) library that selects the best Single Instruction Multiple Data (SIMD) architecture for a given processor and operation to speed up computation [10, 12]. An example of this might be in a dot product, a common signal processing operation, which requires a point-by-point multiplication of two vectors followed by the sum of the products. In many cases, taking advantage of a SIMD architecture yields faster processing time than the equivalent C-style loop [10, 14, 13]. Examples of SIMD architectures are the SSE instructions common in Intel and AMD platforms, while on the ARM architecture the NEON instruction set provides a similar set of SIMD instructions. VOLK, which has been

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 14 JAN 2013	2. REPORT TYPE	3. DATES COVERED 00-00-2013 to 00-00-2013	
4. TITLE AND SUBTITLE Benchmarking GNU Radio Kernels and Multi-Processor Scheduling		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Center for Applied Research in Artificial Intelligence, 4555 Overlook Ave., SW, Washington, DC, 20375		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			
13. SUPPLEMENTARY NOTES New England Workshop for Software Defined Radio (NEWSDR'13), Worcester, MA, 17 May 2013.			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)
			18. NUMBER OF PAGES 23
			19a. NAME OF RESPONSIBLE PERSON

included in GNU Radio since December 2010, attempts to easily take advantage of SIMD instruction sets on all processors without having the application programmer be concerned about using SIMD instructions [10, 11, 12]. The first attempt to benchmark VOLK performance, in February 2012, released tools to time specific math and type-conversion kernels [11]. The current work combines the efforts to benchmark VOLK and multiprocessor scheduling of FFTs and FIRs to a single suite with the intention of comparing potential processors.

Since the interest is in comparing hardware performance, a build framework, called Open Embedded (OE), is used to install Linux with GNU Radio on the different machines. OE is a collection of tools and metadata that can cross-compile a complete Linux system with any applications pre-installed [15]. The OE metadata is separated into layers based on the use and maintainability of the software being built [4]. Most of the required system tools are hosted in a layer called *openembedded-core* (*oe-core*); the kernel, and machine description for real hardware comes from a board support package layer; the applications and development tools come from *meta-openembedded* [4]. There are other layers that describe the file system layout, software to be installed and corresponding versions, and a build system that comes from distribution layers such as Angstrom or Poky [15, 9].

2 Methodology

This study uses a range of hardware to test the effectiveness of benchmarking. The processors tested so far include

- Intel i7 with hyper threading
- Intel Atom with hyper threading
- AMD E350 APU, comparable to Atom
- ARM Cortex A8 running on a Gumstix Overo on an Ettus USRP E110

The general testing procedure consists of

- Build Linux, GNU Radio, and file system for target machine
- Run benchmarking scripts with VOLK enabled
- Run benchmarking scripts with VOLK disabled
- Optionally run oprofile with desired application

The results from each type of benchmark can be compared with the application profile to select the appropriate processor.

2.1 Open Embedded

In order to compare the differences in hardware, the software should be as controlled as reasonably possible. Using OE to build the Linux kernels and GNU Radio provides a reasonably fair platform to run benchmarking tools from [8]. OE is also a reasonable choice for creating the benchmark system since the main interest is in size, weight, and power constrained systems, which will generally be embedded. We use the Poky distribution that is included via the *meta-yocto* layer of OE, and define our custom image which is based on *core-image-minimal*. Our new image installs GNU Radio, which is provided in *meta-openembedded*, and brings in the integrated benchmarking suite.

2.2 Existing Benchmarking Code

We modified the existing benchmarking tools to store data in a consistent format that is text-based to avoid dependencies on databases and make results more portable. The plotting tools for the multi-processor scheduler benchmarking (FFT/FIR arrays) were changed to use Matplotlib so that the whole suite uses the same tools [5]. The plotting for VOLK benchmarking was also changed to retrieve data from the new format and to create consistent coloring of the results for easier comparisons.

2.3 Application Profiling

For benchmarking an application we use oprofile, which samples the CPU either after a certain number of CPU events occur or after a regular time specified, to return the percentage of running time used in various functions [8, 2]. It is important that debug symbols be included in the target image so that applications can be profiled [8]. Since each function call, for example a single VOLK kernel, becomes a compiler symbol, by profiling a running GNU Radio application the functions that use the most time can be easily identified.

3 Results

3.1 Open Embedded

Getting a Poky distribution build with GNU Radio to boot on x86_64 (Intel Atom and AMD E350) and ARM (TI OMAP 3530) created several initial problems since the primary meta-oe development is done by Angstrom developers [1] which focuses primarily on ARM architectures. The primary issue was with the device manager and systemd; specifically, building GNU Radio brings in systemd through a chain of dependencies. To solve the booting issues on x86_64 and build GNU Radio we used the version of udev provided by *oe-core* and use BBMASK variables to mask out everything in the GNU Radio dependency chain that would lead to systemd or *meta-oe* udev versions being built. Falling back on older and more stable *oe-core* tools where necessary resulted in a Linux

image with GNU Radio that could boot and run the exact same software on the AMD-based Gumstix board (Ettus E110), the Intel Atom, and the AMD E350 APU.

3.2 Multi-Processor Scheduling




Figure 1: GFLOPs per second through an FFT array on an Intel i7.

Example output from multi-processor scheduling tests on an Intel i7 are shown in Figures 1, 2, 3, 4. The big performance differences come from using VOLK while doing parallel FFTs. On the FFT plots, Figures 1 and 2, comparing the gradient in the direction of pipes versus stages yields an interesting difference between generic kernels and using VOLK. Using generic kernels adding stages causes a larger increase to FLOPs/s compared to adding pipes; in contrast with VOLK enabled adding pipes causes an equivalent increase in FLOPs/s as adding

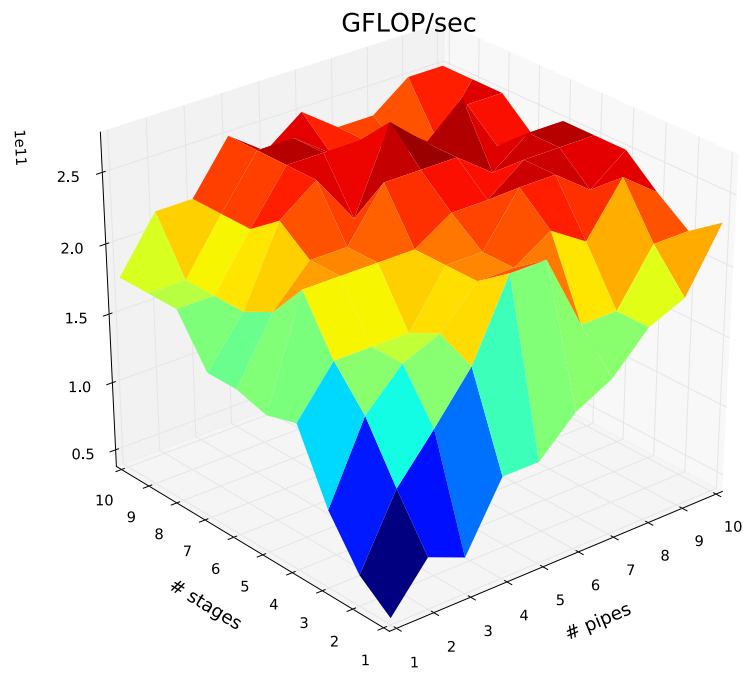


Figure 2: GFLOPs per second through an FFT array with VOLK enabled on an Intel i7.

stages does. Every processor will have a peak in the measured FLOPs where adding more pipes and stages will not yield an increase in FLOPs and past which the measured FLOPs will decrease, likely due to the increased need to store and fetch samples from memory.

The FIR filter results do not seem affected by VOLK kernels being used as opposed to the generic kernels. This is likely because FIR filter blocks have already been hand-tuned to use SIMD instructions without VOLK [6].




Figure 3: GFLOPs per second through an FIR filter array on an Intel i7.

3.3 VOLK kernels

Figures 5 and 6 show the results of benchmarking VOLK math operations implemented as stand-alone GNU Radio blocks using generic kernels and with VOLK optimizations, respectively. Type conversions with generic and VOLK




Figure 4: GFLOPs per second through an FIR filter array with VOLK enabled on an Intel i7.

improvements are shown in Figures 7 and 8. The height of each bar in these four graphs is the total time to repeat the named operation 1 billion times; the black bar shows one standard deviation of those 1 billion measurements.




Figure 5: VOLK math results using a generic VOLK kernel for different processors.

The Intel i7 is obviously much faster than other processors, and the E110 is obviously much slower. The E110 sees no improvement from VOLK across all benchmarks because ARM processors use the NEON architecture that is not well supported in VOLK yet. Some highlights from the VOLK results are that nearly all instructions on the x86 processors do see improvement from generic kernels. It is also clear that some processors have more performance gain from VOLK than others. For example comparing the `multiply_cc`, which executes a complex multiply and outputs a complex result, in Figure 5 the Atom is clearly faster than the E350 using generic kernels. Using VOLK, Figure 6 shows that the E350 is now slightly faster than the Atom. Similar results appear Figures 5 and 6 with `multiply_conjugate_cc`, `complex_to_mag` in Figures 7 and 8.

Some VOLK kernels are actually on par with or slower than the generic kernels. Since this occurs mostly on the simpler instructions such as `add_ff` and `multiply_const_ff` the likely cause is compilers and processors are already fine-tuned to do these instructions efficiently.




Figure 6: VOLK math results using VOLK kernels for different processors.

4 Conclusion

Multi-processor scheduling and speed of vector math will play an important role in processor selection for a software radio application. Being able to pick a processor that matches size, weight, and power constraints that match desired specifications requires knowledge of processor performance. By integrating and enhancing existing tools and working through the bugs to run GNU Radio through a Poky distribution build of OE we have introduced a benchmarking platform which can assist in choosing the best platform for embedded software radios. We have also introduced benchmarking results for a small set of potentially suitable processors for software radio.

References

- [1] Contributors to openembedded/meta-oe, October 2012. <https://github.com/openembedded/meta-oe/graphs/contributors>.
- [2] Oprofile, August 2012. <http://oprofile.sourceforge.net/about/>.
- [3] Eric Blossom. GNU Radio - MP Scheduler Performance. GNU Radio, 9 edition, July 2008. <http://gnuradio.org/redmine/projects/gnuradio/wiki/MPSchedulerPerformance>.




Figure 7: VOLK type conversion results using a generic VOLK kernel for different processors.




Figure 8: VOLK type conversion results using VOLK kernels for different processors.

- [4] Paul Eggleton and Lloyd Chang. Openembedded-core, September 2012. <http://www.openembedded.org/wiki/OpenEmbedded-Core>.
- [5] John Hunter, Darren Dale, Eric Firing, and Michael Droettboom. matplotlib, November 2012. <http://matplotlib.org/contents.html>.
- [6] Sean Nowlan. Discuss-gnuradio, November 2011. <http://lists.gnu.org/archive/html/discuss-gnuradio/2011-11/msg00151.html>.
- [7] George Nychis, Thibaud Hottelier, Zhuocheng Yang, Srinivasan Seshan, and Peter Steenkiste. Enabling mac protocol implementations in software-defined radios. NSDI'09 Proceedings of the 6th USENIX symposium on Networked systems design and implementation, 2009.
- [8] Richard Purdie. The Yocto Project Reference Manual. Linux Foundation, July 2012. <http://www.yoctoproject.org/docs/current/poky-ref-manual/poky-ref-manual.html>.
- [9] Scott Rifenbark. The Yocto Project Development Manual. Intel and Linux Foundation, July 2012. <http://www.yoctoproject.org/docs/current/dev-manual/dev-manual.html>.
- [10] Thomas Rondeau. Volk: Vector-optimized library of kernels. blog, December 2010. <http://www.trondeau.com/blog/2010/12/11/volk-vector-optimized-library-of-kernels.html>.
- [11] Thomas Rondeau. Volk benchmarking. blog, February 2012. <http://www.trondeau.com/blog/2012/2/17/volk-benchmarking.html>.
- [12] Tom Rondeau and Dimitrios Symeonidis. GNU Radio - VOLK. GNU Radio, 11 edition, March 2012. <http://gnuradio.org/redmine/projects/gnuradio/wiki/Volk>.
- [13] Kees van Berkel, Frank Heinle, Patrick P. E. Meuwissen, Kees Moerman, and Matthias Weiss. Vector processing as an enabler for software-defined radio in handheld devices. EURASIP J. Appl. Signal Process., 2005:2613–2625, January 2005.
- [14] M. Woh, Yuan Lin, Sangwon Seo, T. Mudge, and S. Mahlke. Analyzing the scalability of simd for the next generation software defined radio. In Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on, pages 5388 –5391, 31 2008-april 4 2008.
- [15] Yocto Project. Openembedded Core, October 2012. <http://www.yoctoproject.org/projects/openembedded-core>.

Listing 1: Benchmark control script.

```

1  #!/bin/sh

    # 1st parameter is processor description
    # 2nd parameter is gnuradio version
5  ARCH=$1
    GR_VERSION=$2
    FNAME=$1_"$2"_"
    export PYTHONPATH=utils/
9  volk_profile
    mp-sched/run_synthetic_fft.py -m 10 -D 'synth_fft.txt' -L $FNAME"
        fft_volked"
    mp-sched/run_synthetic_fir.py -m 10 -D 'synth_fir.txt' -L $FNAME"
        fir_volked"
    volk/volk_math.py -L $FNAME"volked" -D volk_math.db --all
13  volk/volk_types.py -L $FNAME"volked" -D volk_types.db --all
    # Change all architectures in volk_profile to generic
    # with sed magic
    echo 'old volk_config moved to ~/.volk/volk_config.volked'
17  echo 'generic volk_config being generated'
    mv ~/.volk/volk_config ~/.volk/volk_config.volked
    sed 's/(orc\|neon\|ss\)[a-z0-9_]*\/generic/' ~/.volk/volk_config.volked
        > ~/.volk/volk_config
    mp-sched/run_synthetic_fft.py -m 10 -D 'synth_fft.txt' -L $FNAME"
        fft_generic"
21  mp-sched/run_synthetic_fir.py -m 10 -D 'synth_fir.txt' -L $FNAME"
        fir_generic.raw"
    volk/volk_math.py -L $FNAME"generic" -D volk_math.db --all
    volk/volk_types.py -L $FNAME"generic" -D volk_types.db --all

```

Listing 2: Modified helper functions for testing.

```

1  #!/usr/bin/python
    #!/usr/bin/env python

    from gnuradio import gr
5  import math, sys, os, time, re, pickle

    try:
        import numpy
9  except ImportError:
        sys.stderr.write("Unable_to_import_Numpy\n")
        sys.exit(1)

13  # tables can have names with any letter, number, underscore, period, or
        dash
        table_name_chars = "[a-zA-Z0-9_.-]"

17  def common_args(parser):
        parser.add_argument('-D', '--database', type=str, required=True,
            help='Database_(pickled_file)_to_rw_results')
        parser.add_argument('--listtables',
21         default=False, action='store_true',
            help='print_a_list_of_tables_in_the_database_
                file')

        return parser

```

```

25 def create_connection(fname):
    '''
        return a file object. If it's not createx
    '''
29     try:
        return open(fname,'r+')
    except IOError:
        return open(fname,'w+')
33
def new_table(conn, tablename):
    '''
        Create a new "table" of sorts for the results. Each table
37     should be for a different architecture/machine/type. It's
        best to keep these names unique, but there's no checker
        planned for that.
        You should run list_tables first to make sure that you
41     aren't duplicating a table. I don't know what will happen
    '''
    conn.seek(0,os.SEEK_END) # go to end
    # command is the string to write to file
45     cmd = "\n<{0}>\n".format(tablename)
    conn.write(cmd)

49 def insert_results(conn, res):
    '''
        Insert results that are apparently dictionary values into
        the table. Since this is a text file we pickle the list.
53     '''
    # because mp-sched was designed to be two independent
    # programs (why oh why oh why?! - fix this later)
    # appending to a file is the "best" way to insert results
57     conn.seek(0,2) # go to the last byte in the file
    # the whole dictionary gets pickled
    conn.write("$newdatarow$")
    pickle.dump(res,conn)
61
def list_tables(conn):
    '''
        return a list of all tables in the database
65     '''
    conn.seek(0)
    fstring = conn.read()
    # match as few characters as possible inside angle brackets
69     table_names = re.findall("<("+table_name_chars+"*)>", fstring)
    return table_names

def get_results(conn, tname):
73     '''
        gets all results in tname. tname should match an arch+gr version
        No testing is done on originality of tname
    '''
77     pickle_string = '(\([a-zA-Z0-9\'\n]*s\.)?)'
    conn.seek(0)
    fstring = conn.read()
    # match the < to begin a table

```

```

81 tables = re.split('<.*>?', fstring)
    for i, table in enumerate(tables):
        if table.startswith('<' + tname + '>'):
            res = list()
85         entries = re.split("\$newdatarow\$", tables[i+1])
            for row in entries:
                try:
                    rdata = pickle.loads(row)
89                 except:
                    # yikes! I hope that wasn't data!
                    print
                else:
93                 rdata
                    res.append(rdata)
            return res

97 def close_connection(conn):
    conn.close()

```

Listing 3: Modified Matplotlib for VOLK results plotting.

```

1  #!/usr/bin/python
   #!/usr/bin/env python

   import sys, math
5  import argparse
   from common_test_funcs import *
   import collections

9  try:
       import matplotlib
       import matplotlib.pyplot as plt
   except ImportError:
13     sys.stderr.write("Could not import Matplotlib (http://matplotlib.
        sourceforge.net/)\n")
        sys.exit(1)

   def main():
17     desc='Plot_Volk_performance_results_from_a_SQLite_database.' + \
        'Run_one_of_the_volk_tests_first_(e.g.,_volk_math.py)'
        parser = argparse.ArgumentParser(description=desc)
        # parser.add_argument('-D', '--database', type=str,
21         #                   default='volk_results.db',
        #                   help='Database file to read data from [default
        : %(default)s]')
        parser.add_argument('-E', '--errorbars',
25         action='store_true', default=False,
        help='Show_error_bars_(1_standard_dev.)')
        parser.add_argument('-P', '--plot', type=str,
        choices=['mean', 'min', 'max'],
        default='mean',
29         help='Set_the_type_of_plot_to_produce_[default:
        _%(default)s]')
        parser.add_argument('-%', '--percent', type=str,
        default=None, metavar="table",
        help='Show_percent_difference_to_the_given_type
        _[default:_%(default)s]')
33     parser.add_argument('-T', '--tables', type=str, nargs='*',

```



```

                                default=None,
                                help='select_the_tables_to_plot')
37 parser.add_argument('-o', '--output', type=str, default="",
                                help='file_to_save_output_to_(svg)')
parser = common_args(parser)
args = parser.parse_args()

41
conn = create_connection(args.database)
if(args.listtables):
    tables = list_tables(conn)
45     for t in tables:
        print t
        exit(0)

49
# Set up global plotting properties
matplotlib.rcParams['figure.subplot.bottom'] = 0.2
matplotlib.rcParams['figure.subplot.top'] = 0.95
53 matplotlib.rcParams['figure.subplot.right'] = 0.98
matplotlib.rcParams['ytick.labelsize'] = 16
matplotlib.rcParams['xtick.labelsize'] = 16
matplotlib.rcParams['legend.fontsize'] = 18

57
# Get list of tables to compare
tables = list_tables(conn)
print args.tables
61 try:
    if set(args.tables) & set(tables) == set(args.tables):
        tables=args.tables
    else:
65     print 'sorry_couldnt_find_all_of_your_tables,_try_a_--
        listtables'
        exit(0)
except TypeError:
69     print 'Empty_list_of_tables_provided,_plotting_all'

M = len(tables)

73
# Colors to distinguish each table in the bar graph
# More than 5 tables will wrap around to the start.
colors = ['b', 'r', 'g', 'm', 'k']

77
# Set up figure for plotting
f0 = plt.figure(0, facecolor='w', figsize=(14,10))
s0 = f0.add_subplot(1,1,1)

81
# Create a register of names that exist in all tables
tmp_regs = []
for table in tables:
    # Get results from the next table
    res = get_results(conn, table)

85
    tmp_regs.append(list())
    for r in res:
        try:
89         tmp_regs[-1].index(r['kernel'])

```

```

        except ValueError:
            tmp_regs[-1].append(r['kernel'])

93     # Get only those names that are common in all tables
name_reg = tmp_regs[0]
    for t in tmp_regs[1:]:
        name_reg = list(set(name_reg) & set(t))
97     name_reg.sort()

    # Pull the data out for each table into a dictionary
    # we can ref the table by it's name and the data associated
101    # with a given kernel in name_reg by it's name.
    # This ensures there is no sorting issue with the data in the
    # dictionary, so the kernels are plotted against each other.
table_data = collections.OrderedDict()
105    for i,table in enumerate(tables):
        # Get results from the next table
        print 'results_from_' + table
        res = get_results(conn, table)

109        data = dict()
        for r in res:
            data[r['kernel']] = r

113        table_data[table] = data

    if args.percent is not None:
117        for i,t in enumerate(table_data):
            if args.percent == t:
                norm_data = []
                for name in name_reg:
121                    if(args.plot == 'max'):
                        norm_data.append(table_data[t][name]['max'])
                    elif(args.plot == 'min'):
                        norm_data.append(table_data[t][name]['min'])
125                    elif(args.plot == 'mean'):
                        norm_data.append(table_data[t][name]['avg'])

129    # Plot the results
x0 = xrange(len(name_reg))
i = 0
    # put in to an ordered dict (and order by key) so similar tables
    # come out
133    # with matching colors -- could probably be improved by using
    # orderdict to start
    # table_data = collections.OrderedDict(sorted(table_data.items(),
    # key=lambda t: t[0]))
    for t in (table_data):
        ydata = []
137        stds = []
        for name in name_reg:
            stds.append(math.sqrt(table_data[t][name]['var']))
            if(args.plot == 'max'):
141                ydata.append(table_data[t][name]['max'])
            elif(args.plot == 'min'):
                ydata.append(table_data[t][name]['min'])

```

```

145         elif(args.plot == 'mean'):
            ydata.append(table_data[t][name]['avg'])

    if args.percent is not None:
        ydata = [-100*(y-n)/y for y,n in zip(ydata,norm_data)]
149         if(args.percent != t):
            # makes x values for this data set placement
            # width of bars depends on number of comparisons
            wdh = 0.80/(M-1)
153             x1 = [x + i*wdh for x in x0]
                i += 1

            s0.bar(x1, ydata, width=wdh,
157                 color=colors[(i-1)%M], label=t,
                 edgcolor='k', linewidth=2)

        else:
161             # makes x values for this data set placement
            # width of bars depends on number of comparisons
            wdh = 0.80/M
            x1 = [x + i*wdh for x in x0]
165             i += 1

            if(args.errorbars is False):
                s0.bar(x1, ydata, width=wdh,
169                     color=colors[(i-1)%M], label=t,
                     edgcolor='k', linewidth=2)

            else:
                s0.bar(x1, ydata, width=wdh,
173                     yerr=stds,
                     color=colors[i%M], label=t,
                     edgcolor='k', linewidth=2,
                     error_kw={"ecolor": 'k', "capsize":5,
177                             "linewidth":2})

        nitems = res[0]['nsamples']
        if args.percent is None:
181             s0.set_ylabel("Processing_time_(sec)_[{0:G}_items]".format(
                nitems),
                fontsize=22, fontweight='bold',
                horizontalalignment='center')

        else:
185             s0.set_ylabel("%_Improvement_over_{0}_[{1:G}_items]".format(
                args.percent, nitems),
                fontsize=22, fontweight='bold')

189         s0.legend()
        s0.set_xticks(x0)
        s0.set_xticklabels(name_reg)
        for label in s0.xaxis.get_ticklabels():
193             label.set_rotation(45)
                label.set_fontsize(16)

        if args.output == "_":
197             plt.show()
        else:
            plt.savefig(args.output, format='pdf')

```

```

201 if __name__ == "__main__":
    main()

```

Listing 4: Matplotlib plotting tools for mp-sched results.

```

#!/usr/bin/python
2 #!/usr/bin/env python

import sys, math
import argparse
6 from common_test_funcs import *
from mpl_toolkits.mplot3d import Axes3D
from matplotlib import cm

10 try:
    import matplotlib
    import matplotlib.pyplot as plt
    except ImportError:
14     sys.stderr.write("Could not import Matplotlib_(http://matplotlib.
        sourceforge.net/)\n")
    sys.exit(1)

def main():
18     desc='Plot_synthetic_FFT/FIR_filters_from_a_pickled_file_of_results
        .' + \
        'Run_one_of_the_synthetic_tests_first_(preferbly_with_
        run_benchmarking)'
    parser = argparse.ArgumentParser(description=desc)
    parser.add_argument('-T', '--tables', type=str, nargs='*',
22         default=None,
        help='select_the_tables_to_plot')
    parser.add_argument('-o', '--output', type=str, default="",
        help='file_to_save_output_to_(svg)')
26

    parser = common_args(parser)
    args = parser.parse_args()
30    conn = create_connection(args.database)

    tables = list_tables(conn)
    if(args.listtables):
34        for t in tables:
            print t
        exit(0)

38

    # Set up global plotting properties
    matplotlib.rcParams['figure.subplot.bottom'] = 0.2
    matplotlib.rcParams['figure.subplot.top'] = 0.95
42    matplotlib.rcParams['figure.subplot.right'] = 0.98
    matplotlib.rcParams['ytick.labelsize'] = 16
    matplotlib.rcParams['xtick.labelsize'] = 16
    matplotlib.rcParams['legend.fontsize'] = 18
46

    # Get list of tables to compare
    tables = list_tables(conn)

```

```

50     try:
        if tables.__contains__(args.tables[0]):
            tables=args.tables[0]
        else:
54         print 'sorry_couldnt_find_all_of_your_tables,_try_a_--
            listtables'
            exit(0)
    except TypeError:
        print 'Empty_list_of_tables_provided!'
58     exit(1)

    # f0 = plt.figure(0, facecolor='w', figsize=(14,14))
    # s0 = f0.add_subplot(1,1,1)
62     f1 = plt.figure(1, facecolor='l', figsize=(14,14))
    s1 = f1.add_subplot(111, projection='3d', azimuth=230)

66     res = get_results(conn, tables)

    max_stages = 0
    max_pipes = 0
70     min_stages = 20
    min_pipes = 20
    results = numpy.zeros((10,10))
    for r in res:
74         max_stages = max(max_stages, r['stages'])
            min_stages = min(max_stages, r['stages'])
            max_pipes = max(max_pipes, r['pipes'])
            min_pipes = min(max_pipes, r['pipes'])
78         results[r['pipes']-1,r['stages']-1] = r['pseudoflopreal']

    # stages = numpy.arange(min_stages, max_stages)
    # pipes = numpy.arange(min_stages, max_stages)
82     stages = numpy.arange(1,11)
    pipes = numpy.arange(1,11)
    pipes3d, stages3d = numpy.meshgrid(pipes, stages)
    # print results.size
86     # print pipes.size
    # print stages.size
    surf = s1.plot_surface(pipes3d, stages3d, results, rstride=1,
        cstride=1,cmap=cm.jet,
            linewidth=0, antialiased=True)
90     # f1.colorbar(surf, shrink=0.5, aspect=5)
    # s1.set_zlim3d(results[min_pipes-1, min_stages-1]/2, results.max()
        *1.05)
    # contour_plot = s0.contour(pipes, stages, results/(10**9))
    plt.title('GFLOP/sec', fontsize=28)
94     plt.xlabel('#_pipes', fontsize=20)
    plt.ylabel('#_stages', fontsize=20)
    # plt.clabel(contour_plot, inline=1, fontsize=14)

98     if args.output == "_":
        plt.show()
    else:
        plt.savefig(args.output, format='pdf')
102

```

```

106 if __name__ == "__main__":
    main()

```

Listing 5: Modified mp-sched testing for FFTs

```

1  #!/usr/bin/python
  #!/usr/bin/env python
  #
  # Copyright 2008 Free Software Foundation, Inc.
5  #
  # This file is part of GNU Radio
  #
  # GNU Radio is free software; you can redistribute it and/or modify
9  # it under the terms of the GNU General Public License as published by
  # the Free Software Foundation; either version 3, or (at your option)
  # any later version.
  #
13 # GNU Radio is distributed in the hope that it will be useful,
  # but WITHOUT ANY WARRANTY; without even the implied warranty of
  # MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
  # GNU General Public License for more details.
17 #
  # You should have received a copy of the GNU General Public License
  # along
  # with this program; if not, write to the Free Software Foundation, Inc
  #
  # 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.
21 #
  """
  Run synthetic.py for npipes in [1,16], nstages in [1,16]
25 """

  import re
  import sys
29 import os
  import tempfile
  from optparse import OptionParser
  from common_test_funcs import *
33
  current_dir = os.path.dirname(__file__)

  def write_shell_script(f, description, ncores, gflops,
                        max_pipes_and_stages, database):
37      """
        f is the file to write the script to
        data_filename is the where the data ends up
        description describes the machine
41      ncores is the number of cores (used to size the workload)
        gflops is the estimated GFLOPS per core (used to size the workload)
        """

45      f.write("#!/bin/sh\n")
        f.write("\n")
        if description:
            f.write("echo '#D_%s'\n" % (description,))

```

```

49     for npipes in range(1, max_pipes_and_stages + 1):
        for nstages in range(1, max_pipes_and_stages + 1):
            # We'd like each run of synthetic to take ~10 seconds
53         desired_time_per_run = 10
            est_gflops_avail = min(nstages * npipes, ncores) * gflops
            nsamples = (est_gflops_avail * desired_time_per_run)/(512.0
                * nstages * npipes)
            nsamples = int(nsamples * 1e9)
57         cmd = "./s/synthetic_fft.py_m-s%d_p%d_N%d_D%s\n"
                % (current_dir, nstages, npipes, nsamples, database)
            f.write(cmd)
            f.write('if_test_?_ge_128;_then_exit_128;_fi\n')
61         f.write("_2>&1_\n" )
            f.flush()
65
def main():
    description = """prog gathers multiprocessor scaling data using
        the ./synthetic.py benchmark.
69 All combinations of npipes and nstages between 1 and --max-pipes-and-
        stages are tried.
        The -n and -f options provides hints used to size the workload. We'd
        like each run
        of synthetic to take about 10 seconds. For the full 16x16 case this
        results in a
        total runtime of about 43 minutes, assuming that your values for -n and
        -f are reasonable.
73 For x86 machines, assume 3 FLOPS per processor Hz. E.g., 3 GHz machine
        -> 9 GFLOPS.
        plot_flops.py will make pretty graphs from the output data generated by
        %prog.
    """
    usage = "usage:_%prog_[options]_D_results.db_L_label"
77     parser = OptionParser(usage=usage, description=description)
    parser.add_option("-d", "--description", metavar="DESC",
        help="machine_description, e.g., \"Dual_quad-core_
        Xeon_3_GHz\"", default=None)
    parser.add_option("-n", "--ncores", type="int", default=1,
81         help="number_of_processor_cores_[default=%default
        ]")
    parser.add_option("-g", "--gflops", metavar="GFLOPS", type="float",
        default=3.0,
        help="estimated_GFLOPS_per_core_[default=%default
        ]")
    parser.add_option("-m", "--max-pipes-and-stages", metavar="MAX",
        type="int", default=16,
85         help="maximum_number_of_pipes_and_stages_to_use_[
        default=%default]")
    parser.add_option("-D", "--database", metavar="CONN", type="str")
    parser.add_option("-L", "--table", metavar="TABLE", type="str")
89     (options, args) = parser.parse_args()

```

```
shell = os.popen("/bin/sh", "w")
93 write_shell_script(shell,
                    options.description,
                    options.ncores,
97                    options.gflops,
                    options.max_pipes_and_stages,
                    options.database)
c=create_connection(options.database)
101 new_table(c, options.table)
close_connection(c)

if __name__ == '__main__':
105     main()
```