


The GNU Radio scheduler and internals

Making samples flow

Marcus Müller

17 June 2019



Structure

Introduction

Constituents of a GNU Radio application


Signal Flow Architecture

Implementation

Problems and Challenges

Conclusion

Questions & Answers

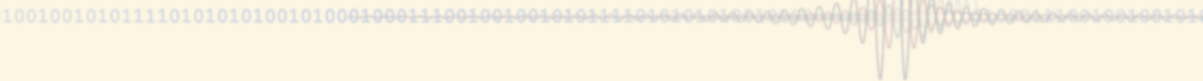



The SDR Problem

... as interpreted by GNU Radio


- ▶ get samples from SDR hardware
- ▶ process
- ▶ get samples back to SDR hardware
- ▶ high rates
- ▶ commodity general-purpose CPUs

→ GNU Radio's flow graph approach



GNU Radio's Signal Flow Graphs

- ▶ Blocks have in- and outputs
- ▶ an output can be connected to many inputs
- ▶ an input can only be connected to one output
- ▶ developers concentrate on writing these blocks
- ▶ GNU Radio takes care of the data marshalling
 - ▶ writing a block: just enter your code into a `work` function



What is a Block?

- ▶ Instance of subclass of `gr::block`
- ▶ most importantly: `general_work` method arguments:
 - ▶ `int noutput_items` (\min_k (available space on output k))
 - ▶ `vector<int> ninput_items` (available items per input)
 - ▶ `vector<const void*> input_items` (pointers to input arrays)
 - ▶ `vector<void*> output_items` (pointers to output arrays)

return value: number of items produced¹

¹or declare yourself using `produce(k, how_many)` and return `WORK_CALLED_PRODUCE`

What is a Block?

- ▶ Instance of subclass of `gr::block`
- ▶ most importantly: `general_work` method arguments:
 - ▶ `int noutput_items` (\min_k (available space on output k))
 - ▶ `vector<int> ninput_items` (available items per input)
 - ▶ `vector<const void*> input_items` (pointers to input arrays)
 - ▶ `vector<void*> output_items` (pointers to output arrays)
- return value: number of items produced¹
- ▶ forecast: “If asked to produce N output items, how much items per input will you need?”
- ▶ not forced to consume all input
 - ▶ can consume less at single-item granularity
 - ▶ “leftovers” beginning of next call’s `input_items`
- ▶ `gr::sync_block`, `decimator`, `interpolator`: Convenience classes that simplify `general_work` to work, implement forecast (in- to output relation known)

¹or declare yourself using `produce(k, how_many)` and return `WORK_CALLED_PRODUCE`

From Blocks to Graphs

Instantiate blocks, connect them:

- ▶ GNU Radio has graph-encapsulating blocks (`hier_block`, `top_block`)
- ▶ these have `connect(src, output_port, sink, input_port)`
- ▶ builds a directed graph internally




From Blocks to Graphs

Instantiate blocks, connect them:

- ▶ GNU Radio has graph-encapsulating blocks (`hier_block`, `top_block`)
- ▶ these have `connect(src, output_port, sink, input_port)`
- ▶ builds a directed graph internally

After “declaration” of abstract flow graph, call `top_block.start()`:

- ▶ allocates fixed buffers acting as transport
- ▶ sets up scheduling (mostly: starts threads)
- ▶ calls each block's `start()` (optionally, sets up to IO etc.)
- ▶ asks around who can begin, and
- ▶ lets the block with zero input requirements produce items



From Blocks to Graphs

Instantiate blocks, connect them:


- ▶ GNU Radio has graph-encapsulating blocks (`hier_block`, `top_block`)
- ▶ these have `connect(src, output_port, sink, input_port)`
- ▶ builds a directed graph internally

After “declaration” of abstract flow graph, call `top_block.start()`:

- ▶ allocates fixed buffers acting as transport
- ▶ sets up scheduling (mostly: starts threads)
- ▶ calls each block’s `start()` (optionally, sets up to IO etc.)
- ▶ asks around who can begin, and
- ▶ lets the block with zero input requirements produce items

→ a lot to unpack, bear with me

Signal Flow Architecture



- ▶ GNU Radio is a backpressure-driven parallel signal processing architecture
- ▶ Blocks produce as much output as they can at once, given
 - ▶ available input data ready at the start of processing
 - ▶ available output data memory
- ▶ Every block runs in its own thread
- ▶ asked to produce $\min(\text{buffer size} / 2, \text{available output buffer})$
- ▶ Block can start working again while downstream block is still consuming
- ▶ → high parallelism

Message Passing



Extension to stream-based architecture

Shortcomings of the stream-only architecture

- ▶ for causality (and historical) reasons: No loops
- ▶ What if we need to sporadically update something (e.g. the equalizer taps from a downstream channel estimator)?

Solution: Asynchronous messaging

- ▶ Message ports get their own connections
- ▶ Insert *handle messages* step in block execution loop
- ▶ Message handlers are executed when work is **not**
 - ▶ No risk of changing block state while doing DSP
- ▶ Should probably replace 90% of block setters



Scheduling Mechanism in detail

Thread-per-Block (TPB)S cheduler (formerly: single-threaded scheduler (STS), now defunct)

- ▶ Each block gets its own executing thread²

When notified³,

- ▶ ask the block (forecast) whether it can produce output, given the available input and output space.

If *READY*

- ▶ call `general_work` **DSP happening here** (this might take some time)
- ▶ notify the upstream block(s) that we've consumed → free output buffer
- ▶ notify the downstream block(s) that we've produced → new input

If *blocked by lack of input*

- ▶ go to sleep for a while and check back later

If *blocked by lack of output space*

- ▶ go to sleep until notification

²`tpb_thread_body.cc`, `block_executor.cc`

³ignoring asynchronous message passing

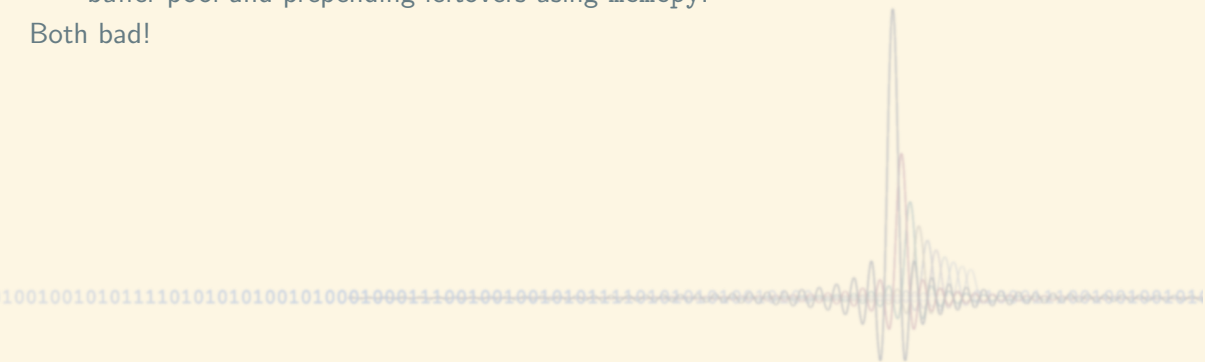
Fixed Buffers

- ▶ Freedom to not consume all input, and
- ▶ Guarantee that memory is contiguous

implies either

- ▶ infinitely large buffers or
- ▶ buffer pool and prepending leftovers using `memcpy`.

Both bad!




Fixed Buffers

- ▶ Freedom to not consume all input, and
- ▶ Guarantee that memory is contiguous

implies either

- ▶ infinitely large buffers or
- ▶ buffer pool and prepending leftovers using `memcpy`.

Both bad!



Ring Buffers

- ▶ Imagine a address space with modulo arithmetic
- ▶ Challenge: x86/ARM doesn't have an Address Generation Unit (AGU) that would would easily allow for that
- ▶ Approach: Use the Memory Management Unit (MMU)!
 - ▶ Designed for memory segmentation; can map 4 kB *pages* to arbitrary userland process memory locations
 - ▶ allows for double-mapping the same logical page to multiple locations
- ▶ Method: Double-mapped memory
 - ▶ Get one desired buffer size of shared memory or file-backed memory (only way for userland to map memory)
 - ▶ map it “back-to-back” in memory → emulate ring buffer



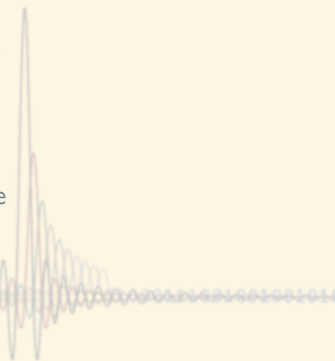
Ring Buffers

- ▶ Imagine a address space with modulo arithmetic
- ▶ Challenge: x86/ARM doesn't have an Address Generation Unit (AGU) that would would easily allow for that
- ▶ Approach: Use the Memory Management Unit (MMU)!
 - ▶ Designed for memory segmentation; can map 4 kB *pages* to arbitrary userland process memory locations
 - ▶ allows for double-mapping the same logical page to multiple locations
- ▶ Method: Double-mapped memory
 - ▶ Get one desired buffer size of shared memory or file-backed memory (only way for userland to map memory)
 - ▶ map it “back-to-back” in memory → emulate ring buffer



Problems and Challenges

- ▶ GNU Radio has ca. 18 years of history
- ▶ Not all decisions made in that period apply to the current architecture
 - ▶ to be completely honest, not even all decisions were good
- ▶ Use cases have evolved
 - ▶ Beginnings: Nearly only stream (TV, audio broadcast) processing
 - ▶ Nowadays: Real-time systems doing packetized data
- ▶ Environment has changed
 - ▶ SDR Hardware that supports bursting
 - ▶ Accelerators (GPUs, FPGAs, even network cards) widely available



Challenges for Scheduling

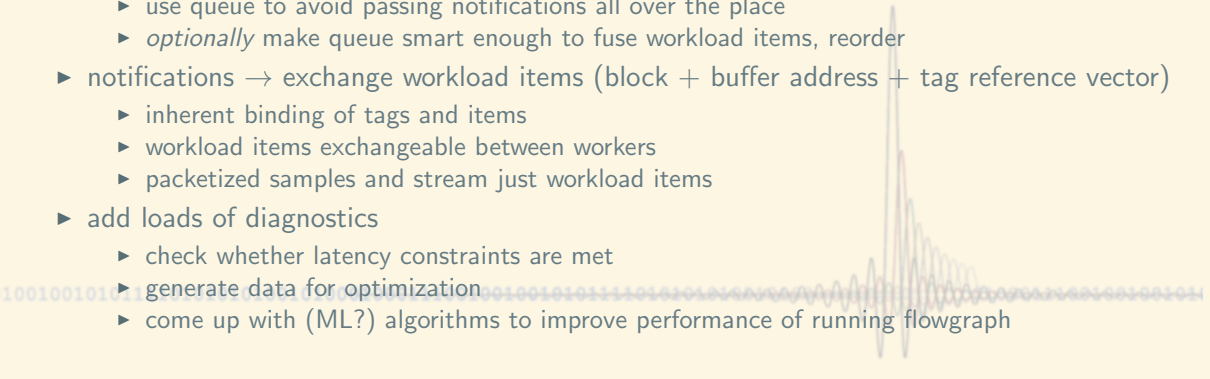
- ▶ Every block gets own thread getting scheduled on randomly (OS-)chosen CPU core
 - ▶ **clearly** suboptimal
- ▶ Memory-mapped Pseudo-Ring Buffers
 - ▶ Many blocks always consume all input → *no need* for ring buffer
 - ▶ DMA'd accelerator memory and doubly-mapped pages technologically mutually exclusive
 - ▶ Inefficient separate handling of tags, which logically belong to chunks of samples



Better Scheduling

Proposed Changes we came up with so far

Not happening in GNU Radio 3.8!

- ▶ Thread-per-Block → Worker Threads (matching CPU core count)
 - ▶ assign graph-sequential blocks to the same worker for memory locality
 - ▶ use queue to avoid passing notifications all over the place
 - ▶ *optionally* make queue smart enough to fuse workload items, reorder
 - ▶ notifications → exchange workload items (block + buffer address + tag reference vector)
 - ▶ inherent binding of tags and items
 - ▶ workload items exchangeable between workers
 - ▶ packetized samples and stream just workload items
 - ▶ add loads of diagnostics
 - ▶ check whether latency constraints are met
 - ▶ generate data for optimization
 - ▶ come up with (ML?) algorithms to improve performance of running flowgraph
- 

Questions & Answers

Ask away!

