

# Click Modular Router

## I. Click Modular Router

Key ideas:

- o Static graph of elements -- can verify well-formed connections
- o Extensibility through interposition or redefinition -- every wire and every element is a place for possible extension
- o Both push and pull processing
- o Kinds of state: local, global (poor), configurations (static) and flow based
- o Flexibility with only minor performance hit (mostly due to virtual functions)

Push vs. Pull:

- o in/out can be push, pull or agnostic
- o pull is call/return from destination, just like database query processing
- o push is driven by the sender, like upcalls
- o pull is good for polling and scheduling
- o push is good for arrivals (and events in general)
- o a queue is a push to pull converter
- o Volcano's "exchange operator" is a pull to push converter, not clear if Click has an analog, although you could build one. Might be useful for multiprocessor version?

Graph properties:

- o connections must match: pull to pull, push to push, or either to agnostic
- o statically checked, and all agnostic ports are resolved (to push/pull) statically
- o elements have configuration parameters and initialization (like "open")
- o no obvious equivalent for "close", although there is an atomic update to the whole graph
- o push outputs have one wire (otherwise implies duplicating output)
- o pull inputs have one wire (otherwise have to decide which one to call, or merge them)
- o push inputs and pull outputs can have multiple wires, first-come first served
- o declarative language (for the graph) => room for optimization (see his PhD)

Scheduling:

## Click Router

- o single thread (!)
- o only some elements need to be scheduled (tasks) -- those that push or pull independent of a caller
- o most just execute when called (pushed or pulled)
- o stride scheduling for tasks (why is this a good idea?)
- o tasks must yield the thread by returning -- not clear how long tasks should run or how it affects future scheduling decisions
- o deadlock? likely if any element can block, since there is no pre-emption and no other threads. exception: block for arrival might work (since that is a form of concurrency)
- o Polling vs. Interrupts: polling is much faster (8x?); processor stays in the sweet spot (no pipeline flushes), better caching, and easy to batch arrivals for lower overhead (and even more locality)
- o “receiver livelock” -- problem in which interrupts are arriving faster than the handling rate, which prevents any real work from getting done. Polling doesn’t have this problem, but you must drop packets at poll time to keep it in the operating range. Note how flat the throughput curves are when they reach saturation -- excellent “graceful degradation”, much like SEDA
- o Note the use of shared memory to get packets, much like U-Net.

### Flow Context:

- o Queried at initialization time to figure out properties about your context (neighbors)
- o Example properties: nearest queue, how many queues downstream or upstream, what interface is this from
- o Enables elements to behave differently based on context
- o Best example is looking at the total downstream queue length for congestion estimation
- o Side note: RED is “random early detection” (Floyd et al. 93). Key idea is to drop packets with (linearly) higher probability as the queue length gets longer, starting at some min. Attacks congestion proactively and earlier, hoping to keep the delays low.
- o In a DBMS, these kinds of things are done by the query optimizer, which typically has lots of specialized operators (and parameters) to use in different contexts

### Initialization:

- o broken into stages to avoid cycle dependencies
- o not clear who decides what goes in each stage; presumably the programmer has to participate

### Annotations:

- o Idea: add metadata to packet to carry information downstream
- o Statically defined fields

## Click Router

- o In theory every producer of an annotation should have exactly one consumer on all possible paths (where discard counts as a consumer). This does not appear to be enforced (or enforcable), since the annotations are not part of the language.
- o This is a bit of a hack: it is its own system that can't be type checked or verified, but it still required for correct operation.
- o Also no support for global information. Not clear why this was avoided. (In general, larger scopes are bad, but if you need one, it is better to create one than to abuse other mechanisms like configurations or annotations.)

### Performance:

- o overall excellent
- o polling really helps
- o flexibility costs about 1us per packet -- mostly due to virtual functions (in C++). This is because every element is a subclass with lots of virtual functions, so every push/pull call is a virtual function call with a dynamic level of indirection. However, since the graph is really static, you can flatten the virtual functions either statically or (in theory) at initialization time.
- o Supports its claims of extensibility and flexibility quite well
- o Supports its claim of low overhead for this flexibility