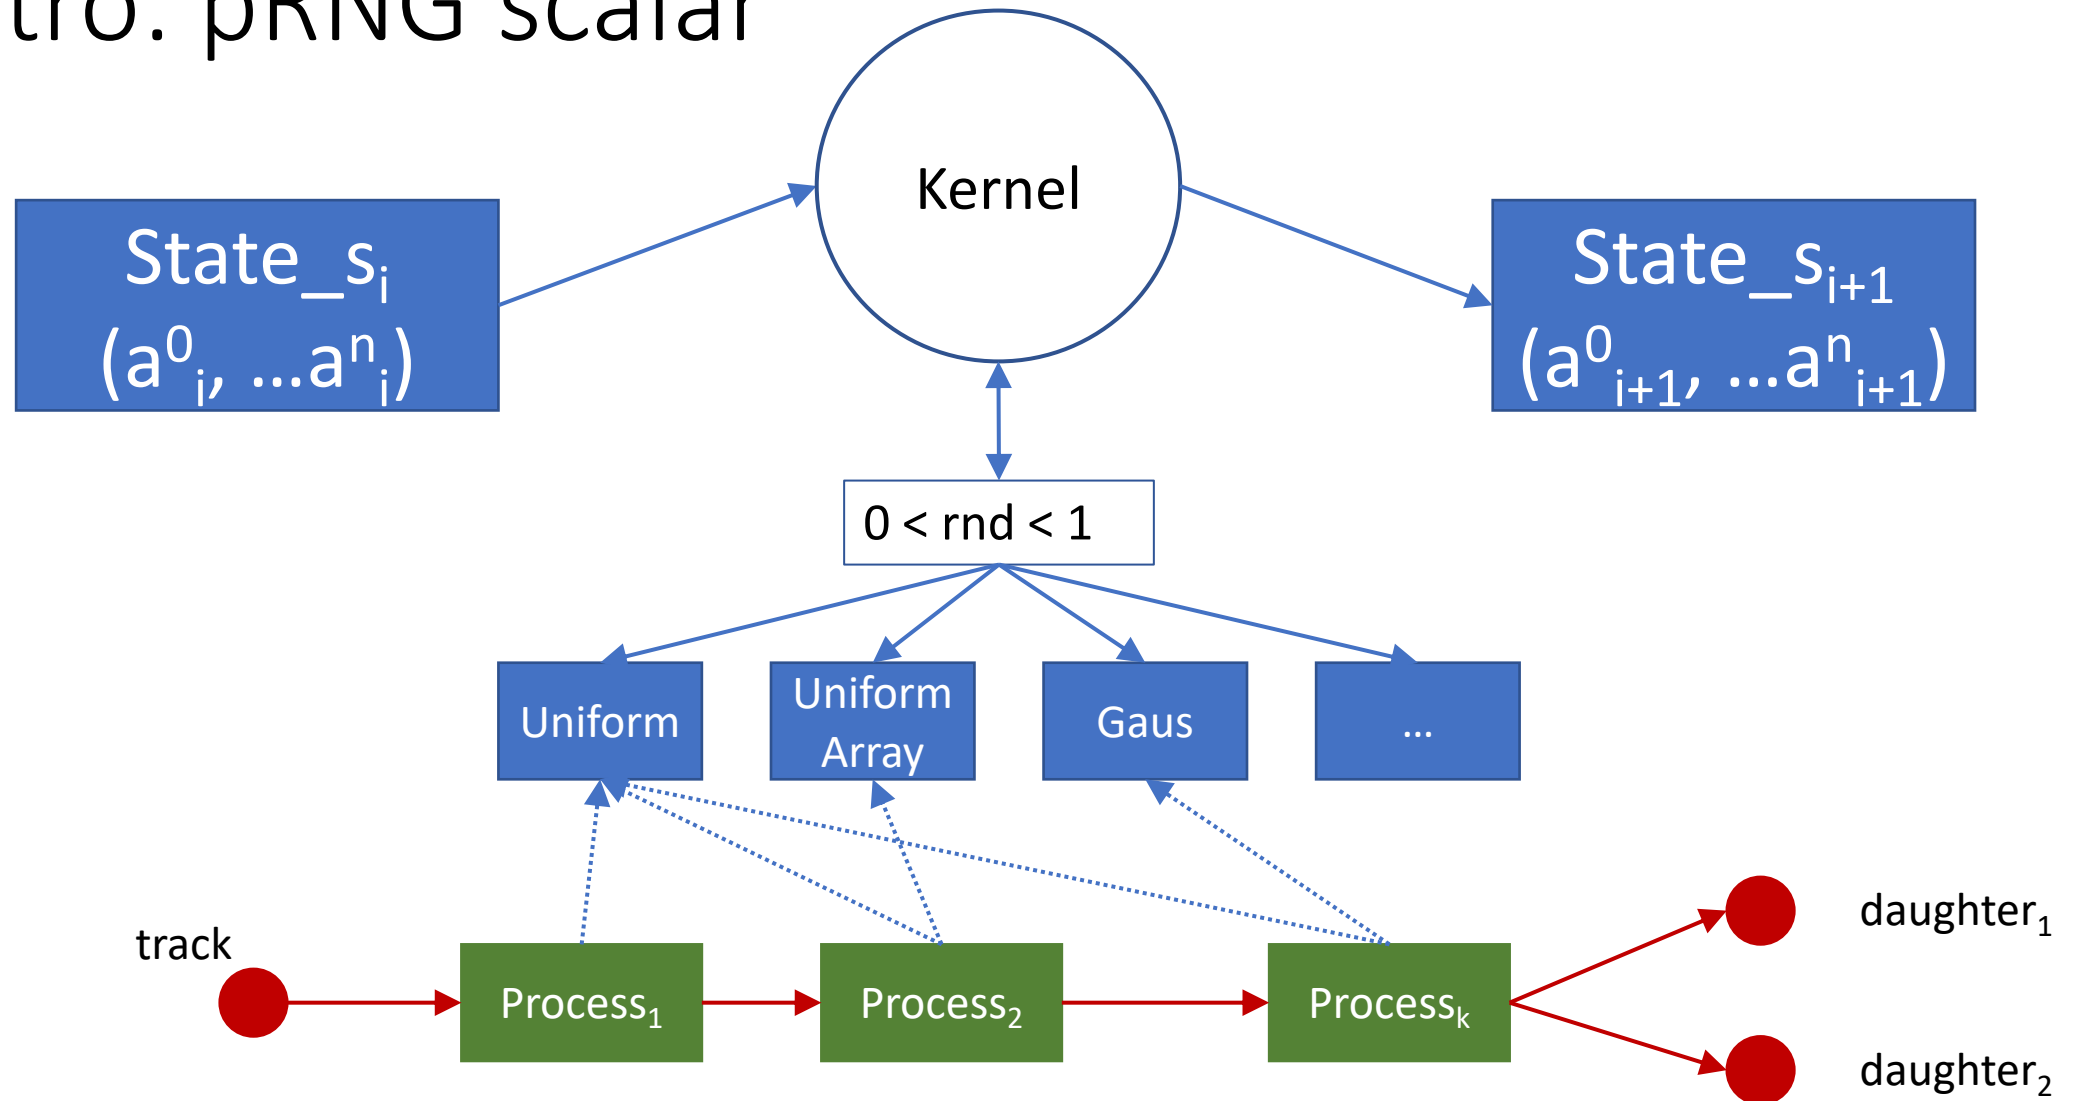


PRNG reproducibility

in multi-thread & scalar & multi-track (basket) case

Intro: pRNG scalar

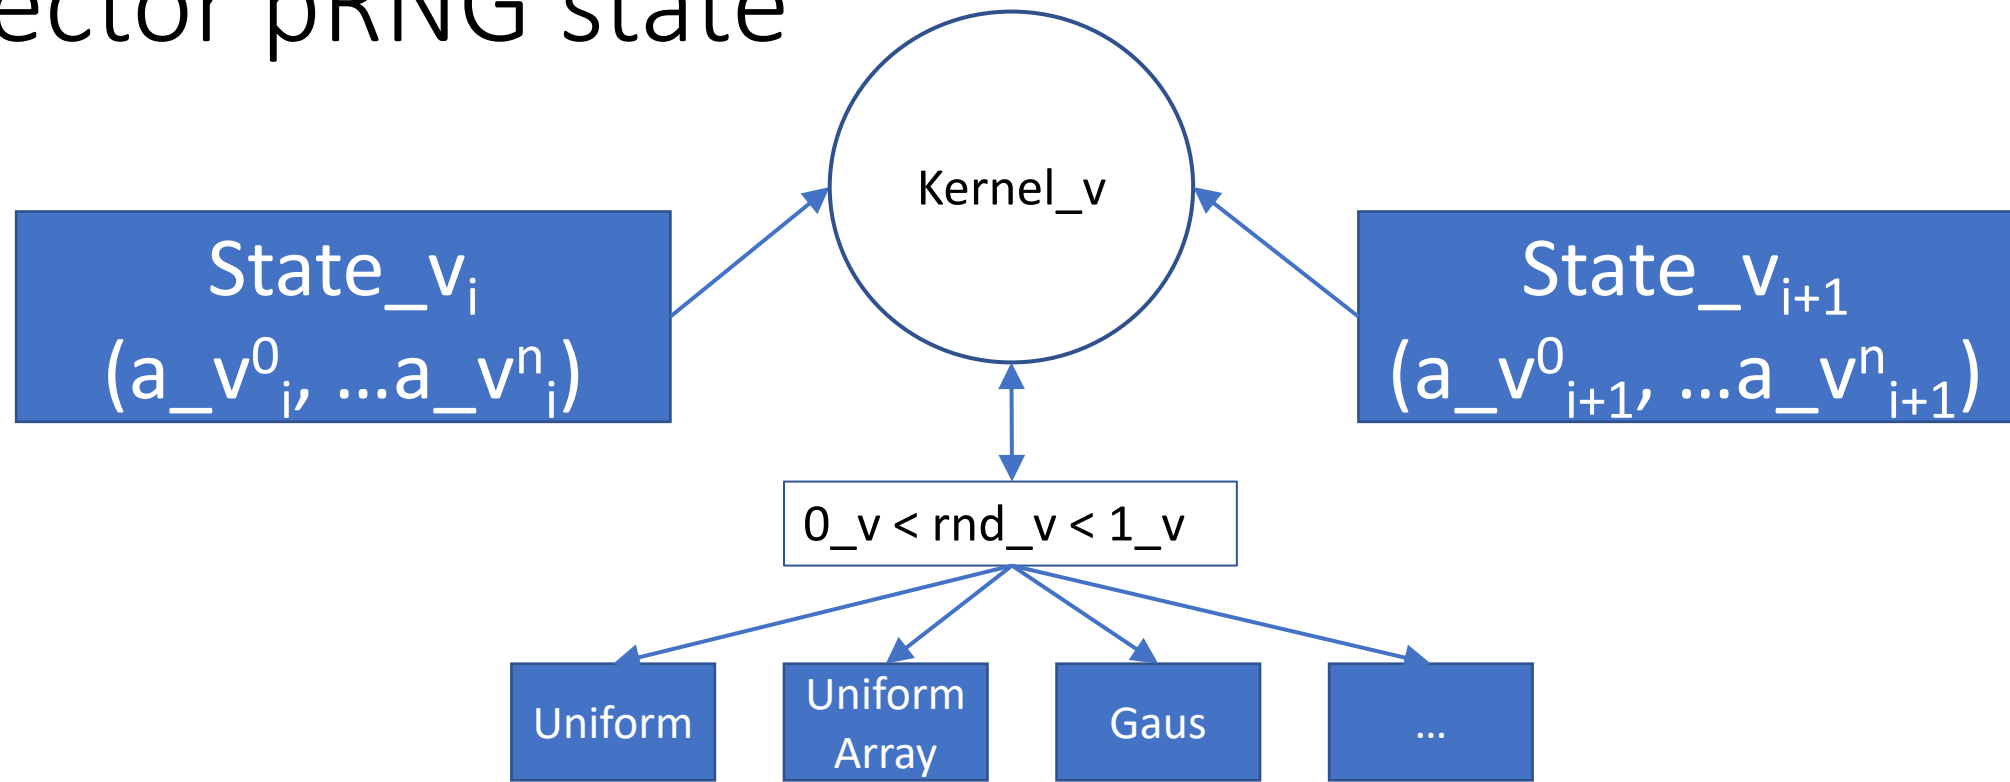


Reproducibility

- Conditions
 - Starting the random sequence from the same seed (state)
 - Preserving the original track sequencing (e.g. basketized vs. non-basketized)
- A global pRNG state is enough if conditions above fulfilled
- Multi-threading and/or basketization change track sequencing for calls to pRNG
 - Need to make the pRNG state intrinsically coupled with the track state, not exposing its evolution to sequencing



Vector pRNG state



where:

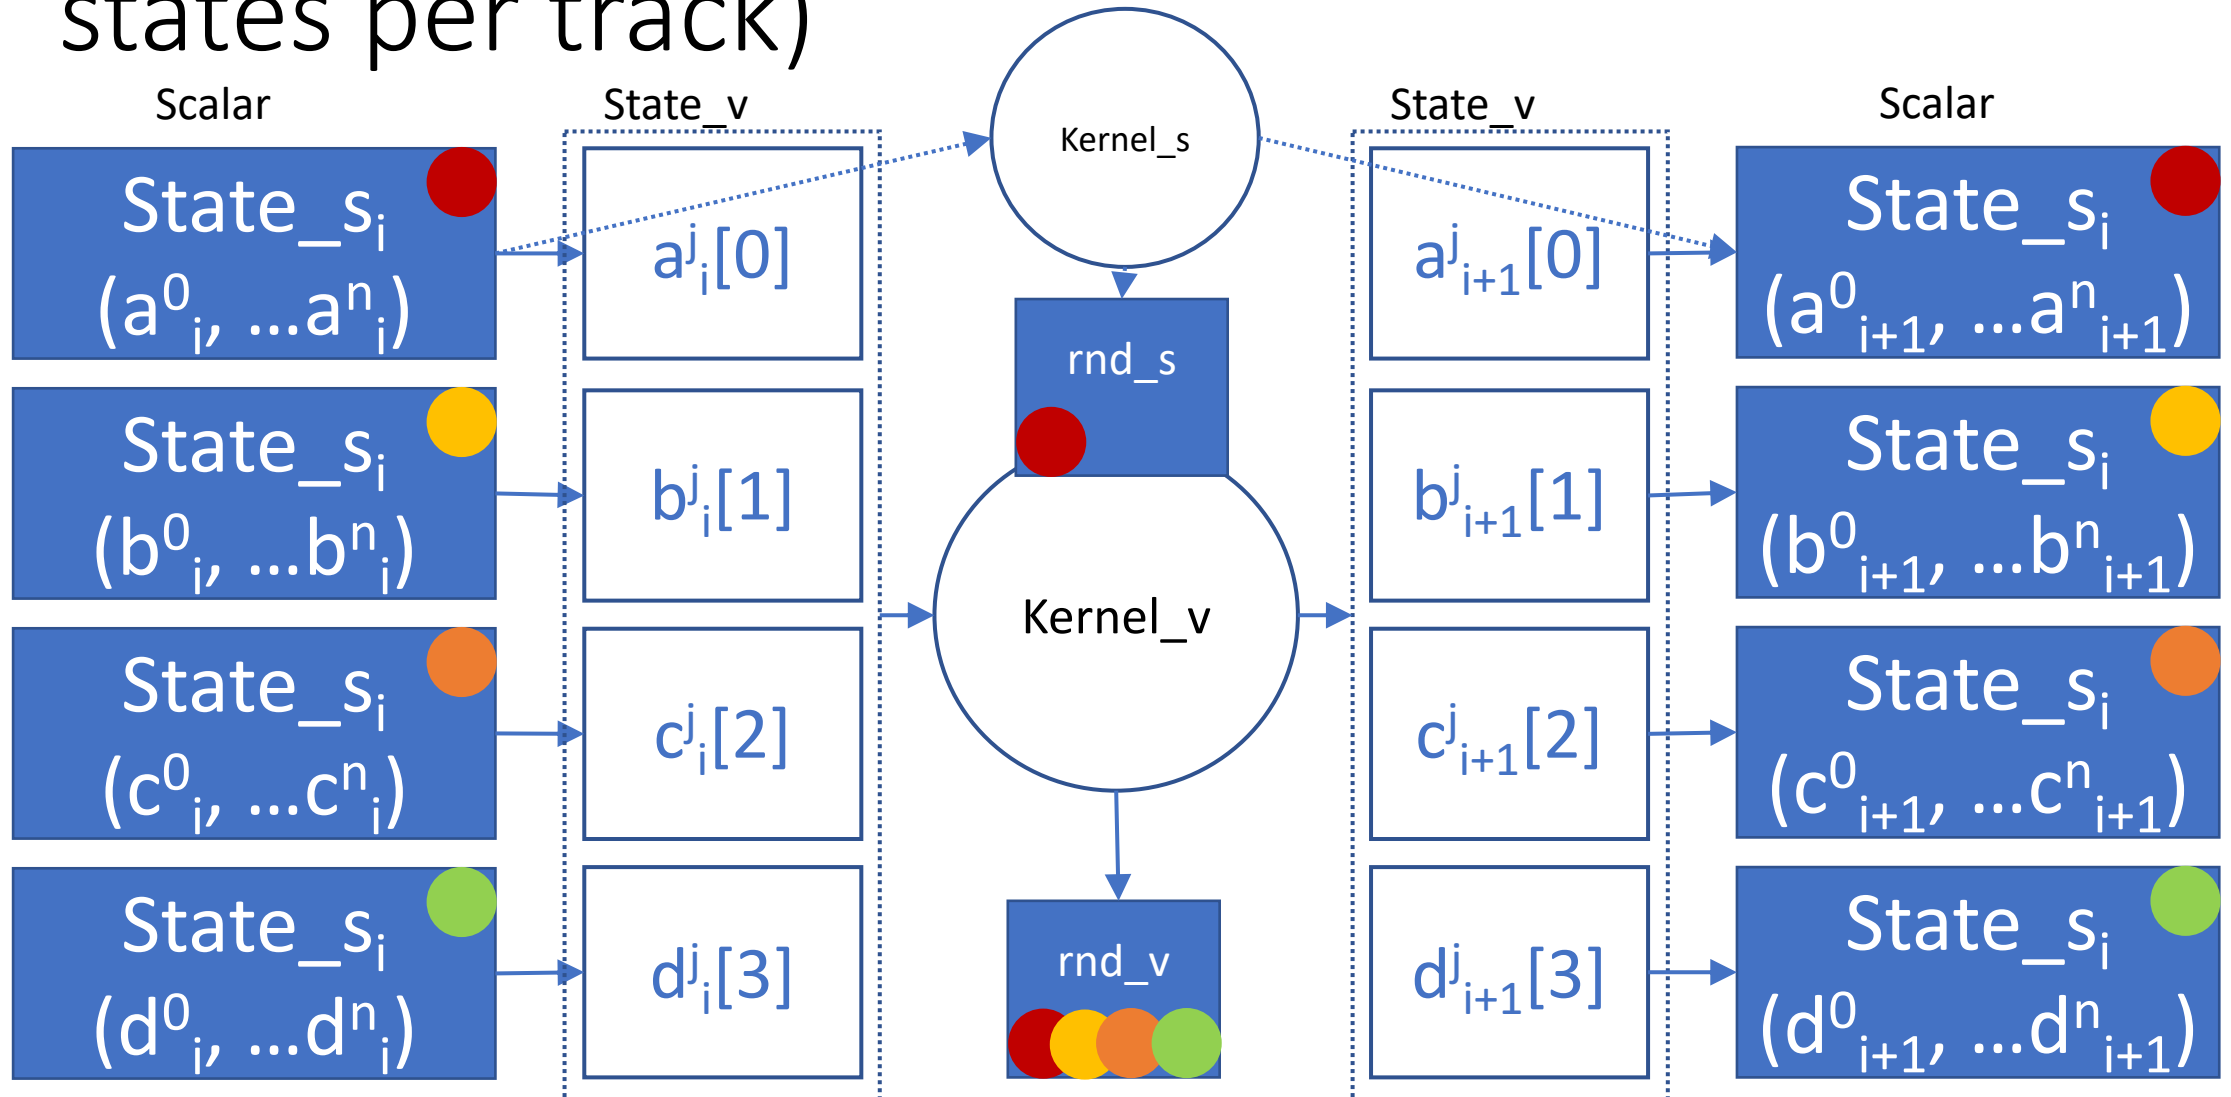
a_v^0

AVX double ($n_{lanes} = 4$)



A vector **State_v** contains n_{lanes} scalar states

Vectorized multi-track approach (scalar pRNG states per track)



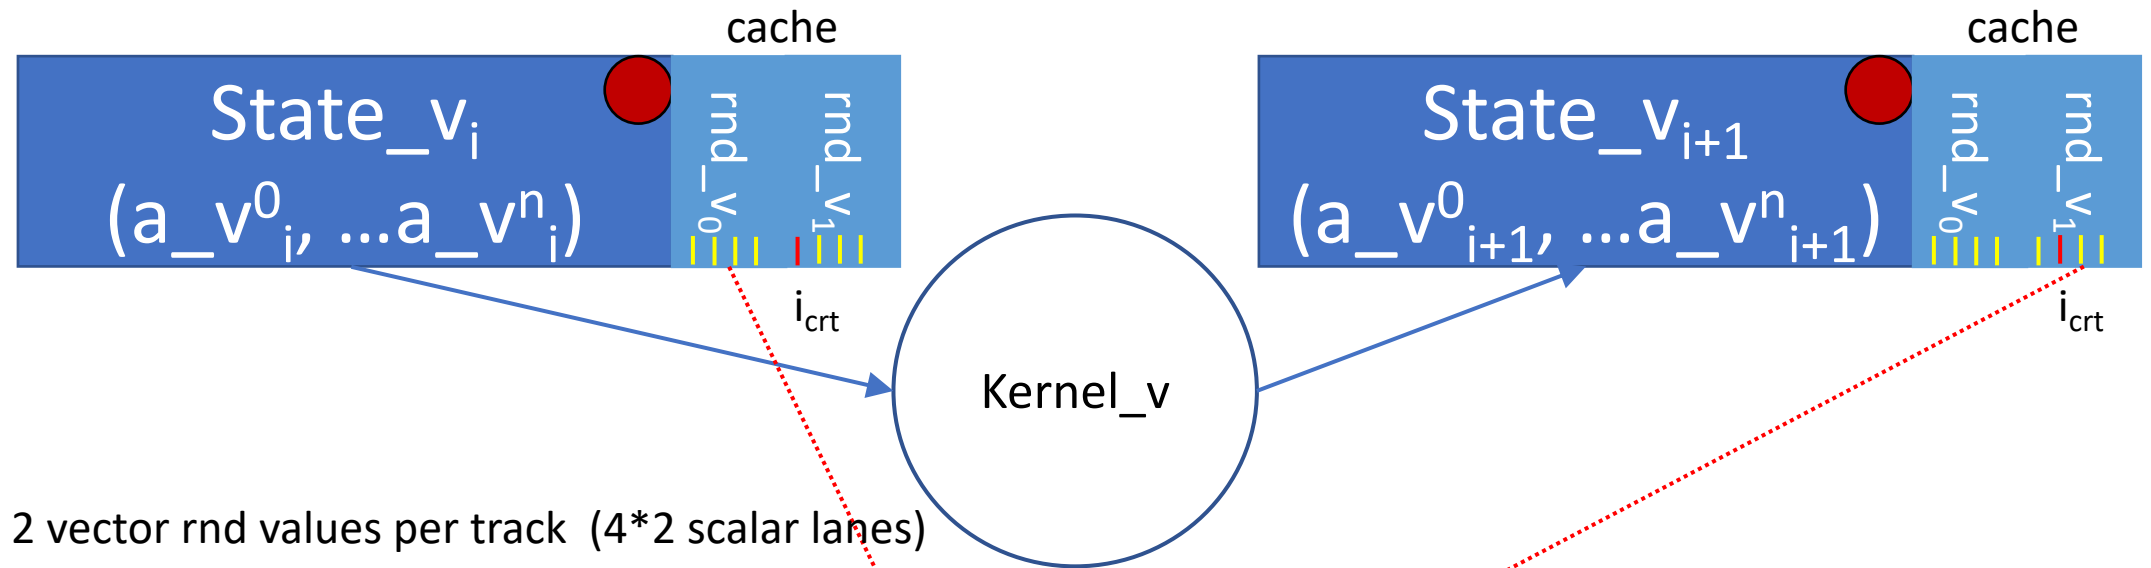
Gather scalar states to individual `State_v` lanes

Scatter lanes into original track scalar states

Remarks

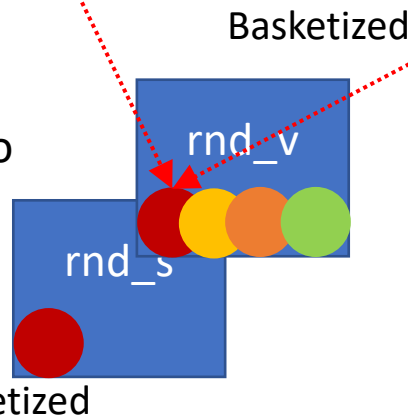
- Easy way to implement the proxy
- Memory footprint: $N_{\text{tracks}} * \text{sizeof}(\text{State_s})$
- $4 * \text{sizeof}(\text{State})$ bytes copied per track
 - track state -> proxy memory -> vector register -> proxy memory -> track state
- Rejection has to be called in single track mode
 - Values computed in vector mode may be dropped for some lanes otherwise, compromising reproducibility
- Single track goes scalar

Vectorized caching approach (vector pRNG states per track)



Cache 2 vector rnd values per track (4*2 scalar lanes)
Keep track of current delivered value in cache

Gather scalar value from cached lanes into $\text{rnd}_s/\text{rnd}_v$ in case: $(i_{\text{crt}} \% 4) > 0$



Call vector kernel to advance the state and refill the cache in case: $(i_{\text{crt}} \% 4) = 0$

Remarks

- More complicated implementation of the proxy
- Memory footprint: $N_{\text{tracks}} * \text{sizeof}(\text{State_v}) + 2 * \text{sizeof}(\text{Double_v})$
- $2 * \text{sizeof}(\text{State}) + \text{sizeof}(\text{double})$ bytes copied per track
 - track state -> vector register -> track state + 1 rnd cached
- All calls to the kernel done in vector mode
 - Run-time benefit in the scalar case
- More complex initialization per track
 - Need N_{lanes} sequences per track instead of one

Discussion

- The 2 approaches cannot be inter-changed during the same run
- There are benefits/overheads in the 2 approaches
- In my opinion we need an implementation for both types of proxies
- Where to implement the proxies?