Partitioning of a Signal Dete Algorithm to a Heterogene Multicomputing Platfor

Michael Vinskus

Mercury Computer Systems, Inc.

High Performance Embedded Computing (HPEC) Conference

September 23, 2003

ʖ The goal of the acquisition system is to detect th presence of new signals in the environment in a manner so that they may be identified and explo

- A homogeneous system based on general purpose processors can be u the problem, but advances in FPGAs allow for a higher processing dens in the channelizer
- Current FPGAs can offer over a 10x computational density improvemen purpose processors for certain applications. Unfortunately, most comm fabrics have not scaled at the same rate

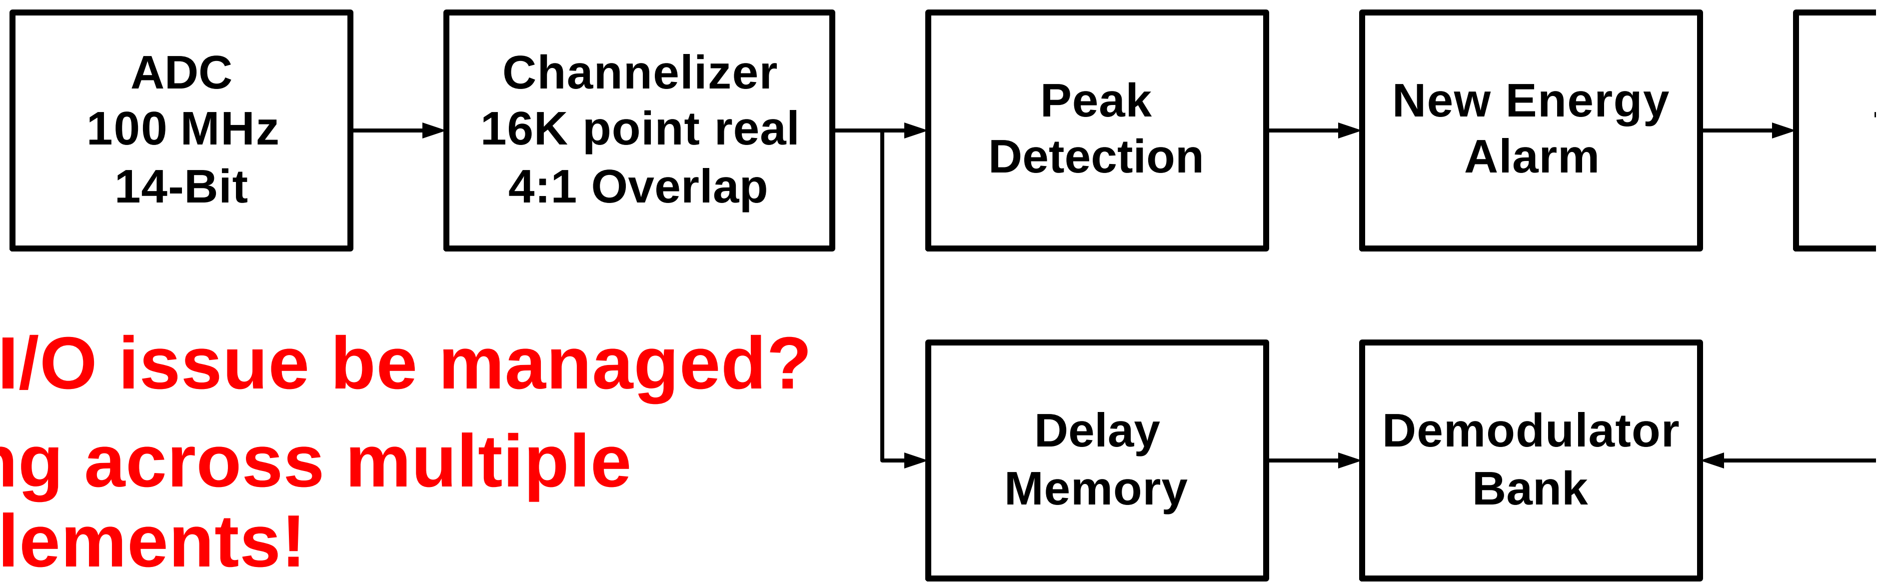

the I/O issue be managed?

oning across multiple ıg elements!

- Typical acquisition systems place the delay memory after the analog to c converter. For demodulation, a digital down converter(DDC) is used to heterodyne and filter the delayed data stream
- When the number of signals to demodulate becomes very large (>100), tl

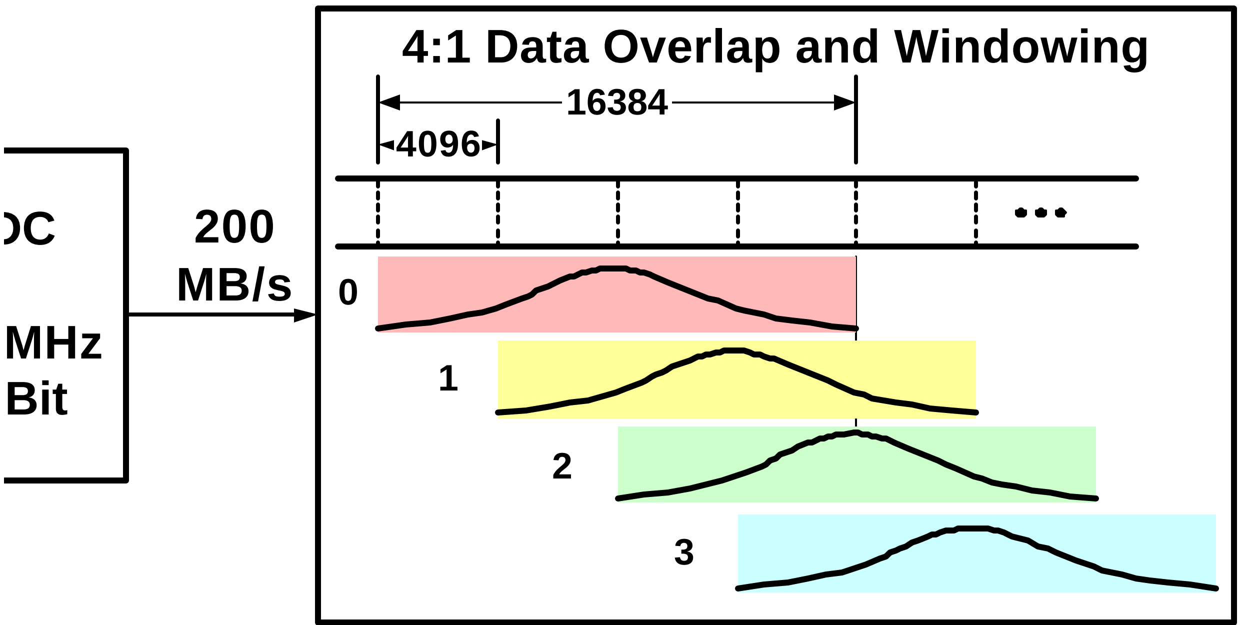

- The channelizer decomposes th sample stream i frequency chan performing over and windowed, time Fourier tra on the input dat

.meters:

- 16384 point real DFT
- 4:1 overlap and windowing (P = 4)
- 16384 input sample maximum latency requirement
- Output 8192 bins, complex, 24 bits per component (1200 MB/s)

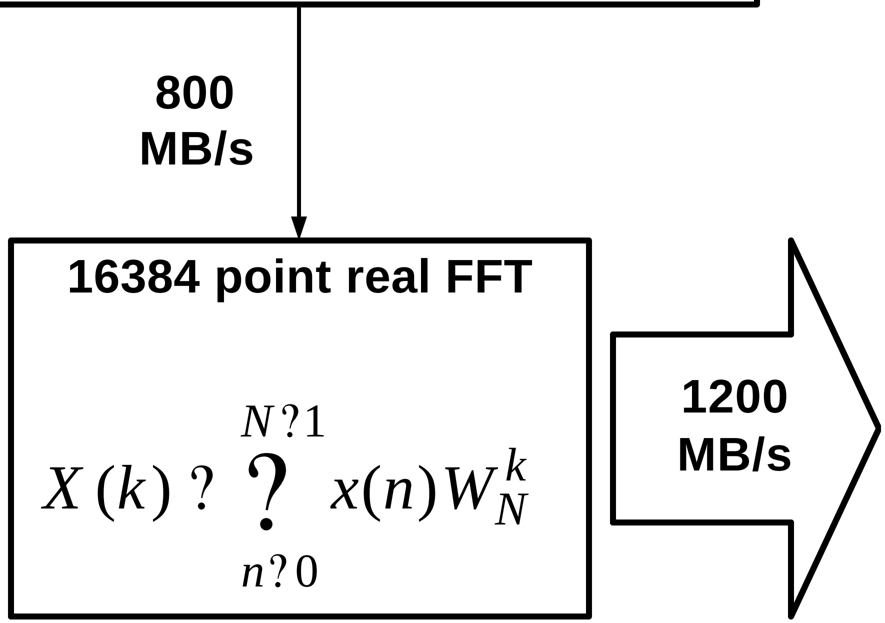

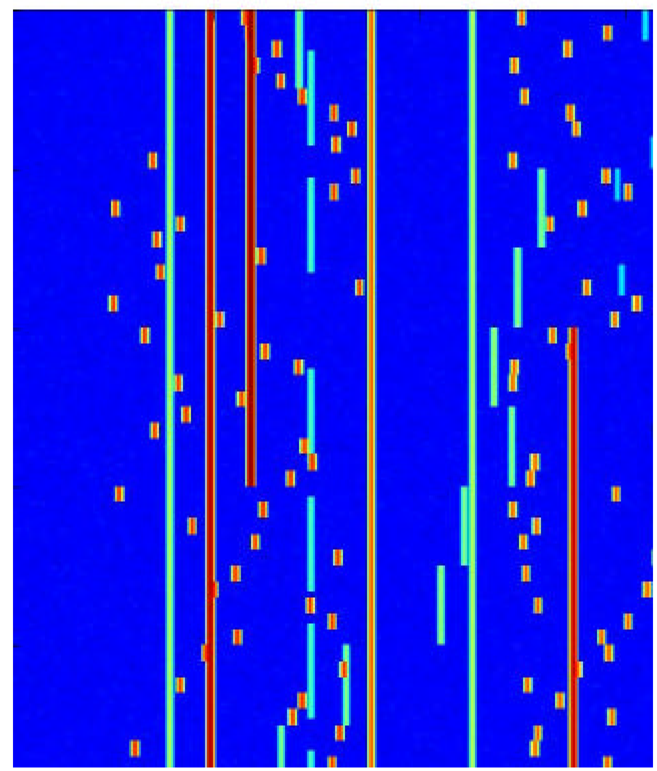

channelizer throughput is greater
most interconnect fabrics can
port

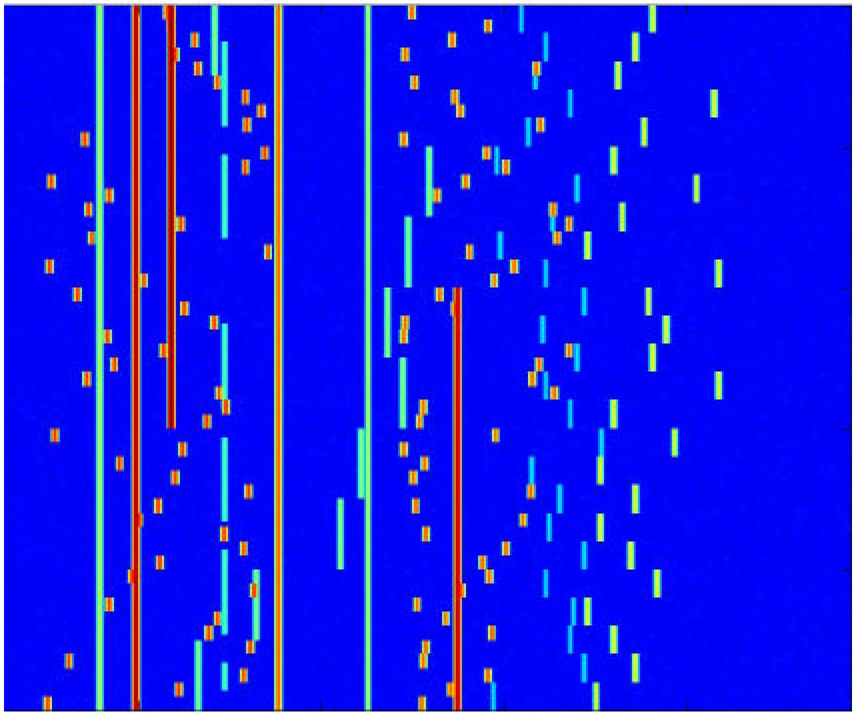

- The new energy alarm d
  the presence of “new” s
  and performs some rudi
  external measurements
- **High-input bandwidth an computational requirem necessitate partitioning**
- Partitioning strategies
  - Commutation
  - Time domain broadcast
  - Frequency division

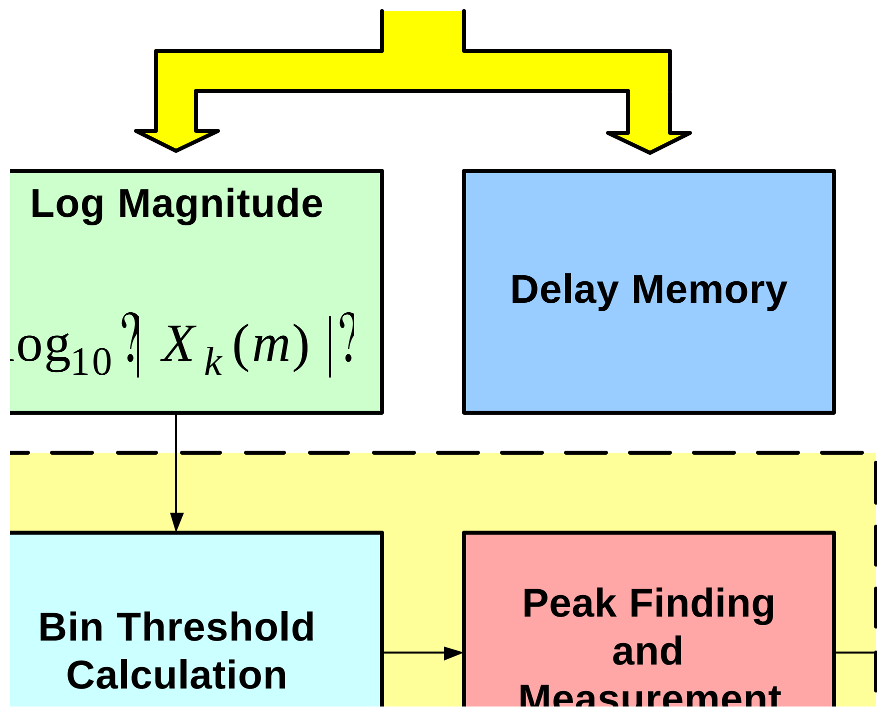

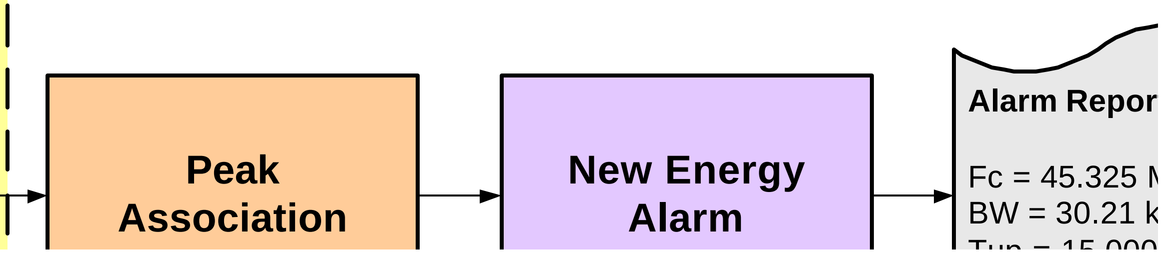

Each processing partition generates the ent frequency sweep for a range of time slices

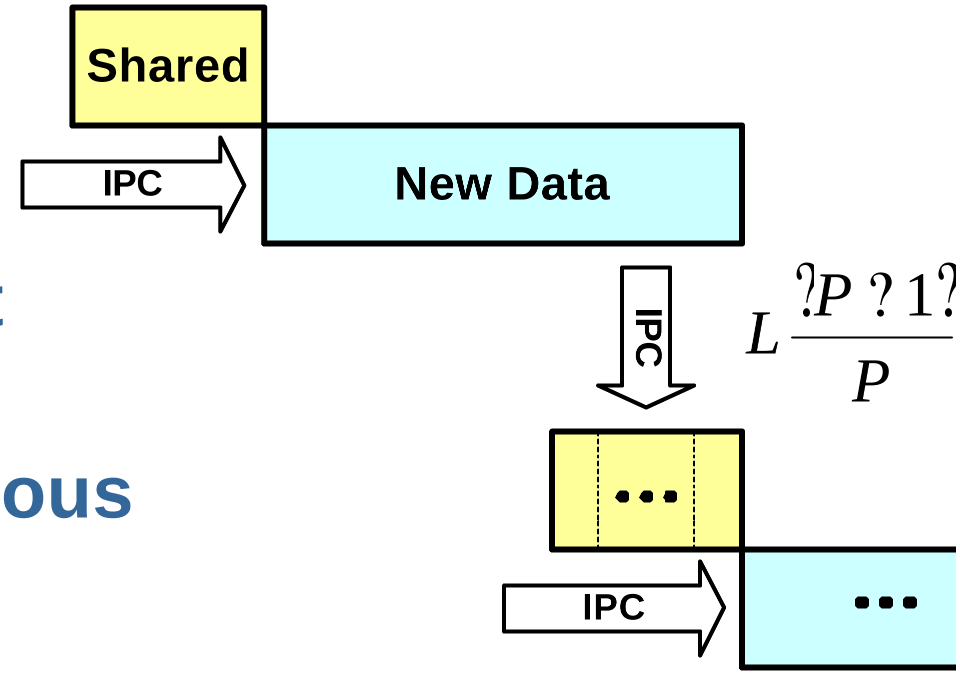

s case, the channelizer is split our partitions

partition processes a contiguous ıent of the input data stream

ırtitioning the problem in this manner, overhead to handle the segment daries is incurred

the system latency ases due to time ısion nature of the

✍ In this example, P-1 old input da need to be received and P-1 bloc to be sent

To meet the one block latency requirement, only block of new data is passed to each partition; 3/4 data comes from the other partitions. This results same data block being transferred an extra 3 time

this example, nput bandwidth eases from 200 s to 800 MB/s. e going the ng way!

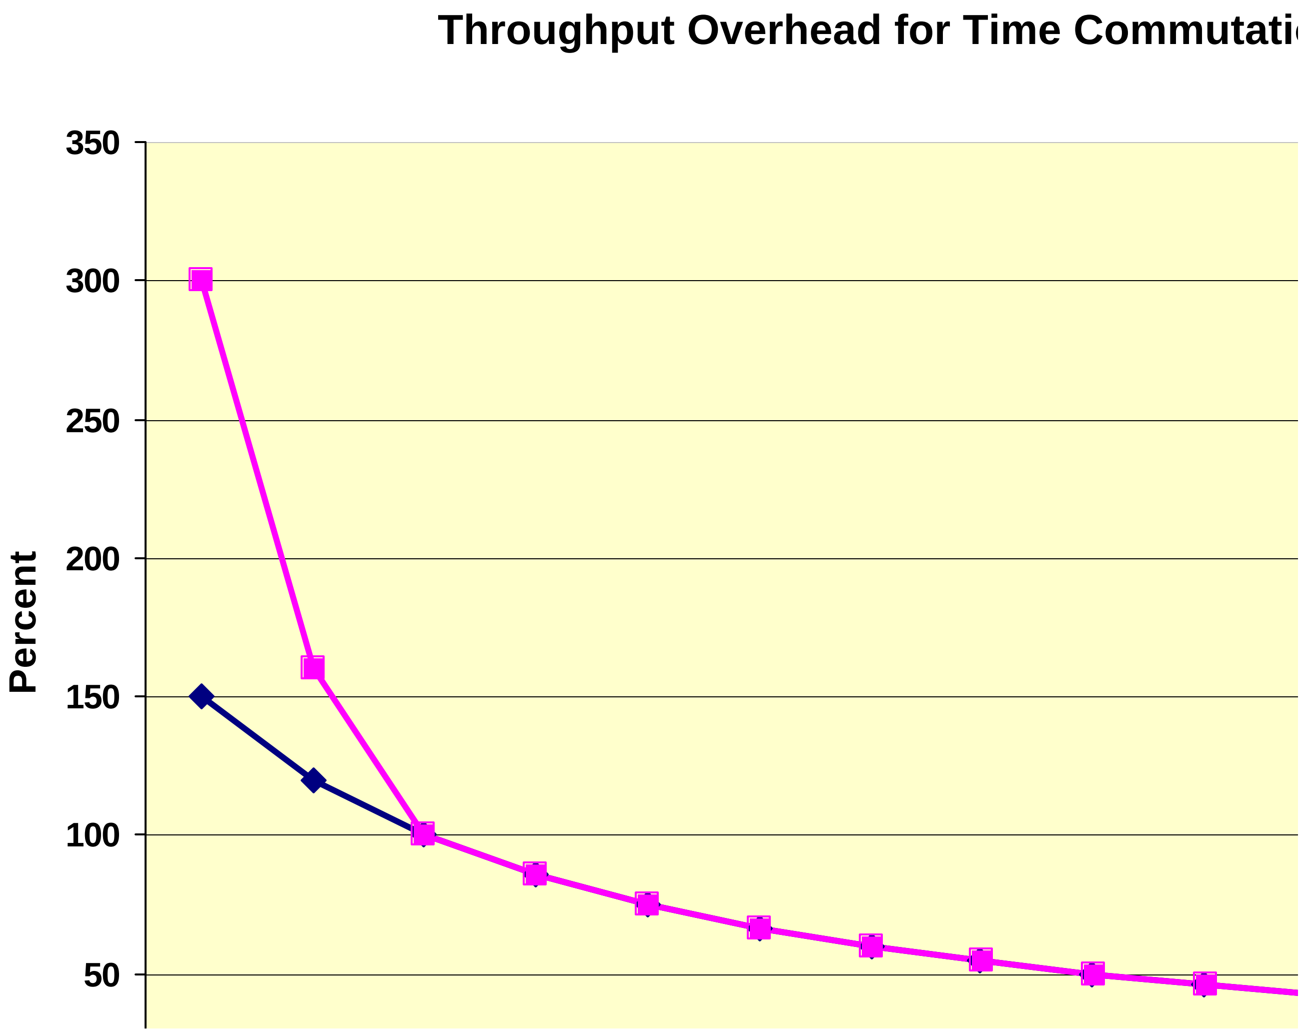

Each partition receives the entire data stream

No extra I/O overhead is incurred

Partitioning does not add latency

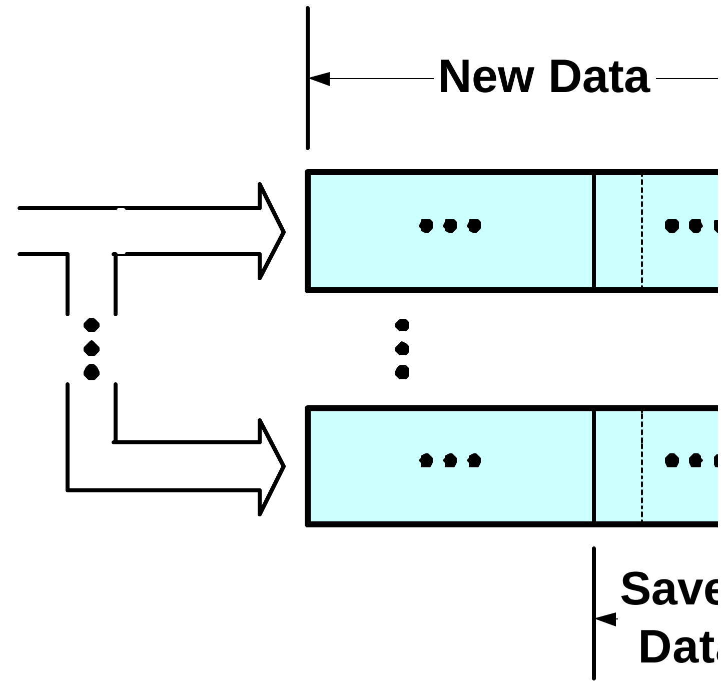

No communication between channelizer partition

Still need to send all of the channelizer output da the detector processors

If there are a sufficient number of detector proce elements allocated and they are located correctly fabric, then I/O bottlenecks can be avoided

Each fabric connection is a 266 MB/s half duplex link to the cross
Typical performance is around 250 MB/s

To accommodate the channelizer output rate, the output stream m
split across multiple connections

This complicates the data flow and fully utilizes the fabric I/O cap
many places

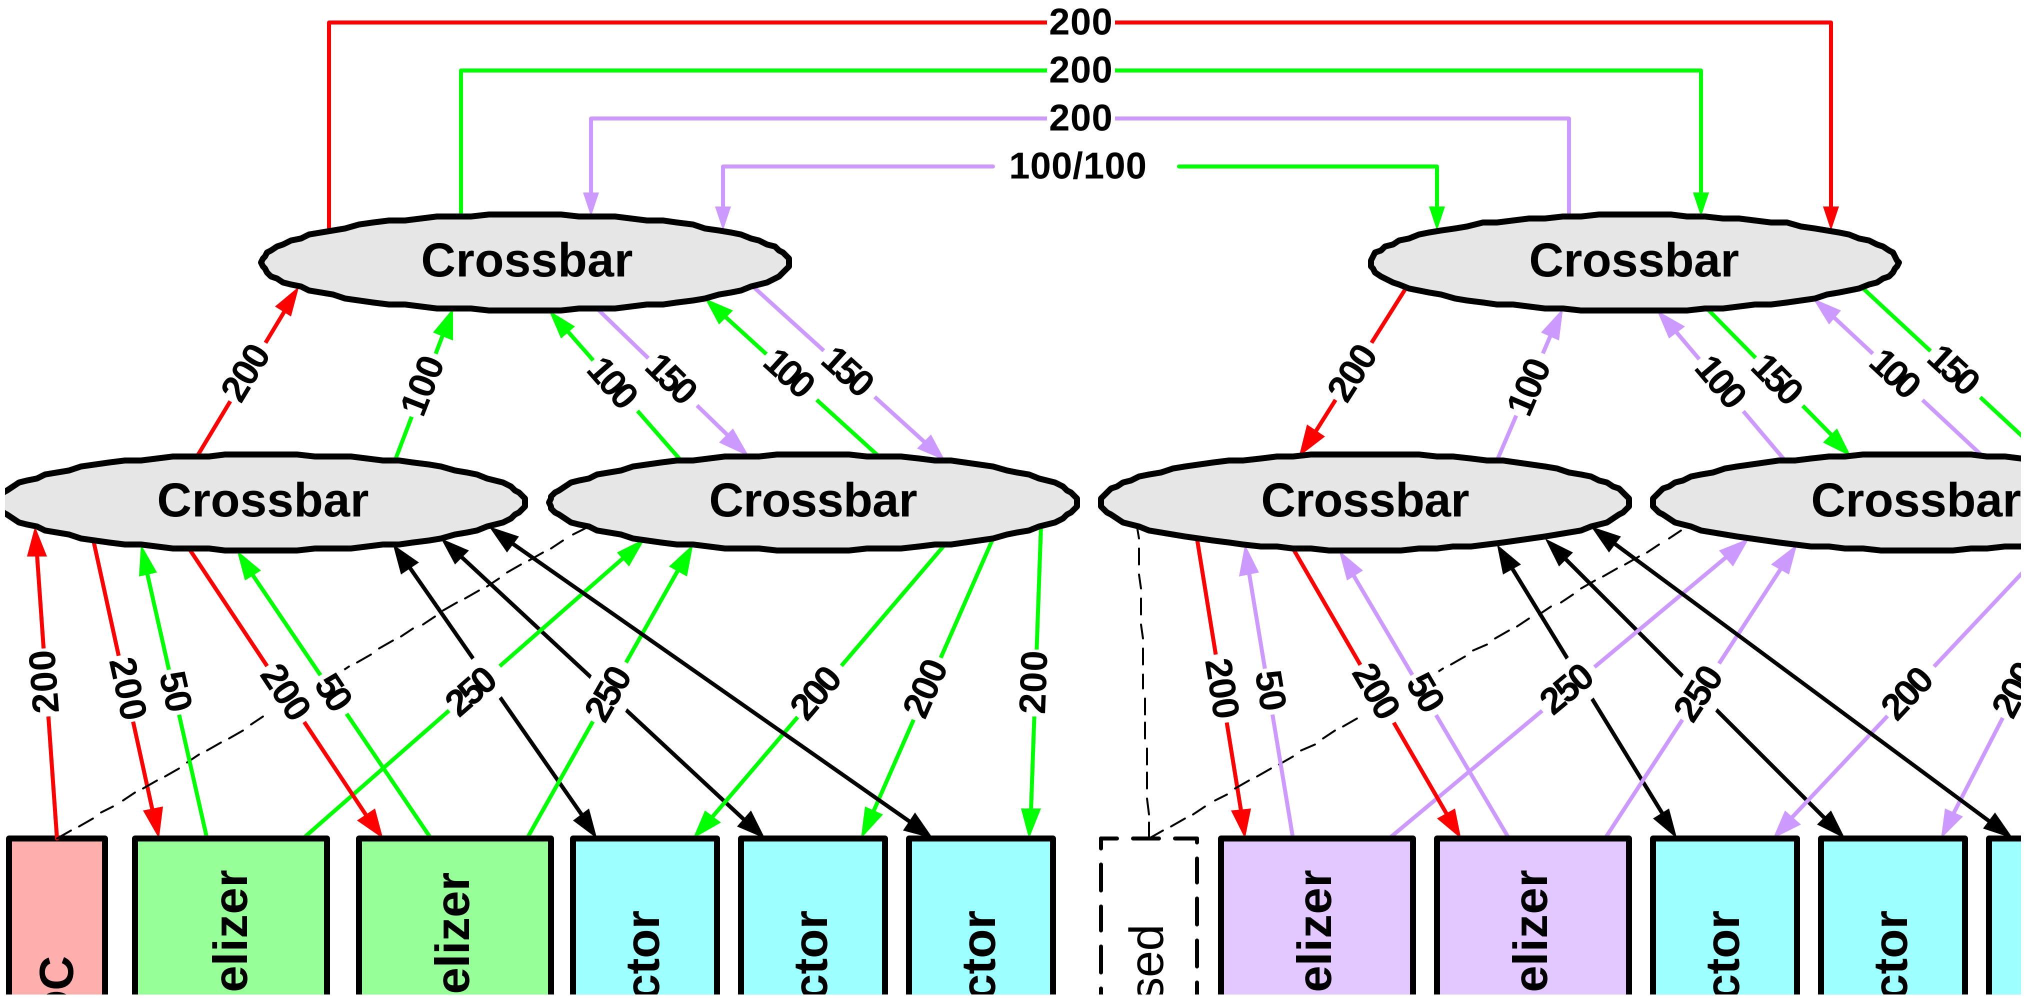

- Each node generates a s
of the frequency bins fo
time slices
- Each partition needs all
time series data for eacl
sweep

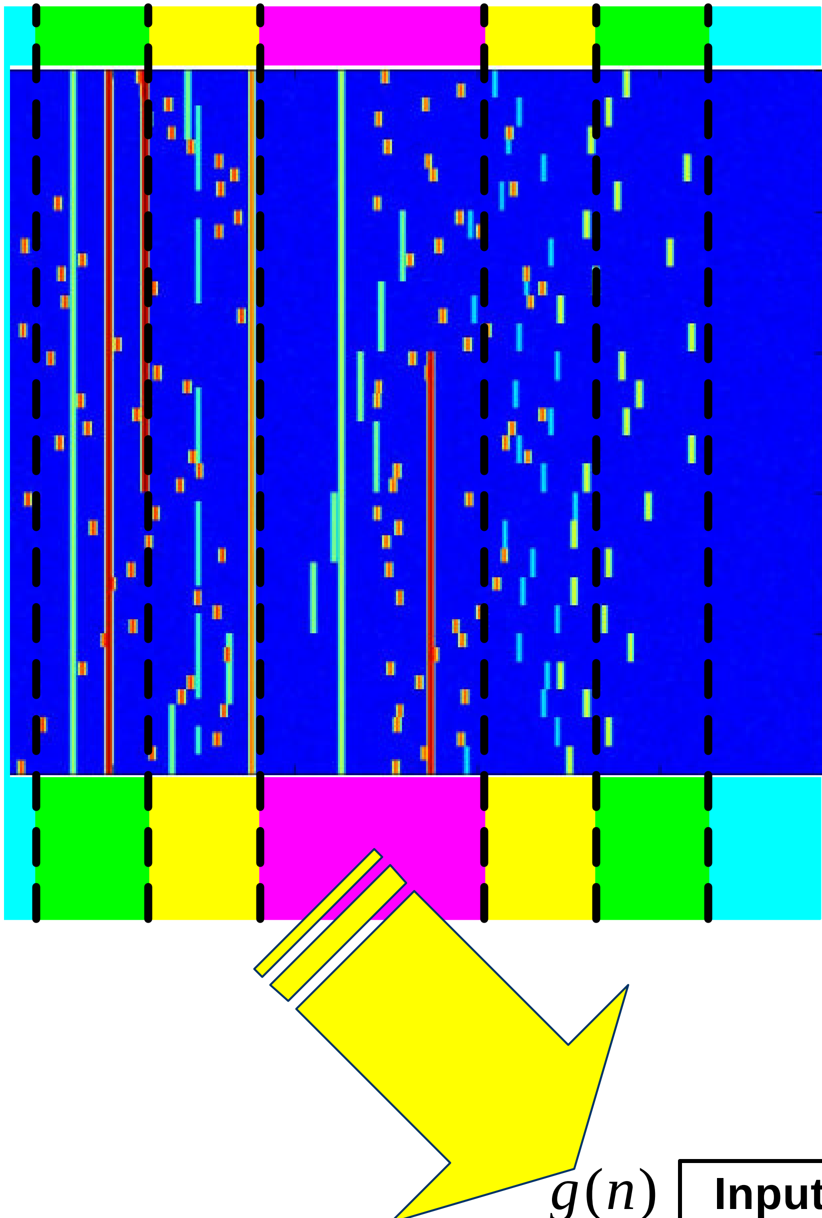

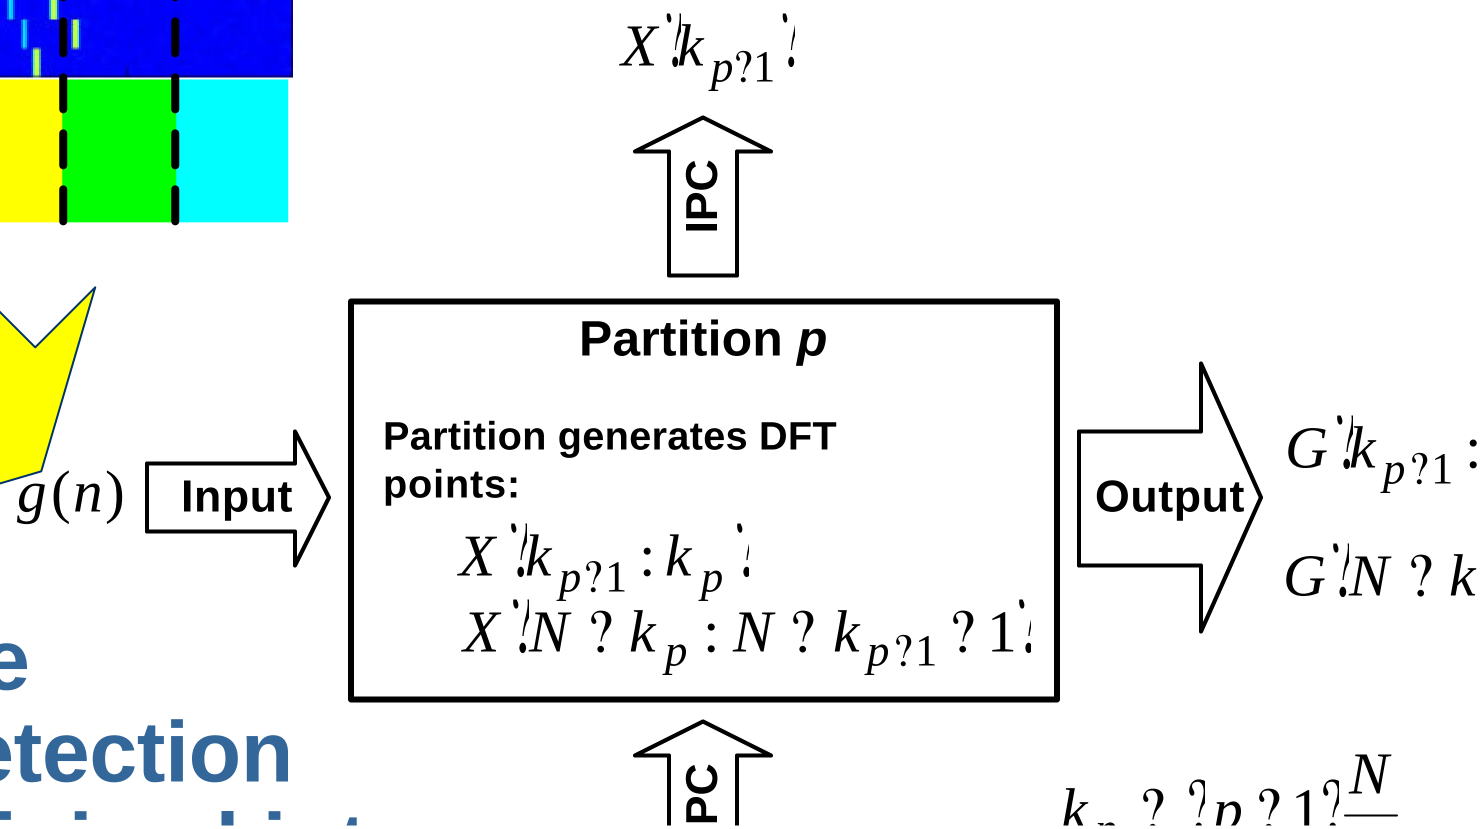

o pipeline the
nelizer and detection

Each partition performs the front part of the FFT computations. This is inefficient, but the I/O issue simpler

Uniform communication between partitions at par boundaries only

- Single element communication between channelizer partitions
- Similar communication between detector partitions!

So why the unusual allocation of the frequency bi

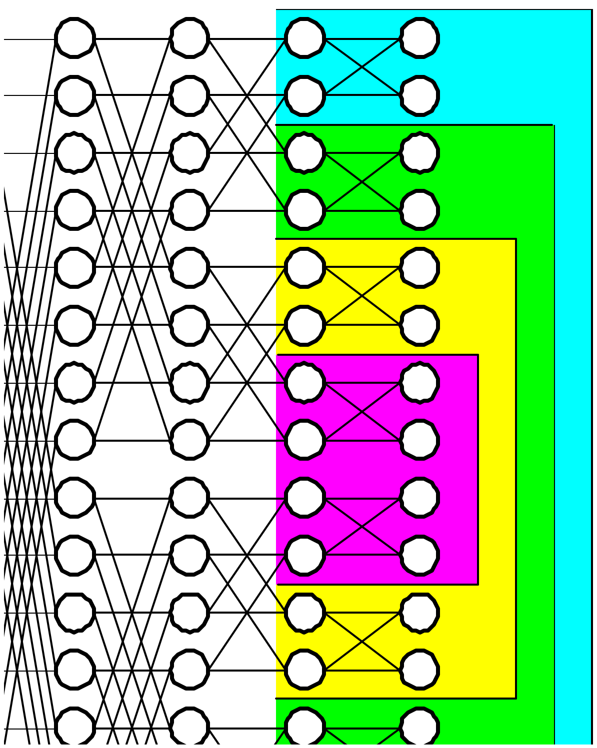

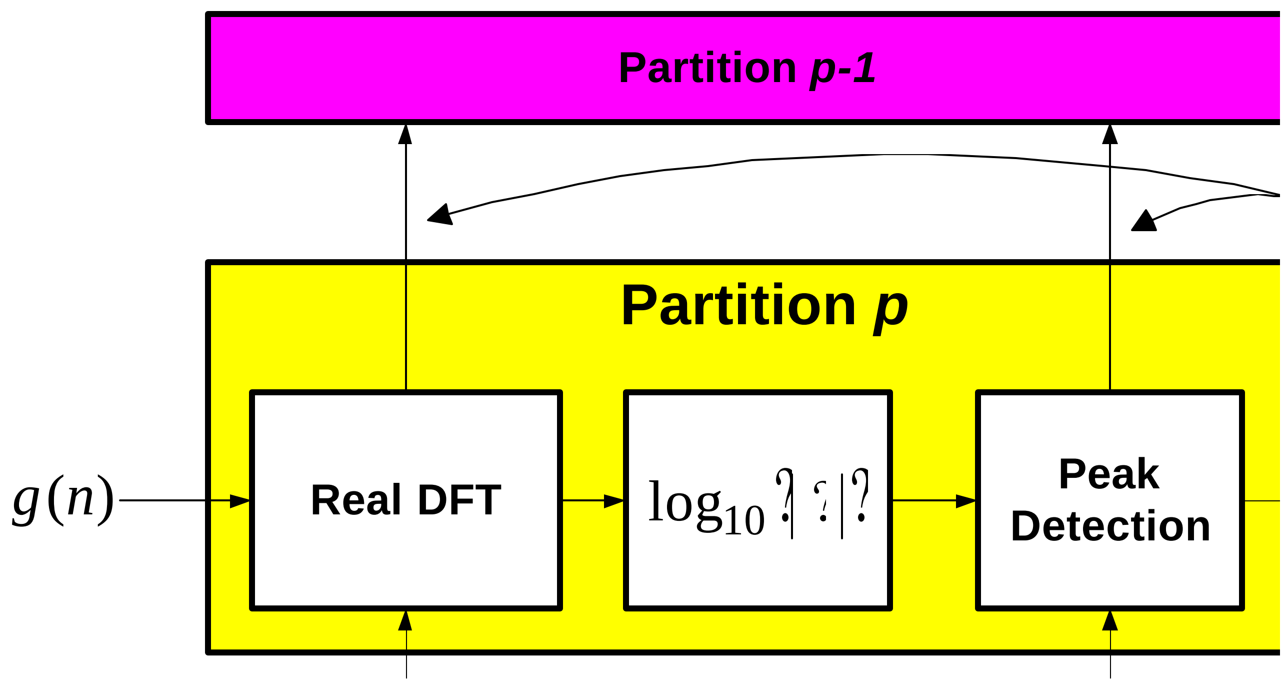

The algorithm exploits the efficient computation of the D 2N-point, real-valued, input sequence with an N-point co transform

Since the input data is real, only one half of the output s needs to be computed

$$G?k?? \frac{1}{2}?X?k?? X^*?N ? k?? j ?W_{2N}^k ??X?k?? X^*?N ? k???$$ **for k = 0, :**

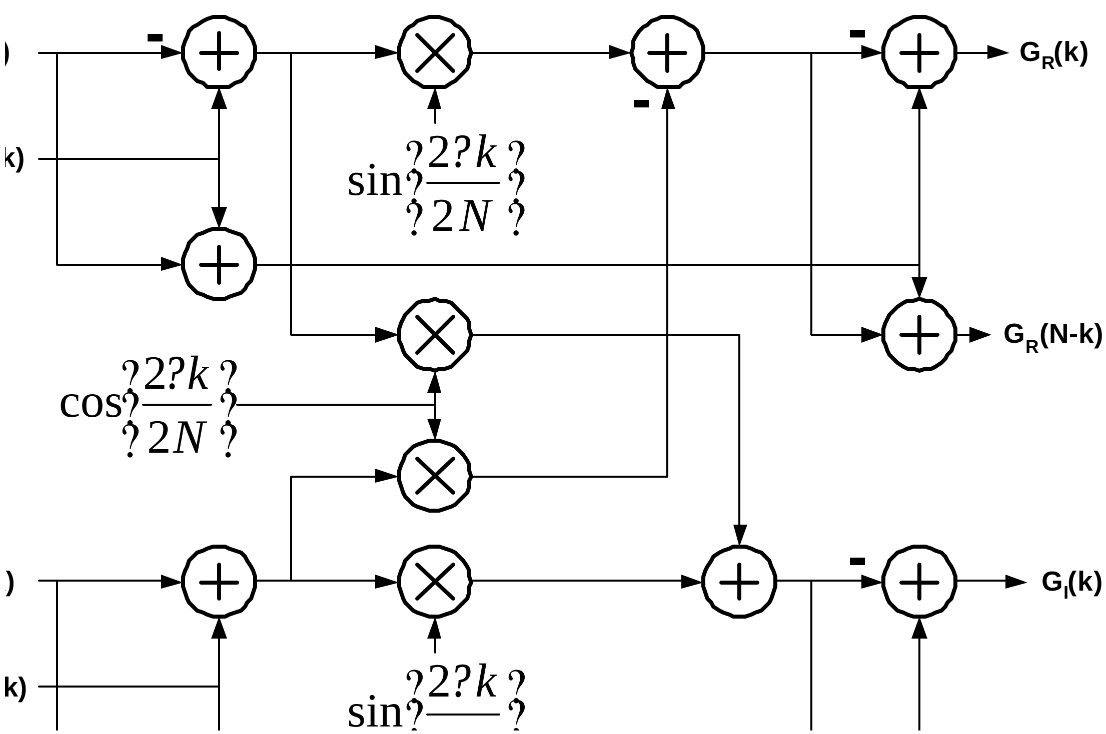

By processing a range X(N-k) frequency pairs, many output points can calculated with only two additions!

Redundant computations are performed in each channelizer partition to reduce the system I/O requirements

Unlike the time partitioning case, most of the I/O movement occurs locally in the fabric

anges
ssing capacity
duced I/O
lexity

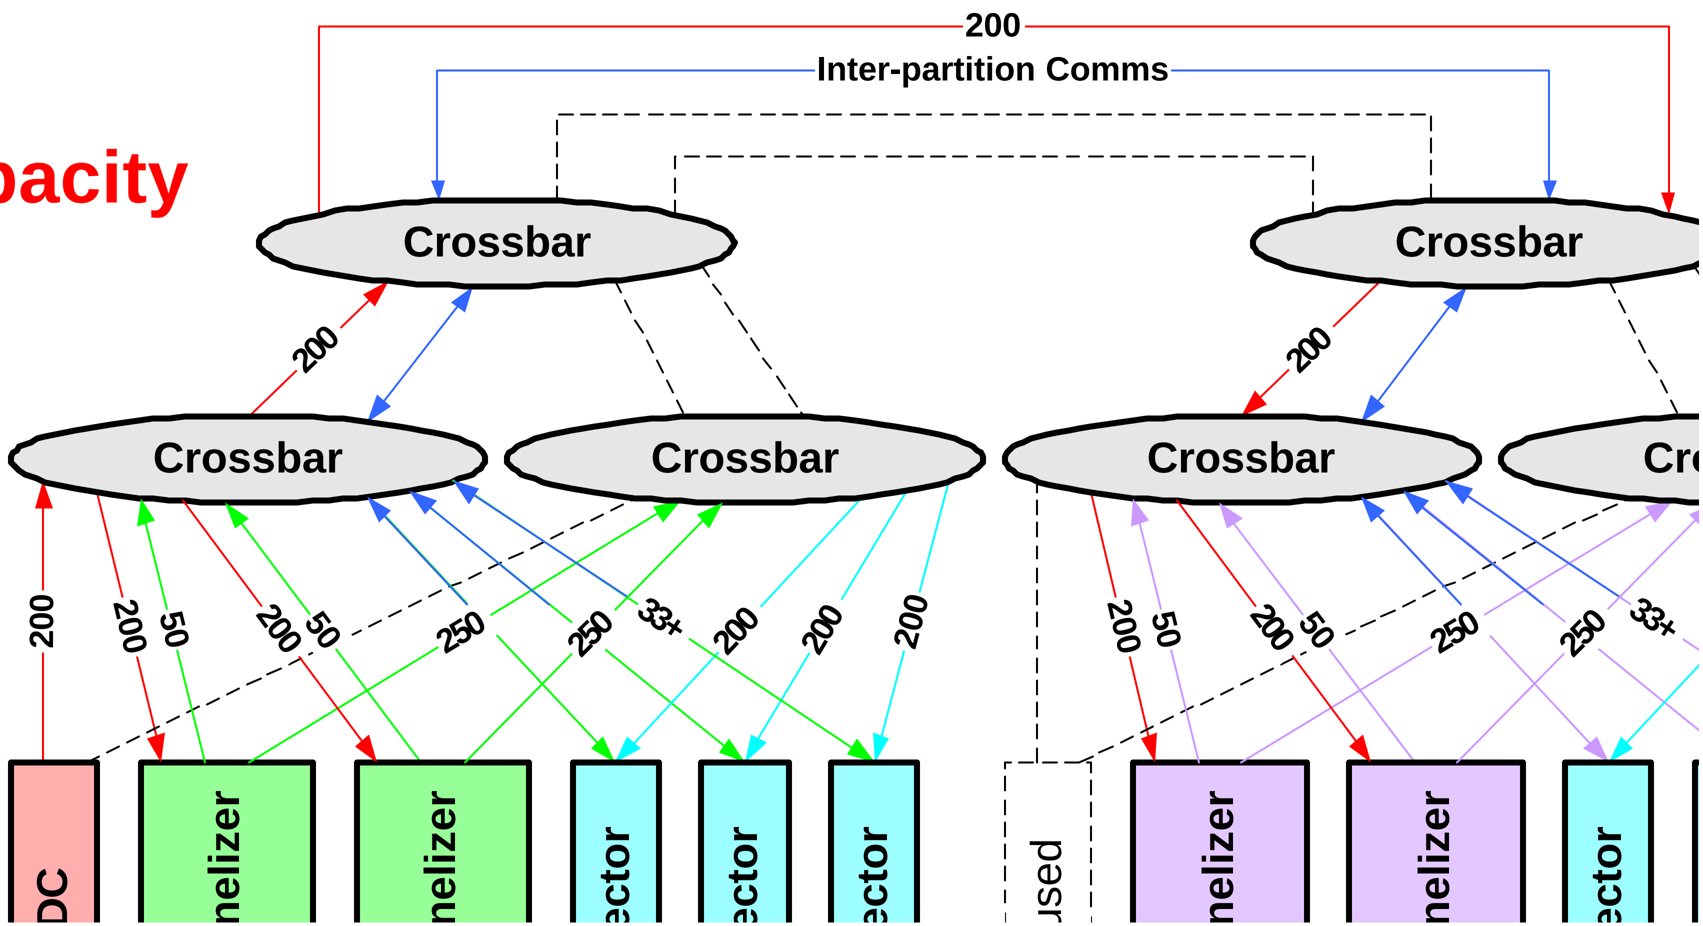

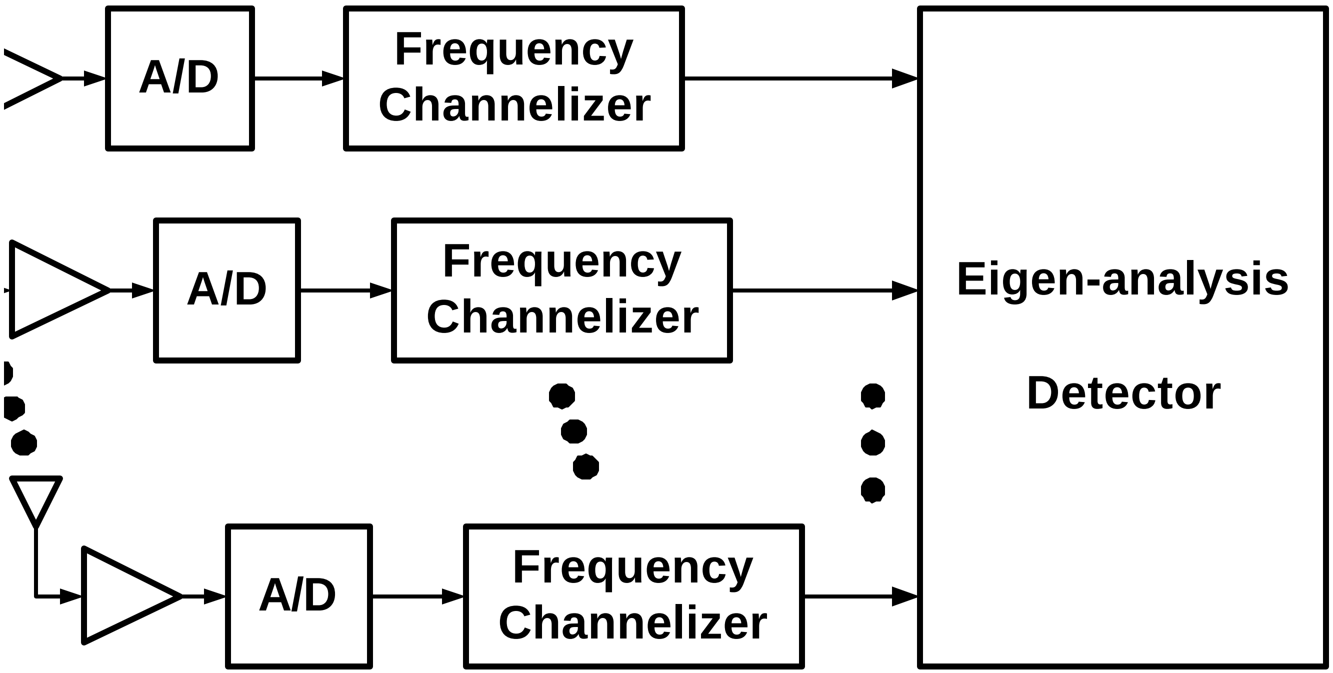

- Adds an extra dimension to t problem
- Typical antenr consist of 4 - 8 elements

letection algorithms are different than the single a . Typically, eigenspace methods are employed

- e channelizer is similar between the single and multi-antenna cases
- tection processing requires the time series data for each frequency bin
- g magnitude computation is not required
- genvalues and eigenvectors are computed for each bin, across all antennas

nown in the previous example, frequency domain tioning has advantages when I/O bound conditions