AMSKepler
ENGENHARIA DE SISTEMAS

High performance satellite image processing with GPUs

Frederico Liporace
Software Development Director
liporace@amskepler.com

Roadmap

- Problem
  - CBERS imagery automatic registration
- Solution
  - GPGPU
    - **G**eneral-**P**urpose computing on **G**raphic **P**rocessing **U**nits
    - Basic concepts
- Results
- New paradigms
- Potential new applications in Mining and Exploration

AMSK and the Brazilian Space Program

- Ground station software development under contract from INPE
  - Since 1998
  - CBERS 1/2/2B, Landsat 1/2/3/5/7, MODIS
  - Already prepared for CBERS 3/4
  - Being extended for Amazonia 1
- All processing chain, from data recording up to image dissemination
  - Includes automatic geometric correction through GCPs and orthorectification

Geometry issue: CBERS 2B Along Track Error

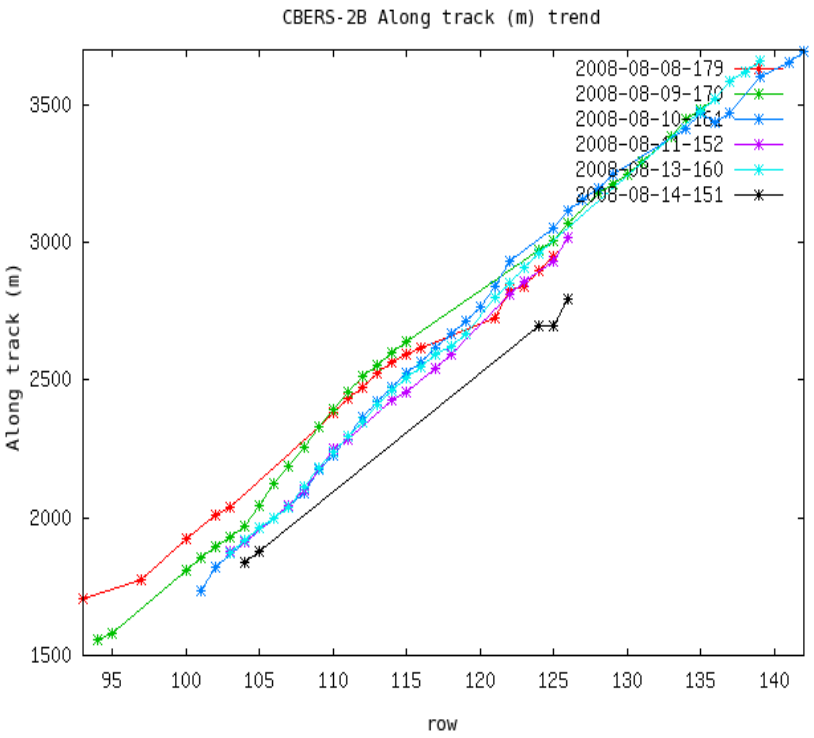

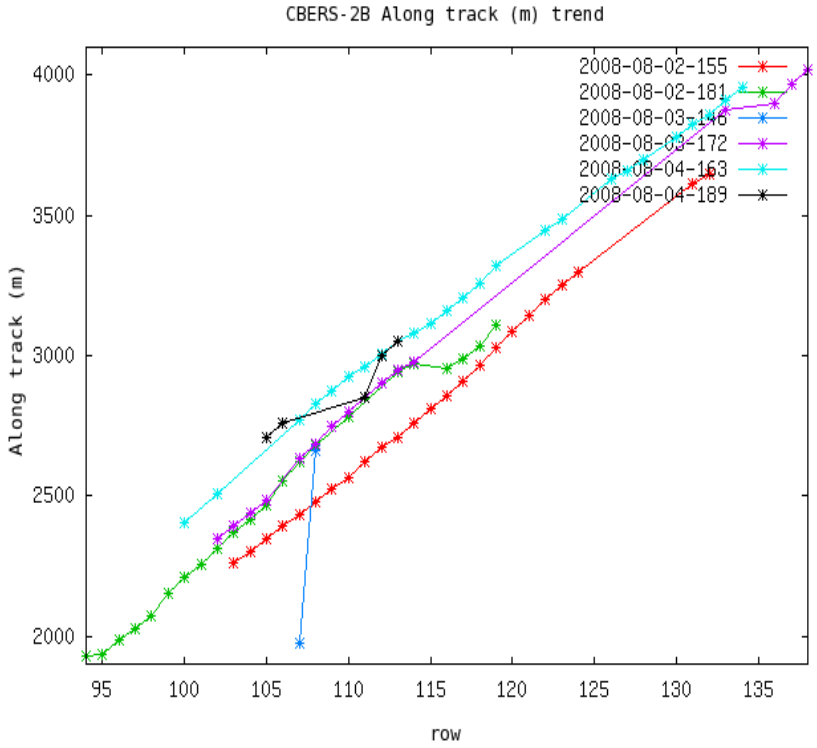

Discussion

- Probably due to wrong attitude information
- Makes automatic registration harder
  - Big search space to deal with general case, long processing time (around 3 hours for each scene)
- CBERS 3,4 and Amazonia 1 will be better…
  - GPS and improved star trackers into the control loop
- … But we need to:
  - Be robust
  - Be prepared to generate mosaics and changes detection products for Amazonia 1
  - Improve the archive quality
- Approach: GPGPU

GPGPU

- **G**eneral **P**urpose computing on **G**raphic **P**rocessing **U**nits
  - GPUs originally developed for high performance graphics applications (games, CAD, virtual reality)
  - Low cost, massive parallel processing environment
    - =~ 400 cores against 4-6 typically found on CPUs
  - Lately the graphic card manufacturers have been pushing SDKs for general purpose programming on GPUs
    - CUDA (NVIDIA)
    - OpenCL (Consortium)
- Ground system processing is being extended to use GPGPU
  - Start with Normalized Cross-Correlation (NCC) computation

Sequential NCC

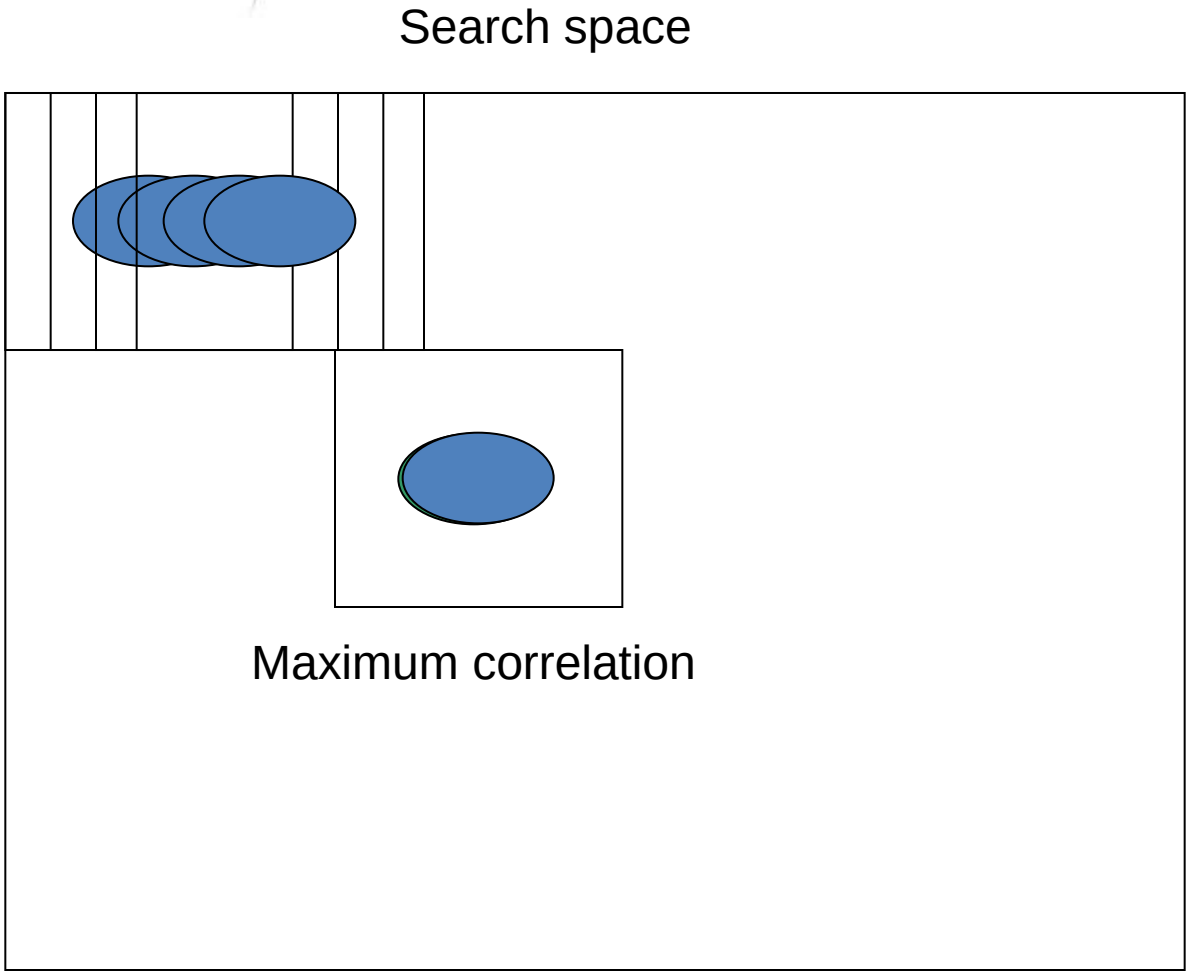

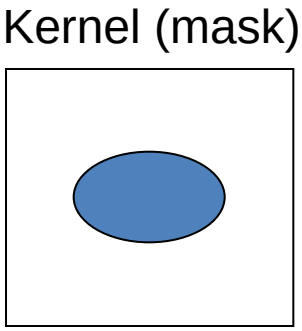

- Kernel center is a control point
  - Landsat GLS
- NCC computed for all search space, highest value is match candidate

Parallel NCC

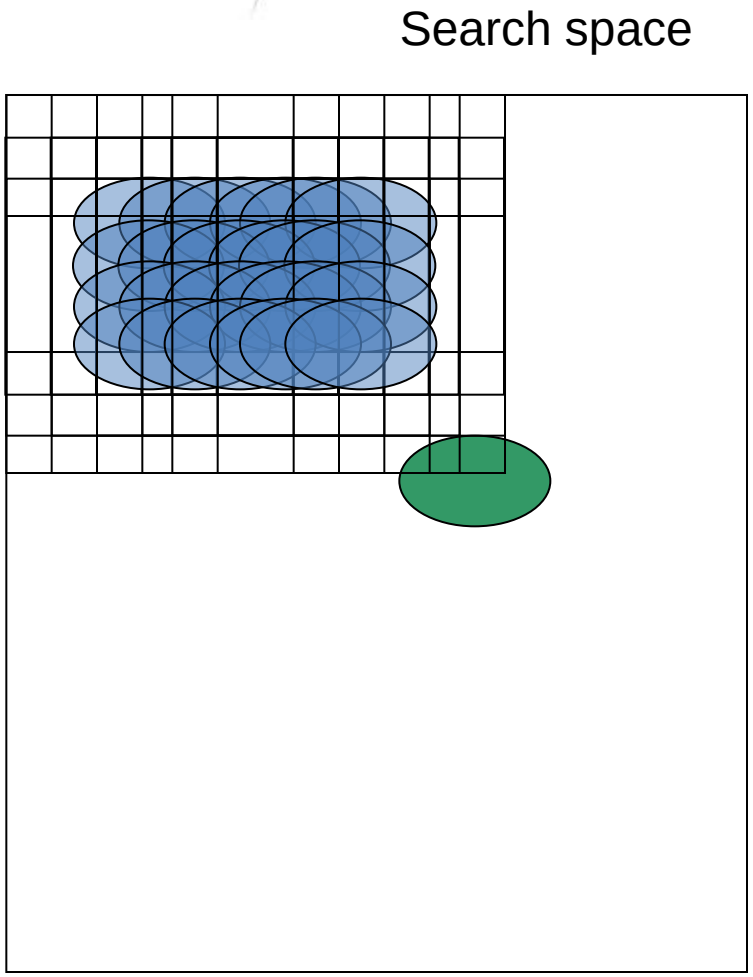

- Data parallelism
- One thread is created for each candidate position
  - Each GPU core handle multiple threads
- “Easy” parallel algorithm
  - No communication or synchronization between threads
  - Single image transfer to the GPU memory, avoid memory transfer bottleneck

AMSKepler

Results

- 3000 x 3000 pixels search space
- 33 x 33 pixels kernel
- NVIDIA GTX 560 Ti
  - 384 cores @ 1.64 GHz
  - Cost: 240 USD
- Performance comparison
  - Sequential: 33 seconds (Intel Xeon E5410 @ 2.33 GHz)
  - GPU: less than 1 second

Nearest Correlation Interpolation (NCI)

- Idea occurred during the development of an automatic image registration assessment tool
  - Image to reference and band to band registration
- Residuals computed in a sub-sampled grid
  - What if computed for **all** pixels?

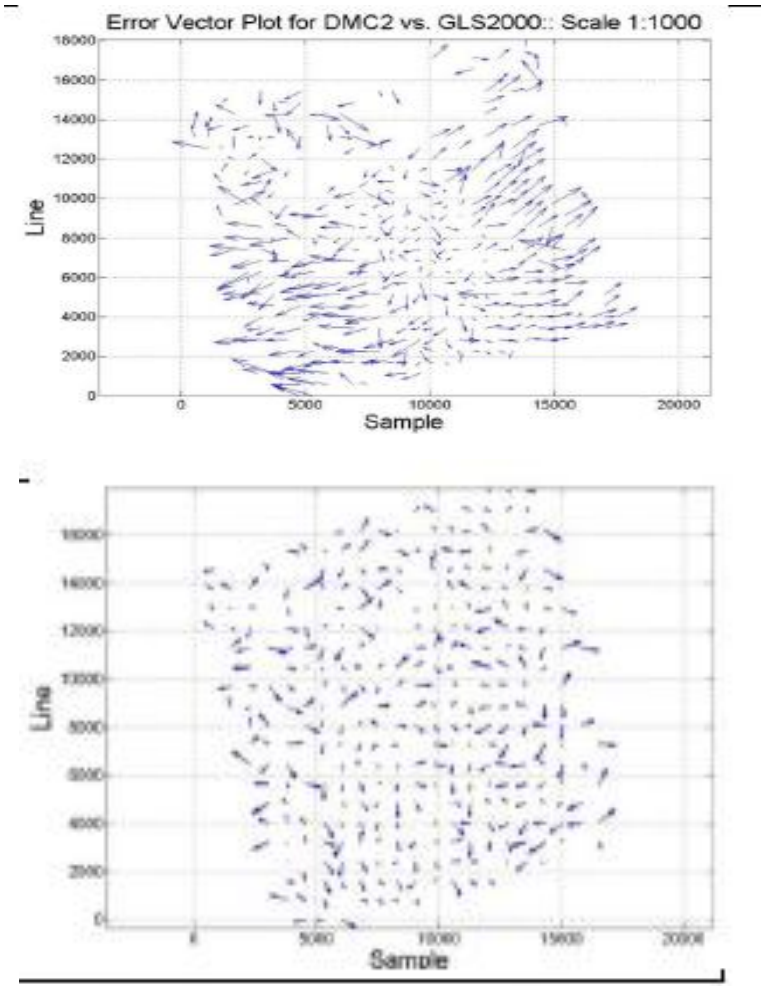

Source: Ajit Sampath, "Geometric Accuracy Assessment of Data from UKDMC-2, SPOT 4/5 and RapidEye" April 2012.

NCI (cont)

Threshold = 0.7

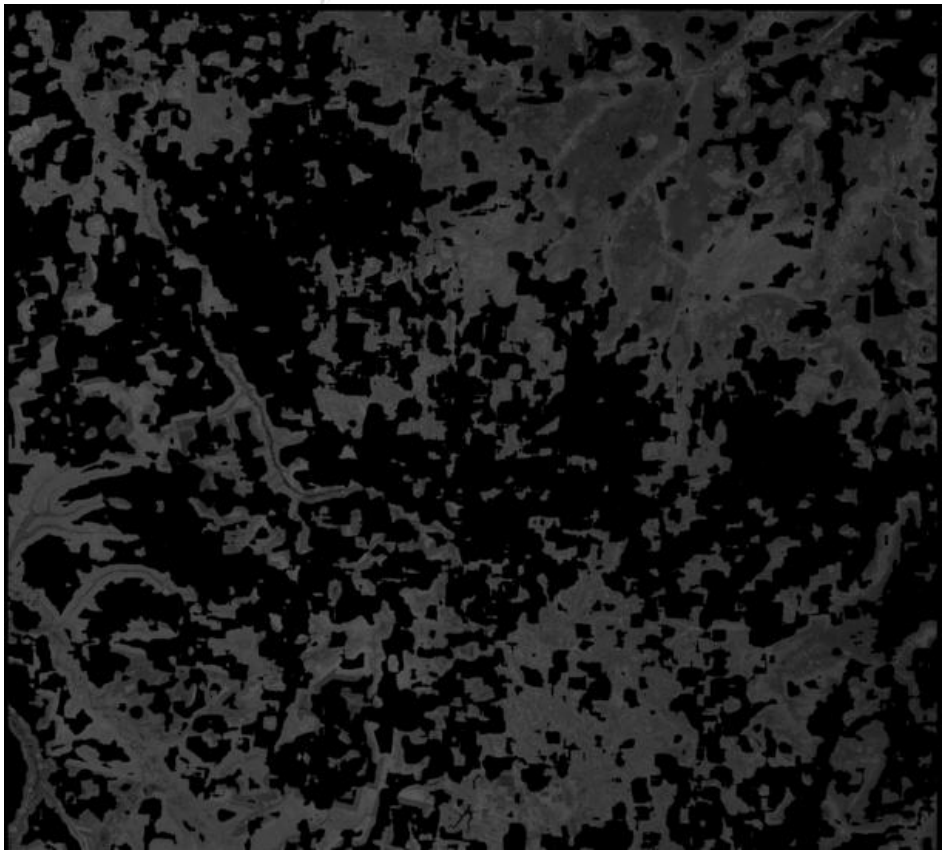

Threshold = 0.1

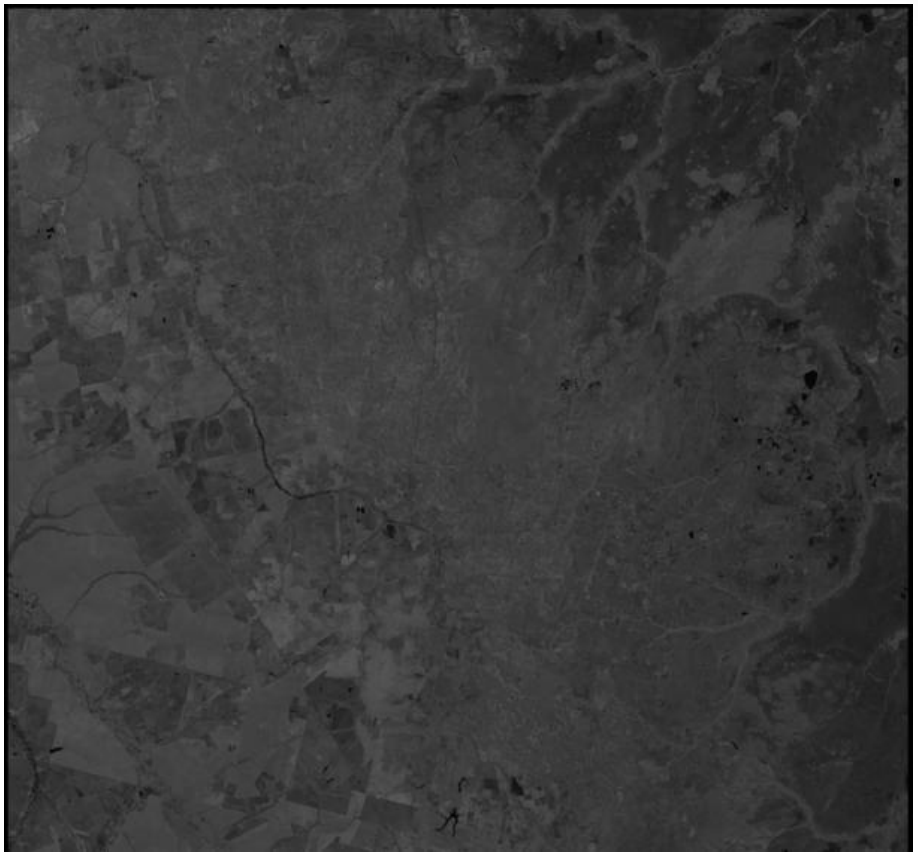

4908 x 4669 image, 11x11 pixels search region, 61x61 pixels correlation mask.
Execution time on GPU (NVIDIA GTX 560 Ti): 4.3 minutes

Potential Applications on Mining and Exploration

- Those depending on registration
  - Changes detection, mosaic generation, etc.
- Any problem for which computing time is currently a bottleneck
  - Image classification of hyperspectral data
    - Test case with Hyperion data being conducted with CVRD, classification procedure (ISOData, K-mean) currently being reported as lasting for days.
- Each case must be individually analyzed
  - Not a panacea
  - It is crucial to develop an efficient parallel algorithm

AMSKepler

Thanks for you attention!