



Meeting timeliness and throughput constraints in IASI L2 operational processing at EUMETSAT

Tim Hultberg



Contents

Brief overview of IASI processing in the EPS ground segment

Optimisation strategies (code and algorithm)

- **Fast radiative transfer model**
- **Use of reconstructed radiances to keep the number of channels to simulate low**
- **Improved first guess / background to keep the number of iterations low**

Parallelization strategies (past, present and future)

- **In principle easy as IASI L2 algorithm is embarrassingly parallel**

Integration of Forli-CO retrieval code from the EUMETSAT SAF on ozone and atmospheric chemistry monitoring (Daniel Hurtmans, ULB)

Legend

- 10/100 1q VLAN trunk
- 10/100 BaseT
- Ethernet
- Serial Line
- Router
- TELENOR equipment
- Remote Entry
- Boundary of Responsibility
- Layer 3 Switch
- Layer 2 Switch

Don't worry





Product Processing Facilities (PPFs)

All Metop products are generated by PPFs, processing 3 minutes chunks of data - so-called PDUs

Most PPF executables are invoked and exits again for each PDU

- Only exception is the IASI L1 PPF (invoked once per orbit)**

Three ground segment instances (GS3=>GS2=>GS1)

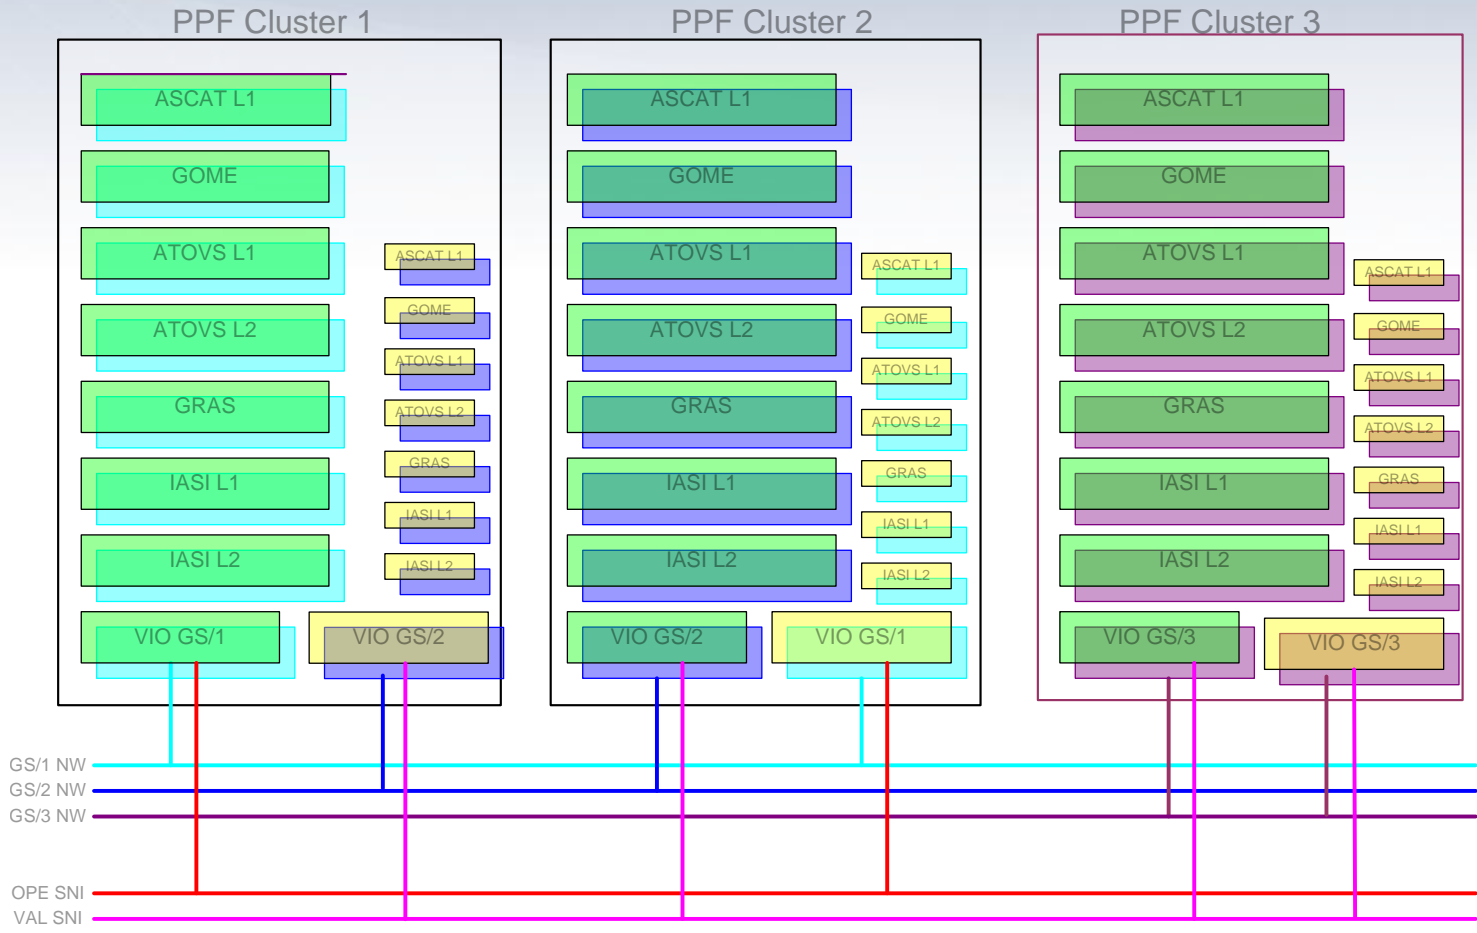
GS1 Operational processing

GS2 Validation

GS3 Testing

PPFs are installed on nodes (virtual machines) on three IBM p575 HPC clusters with 32 4.7 GHz POWER6 processor cores each.

(Operating system is AIX 6.1, 64 bit)





Product Processing Facilities (PPFs)

- **IASI L1 PPF** **One node (4 virtual CPUs, 2Mb Memory)**
- **IASI PCC PPF** **One node (2 virtual CPUs, 2Mb Memory)**
- **IASI L2 PPF** **Three nodes (4 virtual CPUs, 2Mb Memory)**

Per satellite (Metop-A/Metop-B)
per ground segment (GS1/GS2/GS3)
(Nominal/Spare)

Some scope for increasing the number of (virtual) CPUs on the nodes, but not the number of nodes themselves.

- **Integrate PCC PPF into L2 PPF** (PCC processing fast and repeated for noise filtering in the L2 PPF anyway. Get 4 nodes available for IASI L2)
- **IASI L2 on node parallelization** (currently workorders are split in 2 by the start-up script)



IASI L2 operational processing

IASI provides 1.296.000 hyperspectral infrared measurements per day

15 spectra, of 8461 spectral channels each, per second

IASI L2 Product Processing Facility (PPF) turns spectra into atmospheric profiles of temperature, water vapour and ozone, by iteratively applying a fast radiative forward model (FRTM) until the measurement is matched

Processing time is dominated by the FRTM and is roughly linear in the number of iterations, the number of channels used and the number of fields of view

The timeliness requirement for IASI L2 products on EUMETCast is 3 hours (maximum difference between sensing and reception time)

Processing and dissemination is based on PDUs (chunks of data corresponding to about 3 minutes of sensing time)

IASI L2 PDUs can be processed in parallel on N (currently 3) different processing nodes.

Throughput constraint: average processing time per PDU $< N * 3$ minutes

Timeliness constraint: maximum processing time per PDU < 25 minutes

→ Need for parallelization within each individual PDU



Optimization of RTTOV (the FRTM used for IASI L2)

-Imass library (for fast exponential function)

Optimizations performed at EUMETSAT on RTIASI (later merged with RTTOV) and fed back to ECMWF (speed up factor 5 – 10)

- **loop reordering for improved data locality**
- **taking constant expressions out of the loops**

Further RTTOV developments to benefit from

- **new RTTOV coefficients with pre-computed contribution from fixed trace gases**
- **new RTTOV coefficients on 51 levels (instead of 101)**

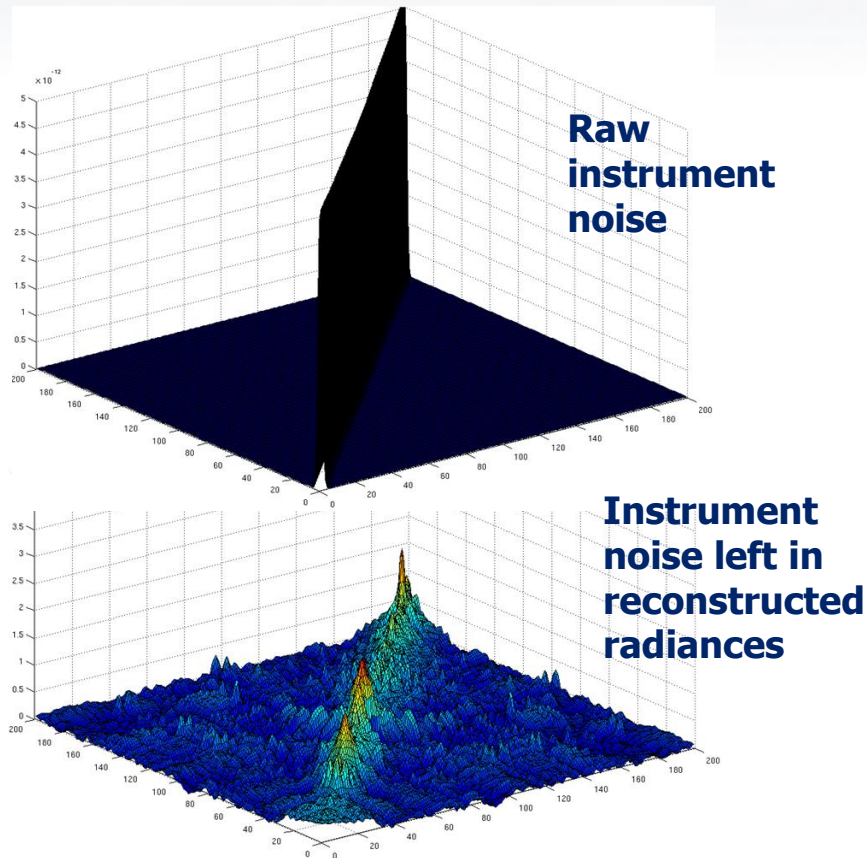
→(wanted) faster implementation of variable trace gases (CO₂, CH₄, N₂O and CO) avoiding multiplications with zeros (“sparse” implementation)

Reconstructed radiances


Virtually all information in the IASI spectra can be captured by a limited number of PC scores – or equivalently a limited number of reconstructed channels →

Use a small subset of channels (200-300 of 8461) while keeping information from the full spectra



No need for PC based forward model (does not harm if not too slow)



Reconstructed BT (K) at 1772.75 cm⁻¹, 20100516_A



BT (K) at 1772.75 cm⁻¹, 20100516_A






Improvement of first guess / background

Training data based on co-located ECMWF forecasts - 393735 cases from 5 days (in different months).

The retrieval statistics are improved considerably by dividing the data into 24 regression classes for which individual regression coefficients are retrieved. The class is determined solely by the AMSU radiance in channel 2 and 4 and the first two PC scores in IASI band 2.


Regression coefficients to predict the sum of the absolute values of retrieval errors of T_a , W_a and T_s are computed, which allows for an estimation of the quality of each individual retrieval.




Regression Class 20120915



MWIR T_a 20120915



Q 20120915



Forli-CO (Processing capacities in Brussels)



232 CPU's dedicated to optimal estimation retrievals of CO, HNO₃ and O₃

26 Nodes running OpenSuse 12.x


21xx: 10 Pentium D @ 3.00GHz (1x Dual core)

82xx: 7 Xeon E5420 @ 2.50GHz (2x Quad core)


83xx: 7 Xeon E5520 @ 2.27GHz (2x Quad core HT)

84xx: 2 Xeon E5620 @ 2.40GHz (2x Quad core HT)

Load



Jobs



Backprocessing PDU's waiting in the queue

~180 PDU's running



Data flux in Brussels



- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU
- 1 CPU = 1 PDU

PBS queuing




BUFR
- L1c
- L2




Forli CPU usage after optimisation at ULB

CO




~3 min

HNO₃



~10 min

O₃



~15 min



Forli-CO rewritten to use thread safe API

File based interface → API (called per each scan-line – about 22 per PDU)

Multithreaded implementation (using pthreads library)

Preliminary execution times (4 threads) on IBM hardware at EUMETSAT:

- **Average time per CPU: 87.5 seconds**
- **Maximum time per CPU: 288 seconds**

More than 5 times slower than reported on Intel Core i7

- **fftw**
- **gsl (→ essl ?)**

Realistic to squeeze in on current hardware, but not with O3 and HNO3 (which are to be added later) → more CPUs need to be allocated for IASI L2



Multithreaded 1DVAR implementation based on RTTOV

Desirable as current on node parallelization scheme is inefficient and does not scale well

“Update to RTTOV v10.2 - 10th April 2012 - to ensure code is thread-safe in all parallel implementations”

This work not started yet. Potential complications:

- memory consumption (separate copies of data structures)**
- does not fit well with current wrapper to C++ (a Fortran Module on top of RTTOV)**

Conclusion:

IASI processing requirements are increasing (more trace gases, more (cloudy) fields of view,....)

→Faster algorithms, software and hardware are always needed

... and then we'll get MTG-IRS providing about 2560 hyperspectral IR spectra per second