ECMWF

20th workshop on high performance computing in meteorology

Enhancing Global Offshore Oil Pollution Monitoring with ECMWF Data

M. Young*, H. Robertson, B. Warrilow, H. McCormack (*Michael.young@cgg.com)

CGG

SeaScope: A real-time platform for global offshore pollution monitoring

Monitoring the temporal and spatial variability of offshore oil pollution is crucial for environmental protection and regulatory compliance. Satellite imagery offers a valuable source of near real-time information on oil pollution, providing global coverage.

SeaScope is a flexible platform for global real-time monitoring of offshore assets, providing Energy companies with a critical environmental awareness and risk mitigation tool. Observations are either made through direct human interpretation from inhouse offshore pollution specialists for the highest accuracy, or through our cutting-edge vision transformation machine learning model which can provide global scale asset monitoring.

The observations provided by SeaScope strengthen situational awareness of Energy companies by maximising their visibility of how active and decommissioned assets (alongside third-party activity close by) interact and impact the local maritime environment.

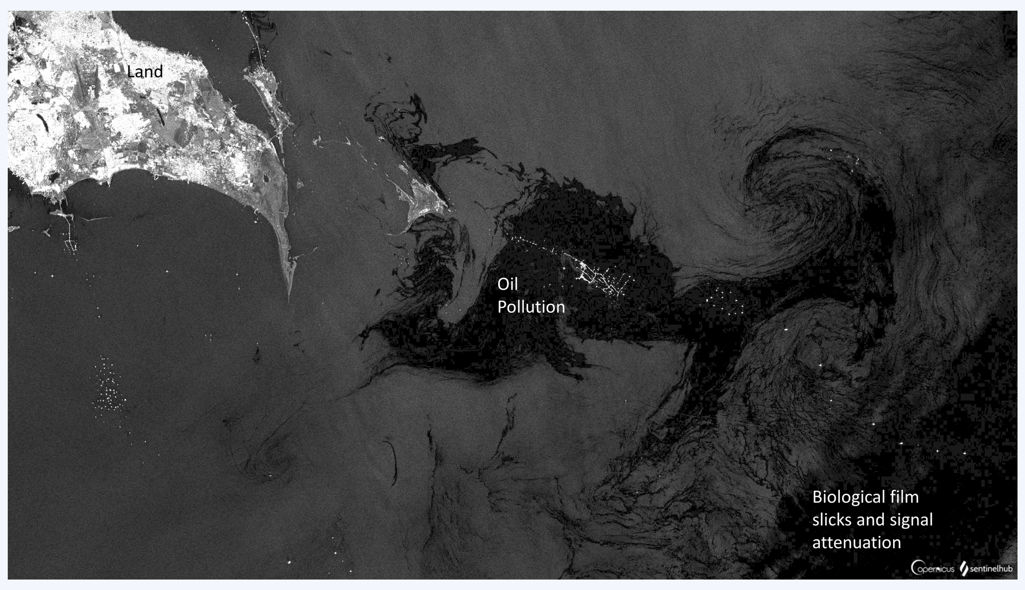

Figure 1: Annotated SAR image containing oil slicks (Contains modified Copernicus Sentinel data 2023)

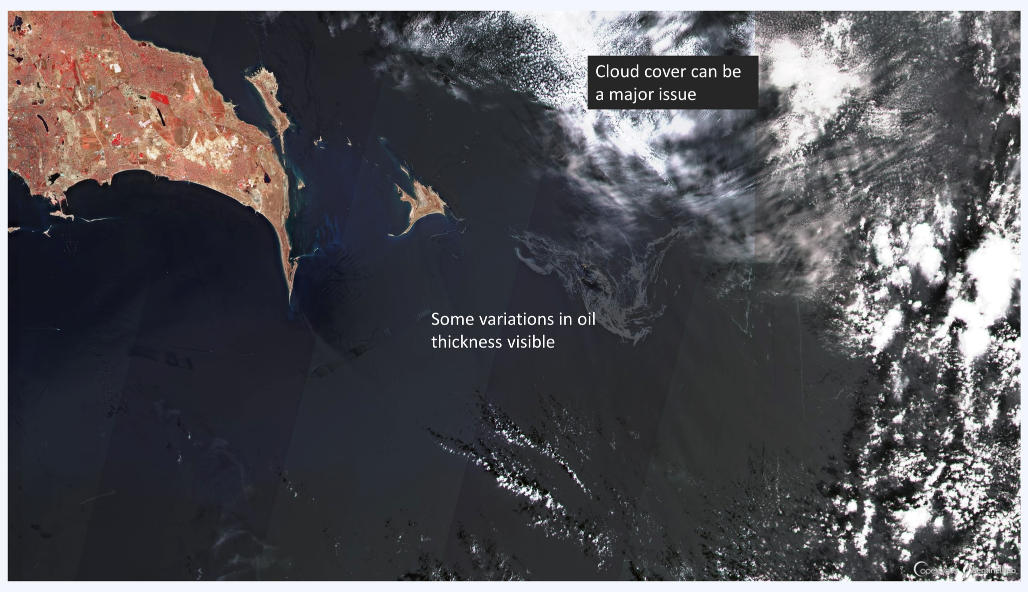

Figure 2: Annotated multispectral image over the same area as Figure 1 (Contains modified Copernicus Sentinel data 2023)

SAR Imagery

SAR is an active remote sensing instrument which utilises radio waves to capture an image of the Earth's surface. Calm, flat seas and slicks cause scattering directly away from the sensor, which appear dark on images. Rougher sea, influenced by wind, returns higher amounts of backscatter, appearing brighter on images. A significant advantage of SAR is the ability for radio waves to penetrate through all but the densest cloud, both night and day.

Multispectral Imagery

Multispectral satellites, such as Sentinel-2, are typically used to supplement oil slick mapping, particularly in coastal regions. Optical imagery relies on sunlight illuminating the Earth's surface. Oil-containing pixels can be detected and characterised because oil has a higher refraction index and absorption coefficient than water. These properties, combined with the various observing conditions in the ocean, form the fundamental basis to detect oil, characterise oil type, and estimate oil concentration or thickness properties from multi-spectral imagery.

ECMWF data and slick observation

The incorporation of weather datasets such as wind speed and ocean currents significantly improve the accuracy of each pollution detection by bringing a level of contextual understanding of the conditions at time of image acquisition. Based on an analysis of over 25,000 Sentinel-1 acquisitions and their corresponding weather layers, we have gained a new understanding of how to observe oil slicks in SAR data. Across all slick categories, the upper limit for mapping slicks at high wind speeds has been pushed higher than previously documented, up to 17.28 m/s (previously published at 13 m/s for pollution slicks). Furthermore, fresher pollution events are sustained at higher wind speeds than fragmenting pollution (UOPS), as fragmenting slicks are often formed of thicker oil.

At wind speeds between 2.5 and 7.5 m/s, the line between normal production water slicks (PWS) and anomalous events is less definitive. Our expert interpretation of PWS shows very little variation throughout the wind speed spectrum, whereas anomalous event classification shows a significant increase at lower wind speeds.

Therefore, in this optimum wind envelope for slick growth and sustenance, more research is needed to reliably distinguish pollution events from legal discharges.

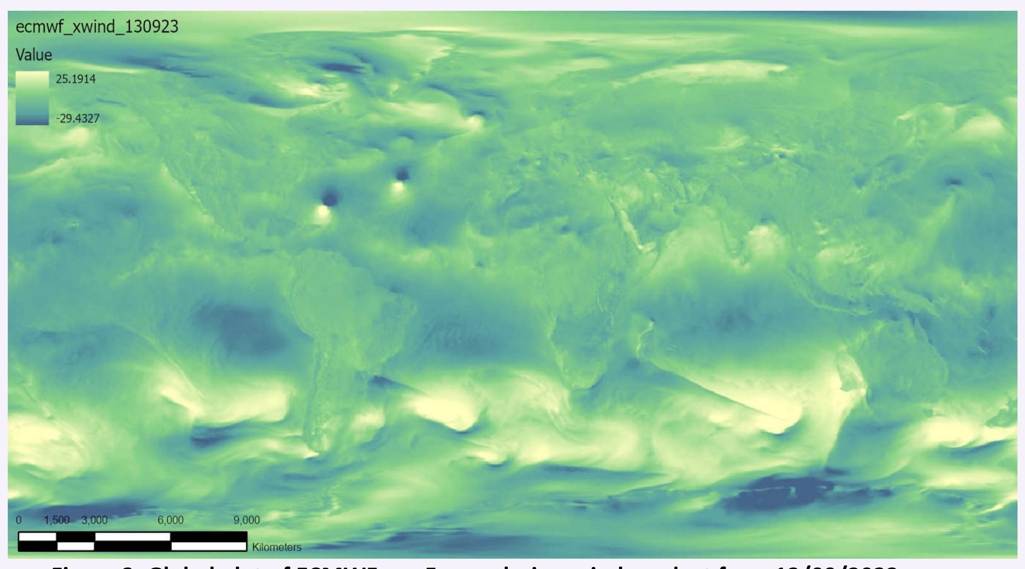

Figure 3: Global plot of ECMWF era-5 reanalysis x-wind product from 13/09/2023

Slick Modelling

These datasets are also integrated into the slick modelling framework. We use a Lagrangian framework to model the slick movements whereby each slick is treated as a cloud of particles which are modelled separately then aggregated.

Forecasting the movements of each slick is crucial for informing clean-up efforts. Given that maritime operations are another major source of oil pollution, AIS vessel tracking data is also integrated into the hindcast modelling. Accurate pollution source estimation is crucial for bringing energy and maritime companies to account for harmful operations.

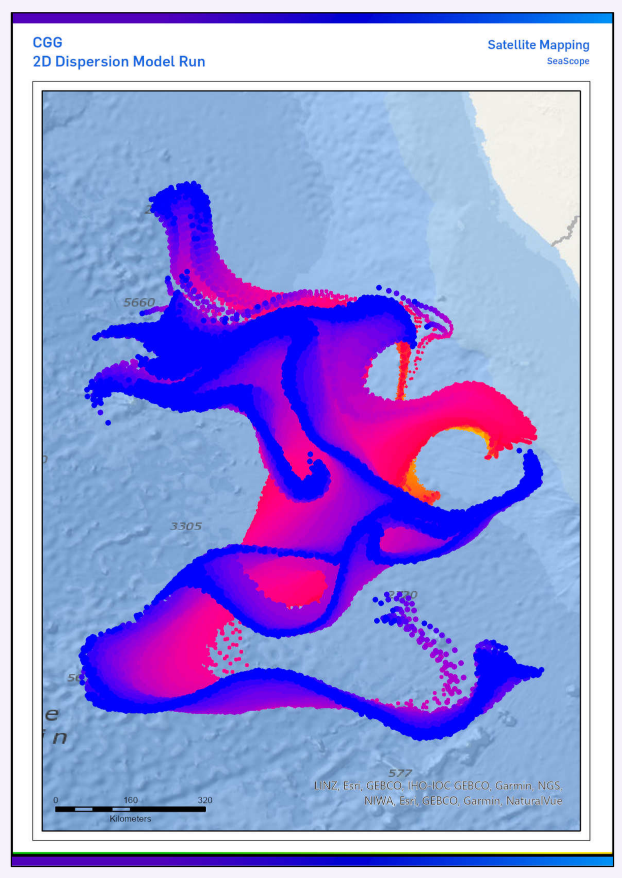

Figure 4: Long-term run testing of the 2D model

High Performance Computing

To achieve the global scale processing of satellite imagery in near real time, we have adopted many cutting-edge technologies in the field of HPC.

Each image that comes through the system requires an amount of automated processing to be performed e.g., atmospheric correction, coregistration, bright spot detection and machine learning interpretation.

Slurm is a workflow manager that is employed to handle the hardware requirements and task dependencies of each processing job, allowing them to be distributed across the compute cluster and performed in parallel to reduce redundancy.

Moreover, an event-driven stream processing architecture, **Kafka**, is designed which can provide a unified, high-throughput, low-latency platform for handling real-time data feeds. It is employed in **SeaScope** to trigger the imagery pipeline for an image as soon as it becomes available from its provider to deliver the interpreted mages in near real time.

Oil Immersion Cooling

Data centres already account for roughly 2% of global electricity consumption – with more than a third of this power used for cooling electrical components. Data centres need more efficient, more climate-friendly ways to drive their operations.

We have always been at the forefront of industrial HPC architecture adoption. At the turn of the millennium, we were pioneering the use of commodity clusters, and started to add accelerators a couple of years later, even before GPGPU programming languages formally emerged.

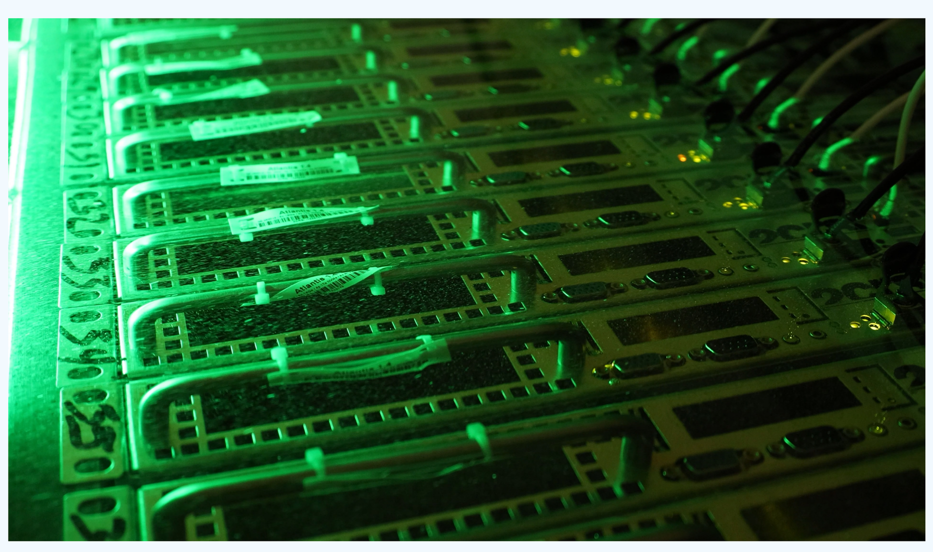

Figure 5: High- density GPU servers submerged in oil

To cater for higher-power densities and improve PUE, we adopted immersion cooling in 2009. Oil immersion has allowed us to push these boundaries further by enabling us to make significant reductions in electricity usage by immersing our hottest servers in oil (and eliminating cooling fans), and also over time to go denser and denser. We are currently using very dense GPUs servers dissipating more than 2 kW per U.

As chips and designs continue to push boundaries with ever higher thermal power densities (TPDs), managing the dissipated heat becomes increasingly challenging.

Through innovative immersion cooling, we can address demanding HPC and ML workloads, maximising performance and ensuring reliable operations while benefiting from unparalleled cooling efficiency, space optimization, scalability, environmental sustainability, and future-proofing capabilities.

CGG