

# HTC on HPC



Parsl's Role in Mapping Earth's  
Topography from Space

Claire Porter  
Polar Geospatial Center  
Sept 2022


# PGC DEM INITIATIVES

REMA, ARCTICDEM, EARTHDEM

1. Stereoscopic high-resolution imagery coverage
2. Scalable terrain extraction algorithm
3. Ridiculous compute resources



# HTC ON HPC

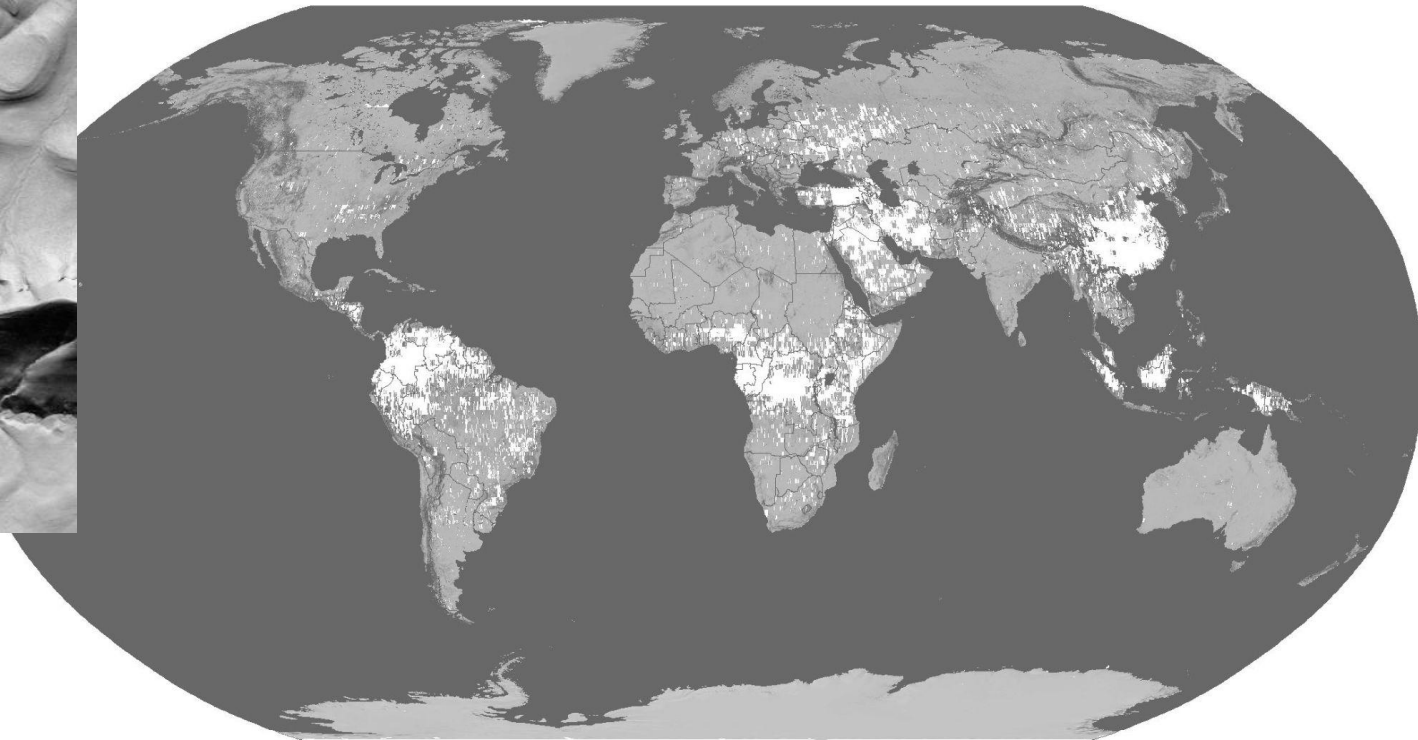


Source: Polar Geospatial Center

## PROBLEM SCALE

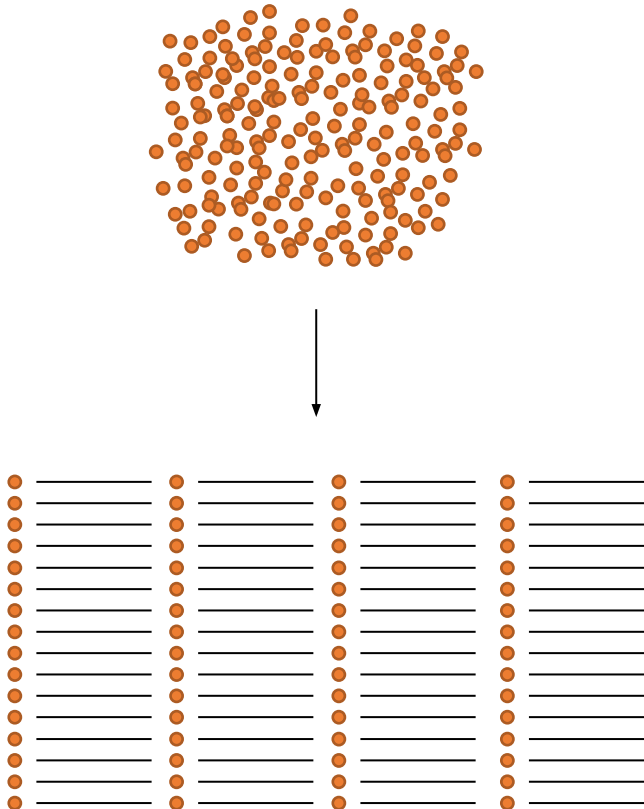
	ArcticDEM	REMA	EarthDEM	Total
<i>Strip DEMs</i>	365,766	232,898	950,070	1,548,734

— **15,177,593** DEM extraction tasks  
— **1,548,734** DEM assembly tasks




# HTC ON HPC

Direct Job Submission:  
1 to 1 task-job bundling

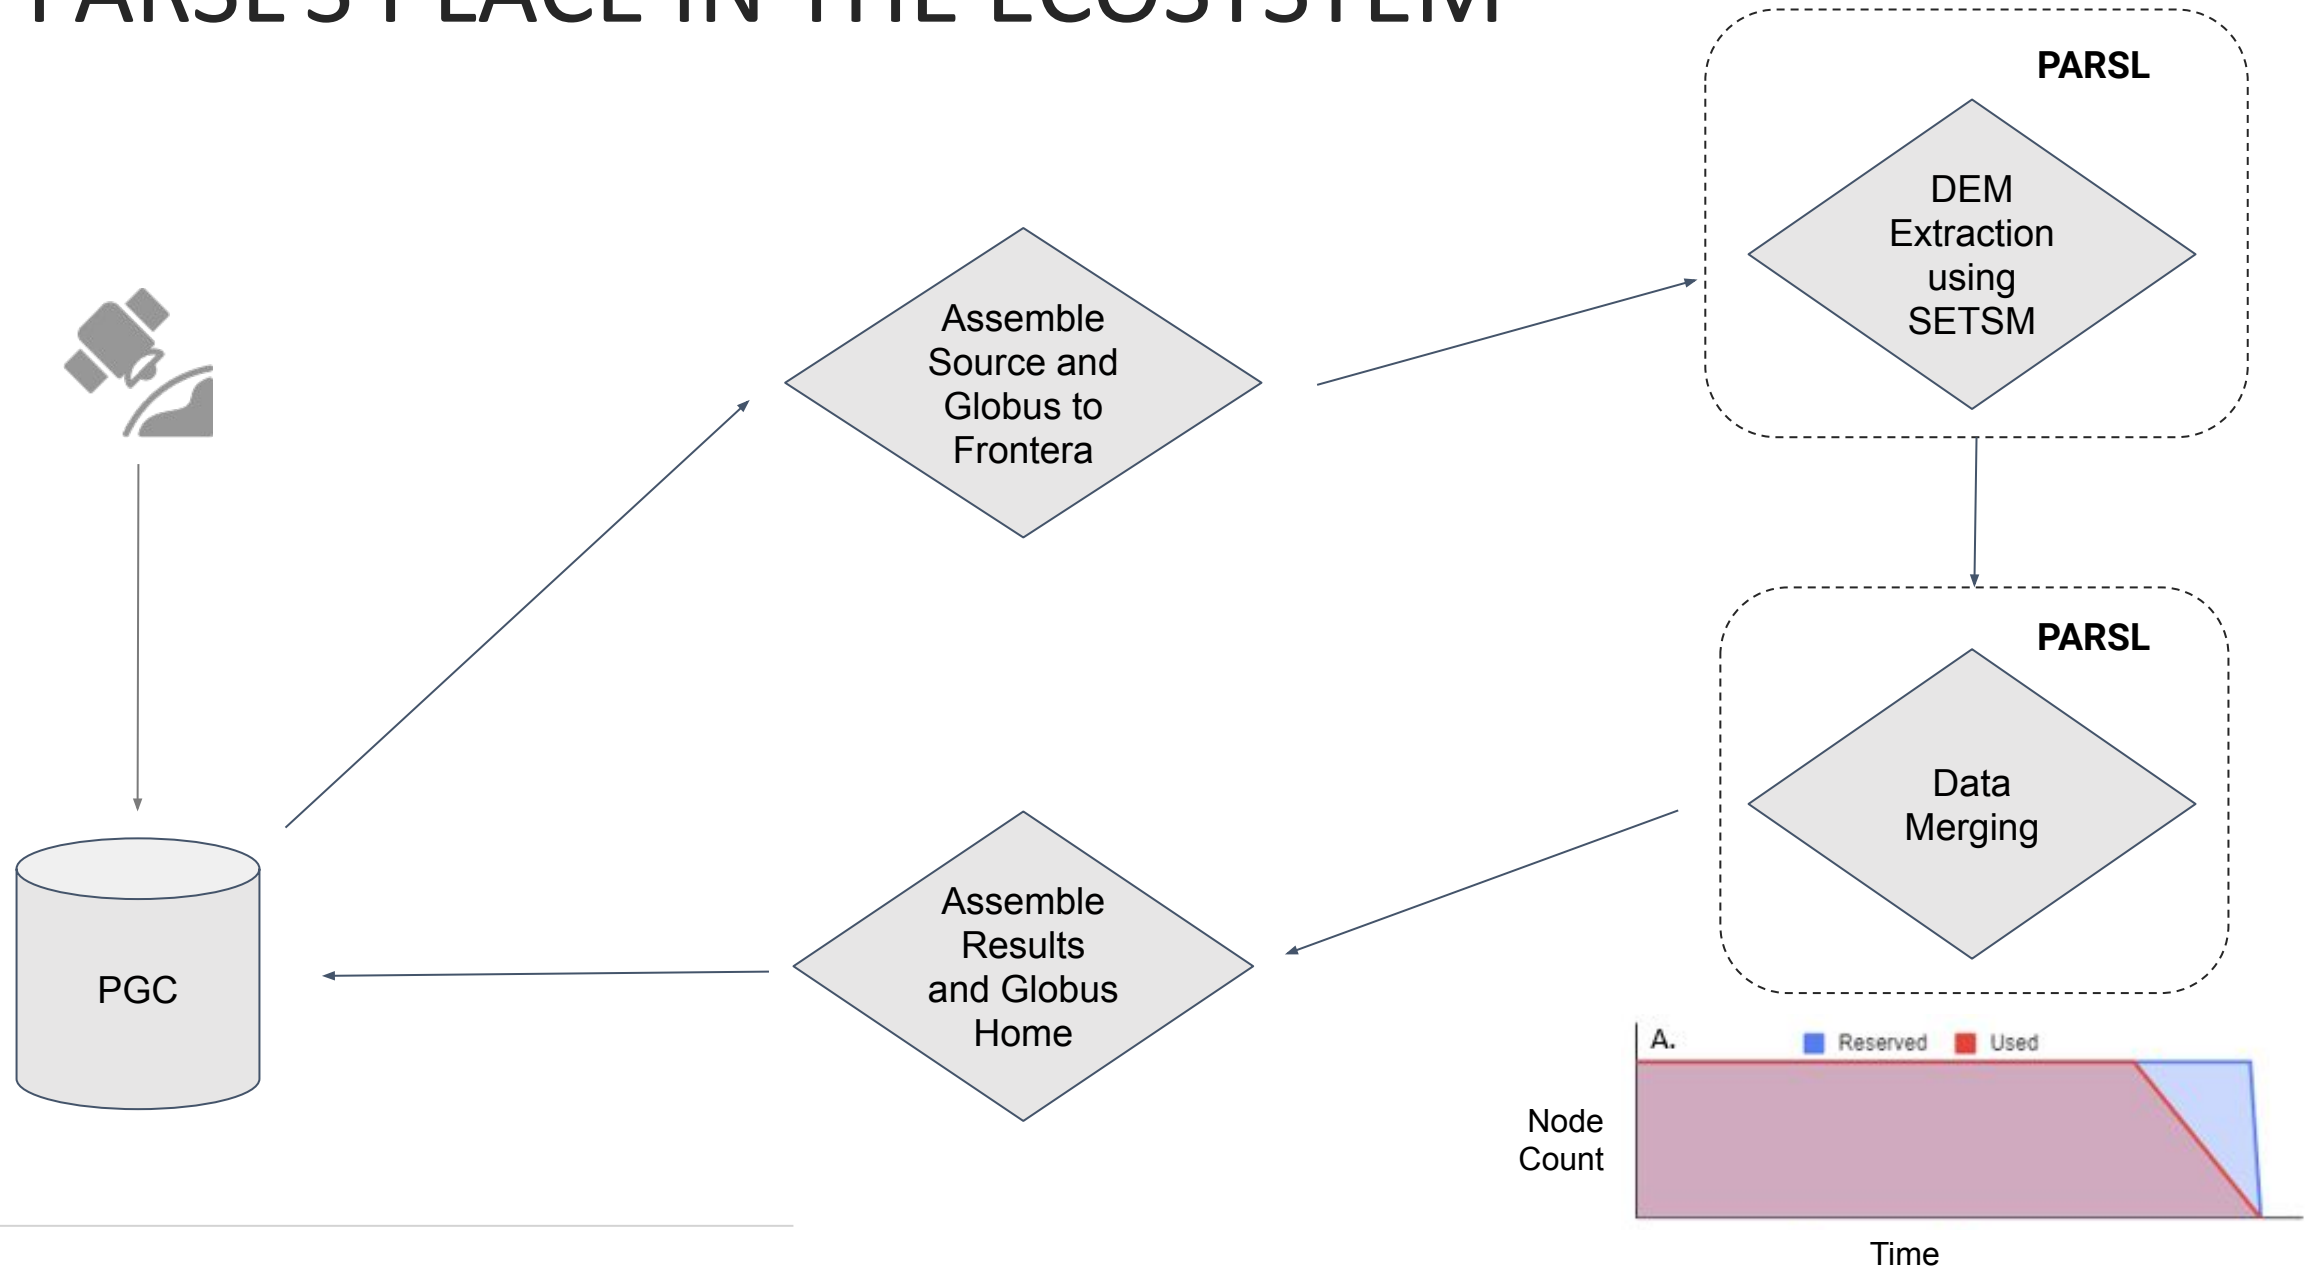


Swift/Parsl:  
100+ to 1 task-job bundling

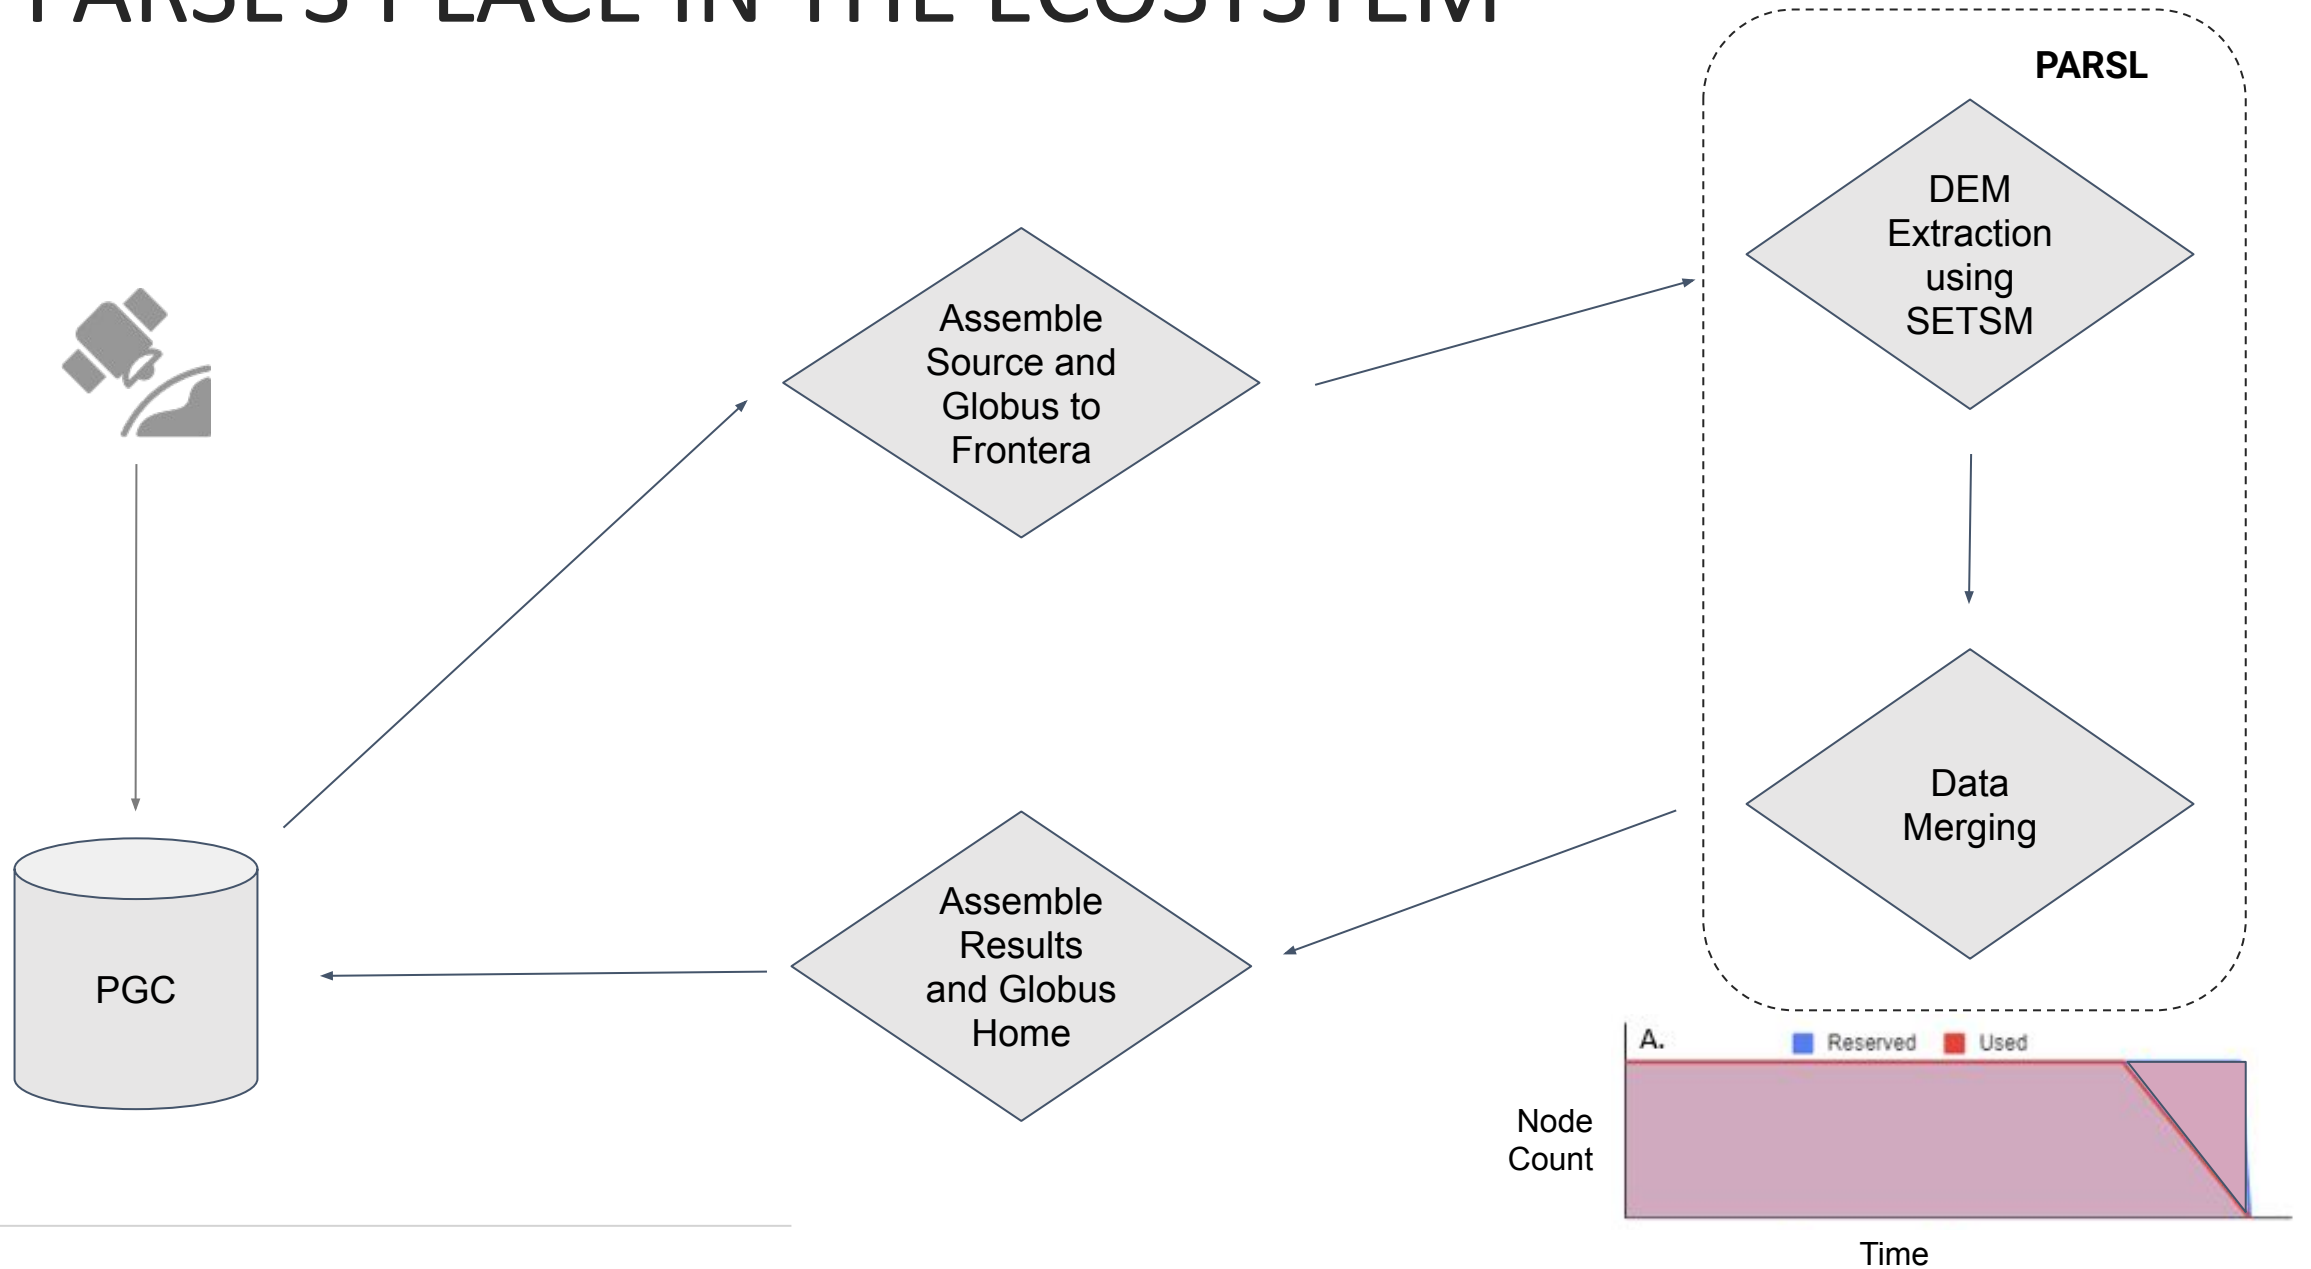


- A set of tasks can number over 10,000, each using  $\frac{1}{2}$  a single node
- If project queue limits are high, jobs cripple the scheduler by raising iteration time
- If project queue limits are low, jobs can't effectively use the available compute nodes
- Parsl allows tasks to be bundled into groups, reducing the impact of the scheduler while achieving high compute node usage
- Parsl also can dynamically allocate tasks to open nodes within a running job – they are not limited to a static list of tasks to run (inefficient if run time varies greatly)

# PARSL'S PLACE IN THE ECOSYSTEM



# PARSL'S PLACE IN THE ECOSYSTEM



# PARSL'S PLACE IN THE ECOSYSTEM

