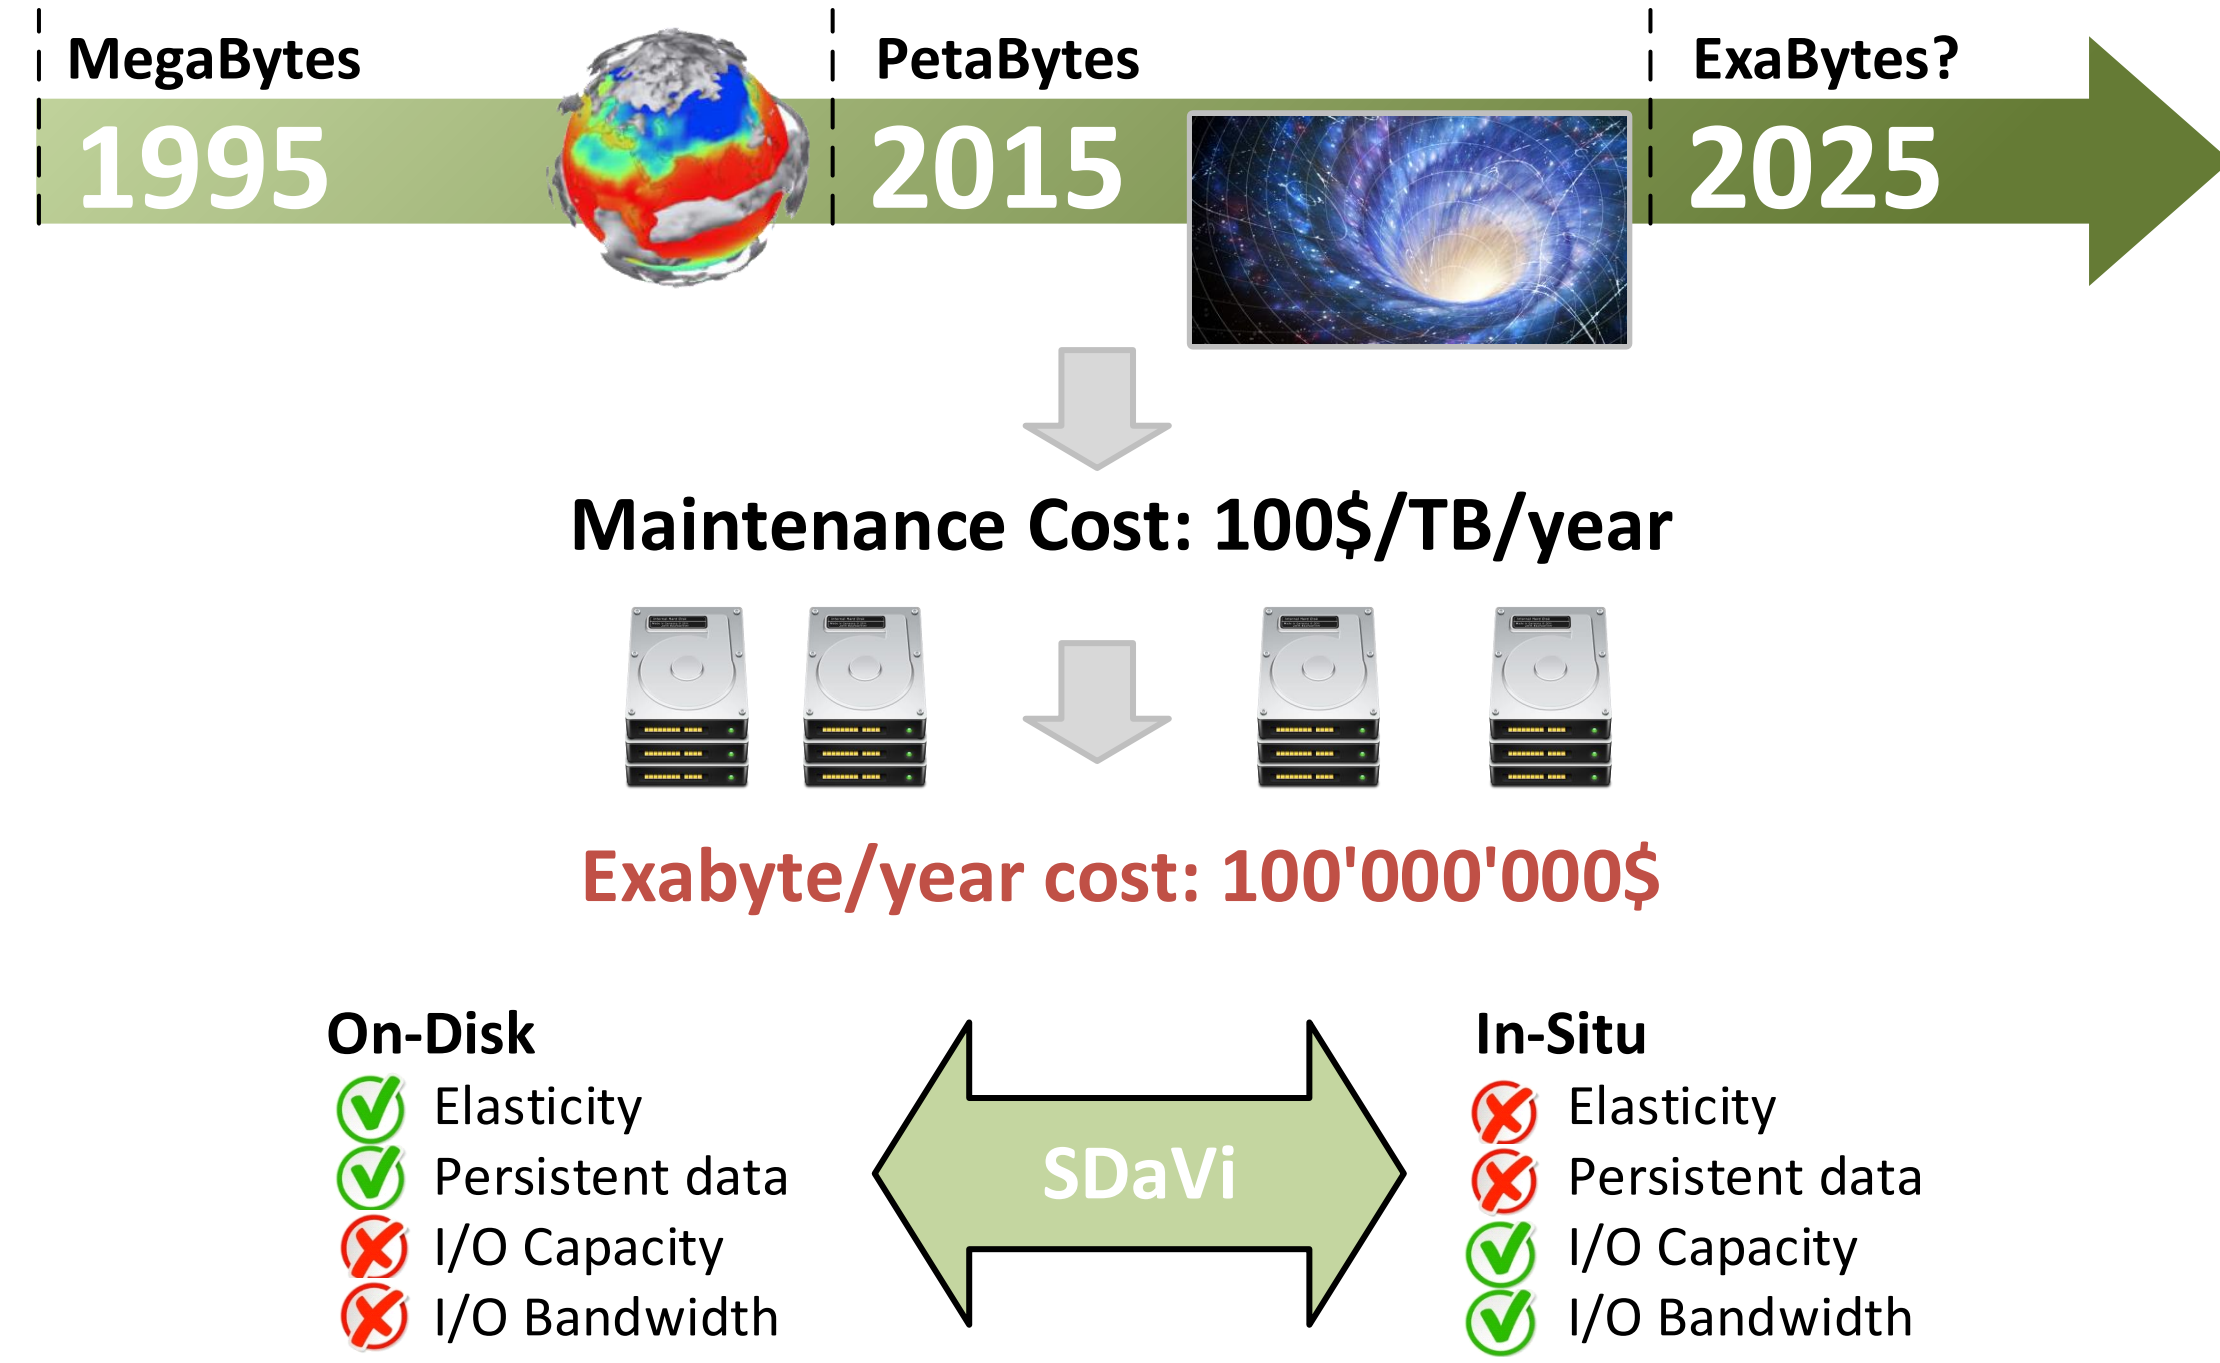


Teaser & Motivation

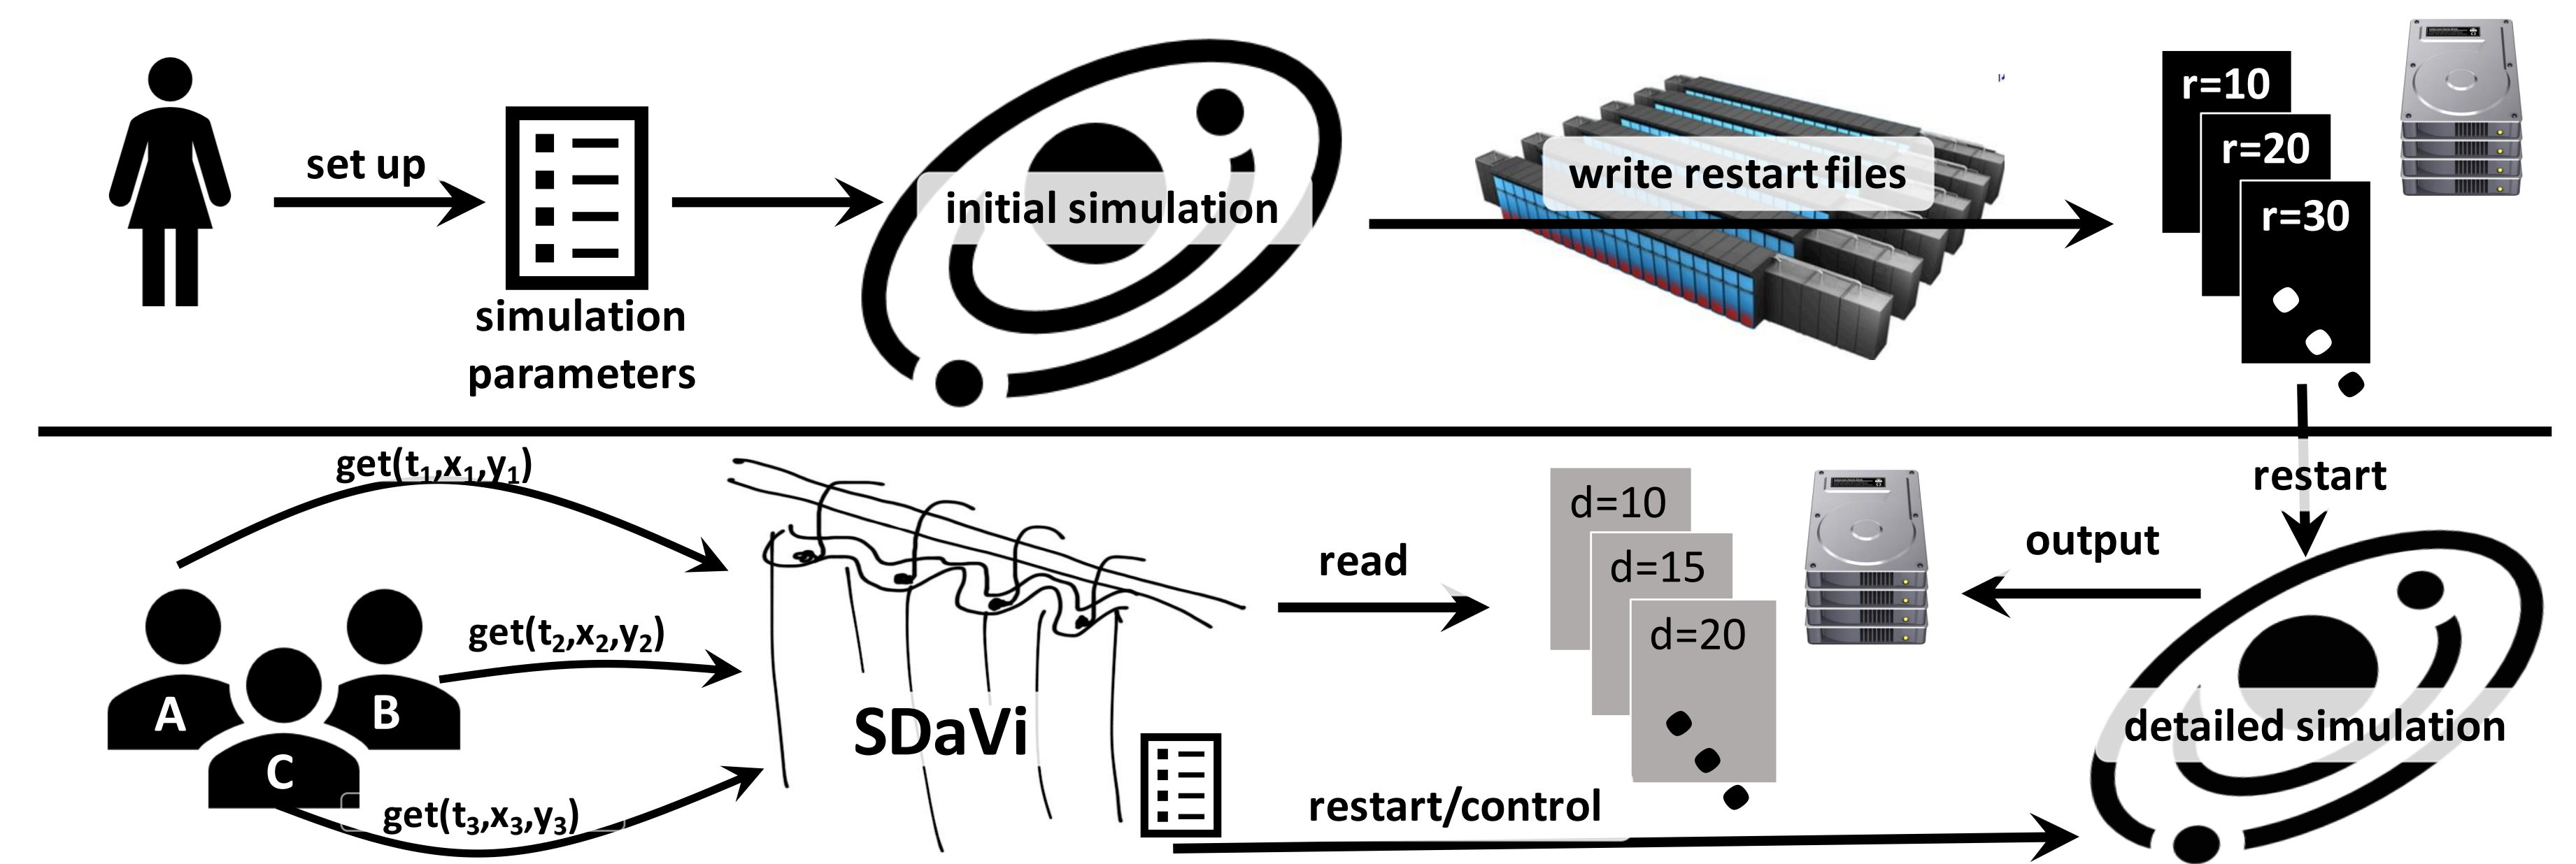


Abstract
 Scientific simulations are being pushed to the extreme in terms of size and complexity of the addressed problems, producing astonishing amount of data. If the data is stored on disk, analysis applications can randomly access simulation output. Yet, storing the massive amounts of ever-growing simulation output data is a large part of the big data challenge in scientific computing. This is primarily due to the high storage costs and the fact that compute capabilities grow faster than storage capacities and bandwidths. In situ analysis removes the storage costs by processing the data as it is generated but applications lose random access.

We propose to not store the full simulation output data but to produce it on demand. Our system – SDAVi, Simulation Data Virtualizer - intercepts I/O requests of both analysis tools and simulators, enabling data virtualization. This new paradigm allows us to explore the computation-storage tradeoff, by trading computation power for storage space. Hence, SDAVi can shift towards the on-disk or in-situ approach according to the available/assigned storage space.

We apply SDAVi to applications employing typical data traversal patterns in two different contexts: climate modeling (COSMO) and multiphysics simulations (FLASH). SDAVi offers a viable path towards exa-scale scientific simulations, by exploiting the growing computing power and relaxing the storage capacity requirements.

SDaVi: Simulation Data Virtualizer



Interfacing the Analysis Tools

Virtualization-Oblivious:
 Analysis tool is not aware of the virtualized environment. Requests are expressed as calls to common I/O libraries and served one-by-one.

Operation	(P)netCDF	(P)HDF5	ADIOS
open	nc(mpi)_open	H5Fopen	adios_open
put	nc(mpi)_vara_put_type	H5Dwrite	adios_write
get	nc(mpi)_vara_get_type	H5Dread	adios_schedule_read
close	nc(mpi)_close	H5Fclose	adios_close

Virtualization-Aware:
 Additional hints on the access pattern are provided.

```
int sdavi_request_range(char * begin, char * end, int stride)
```

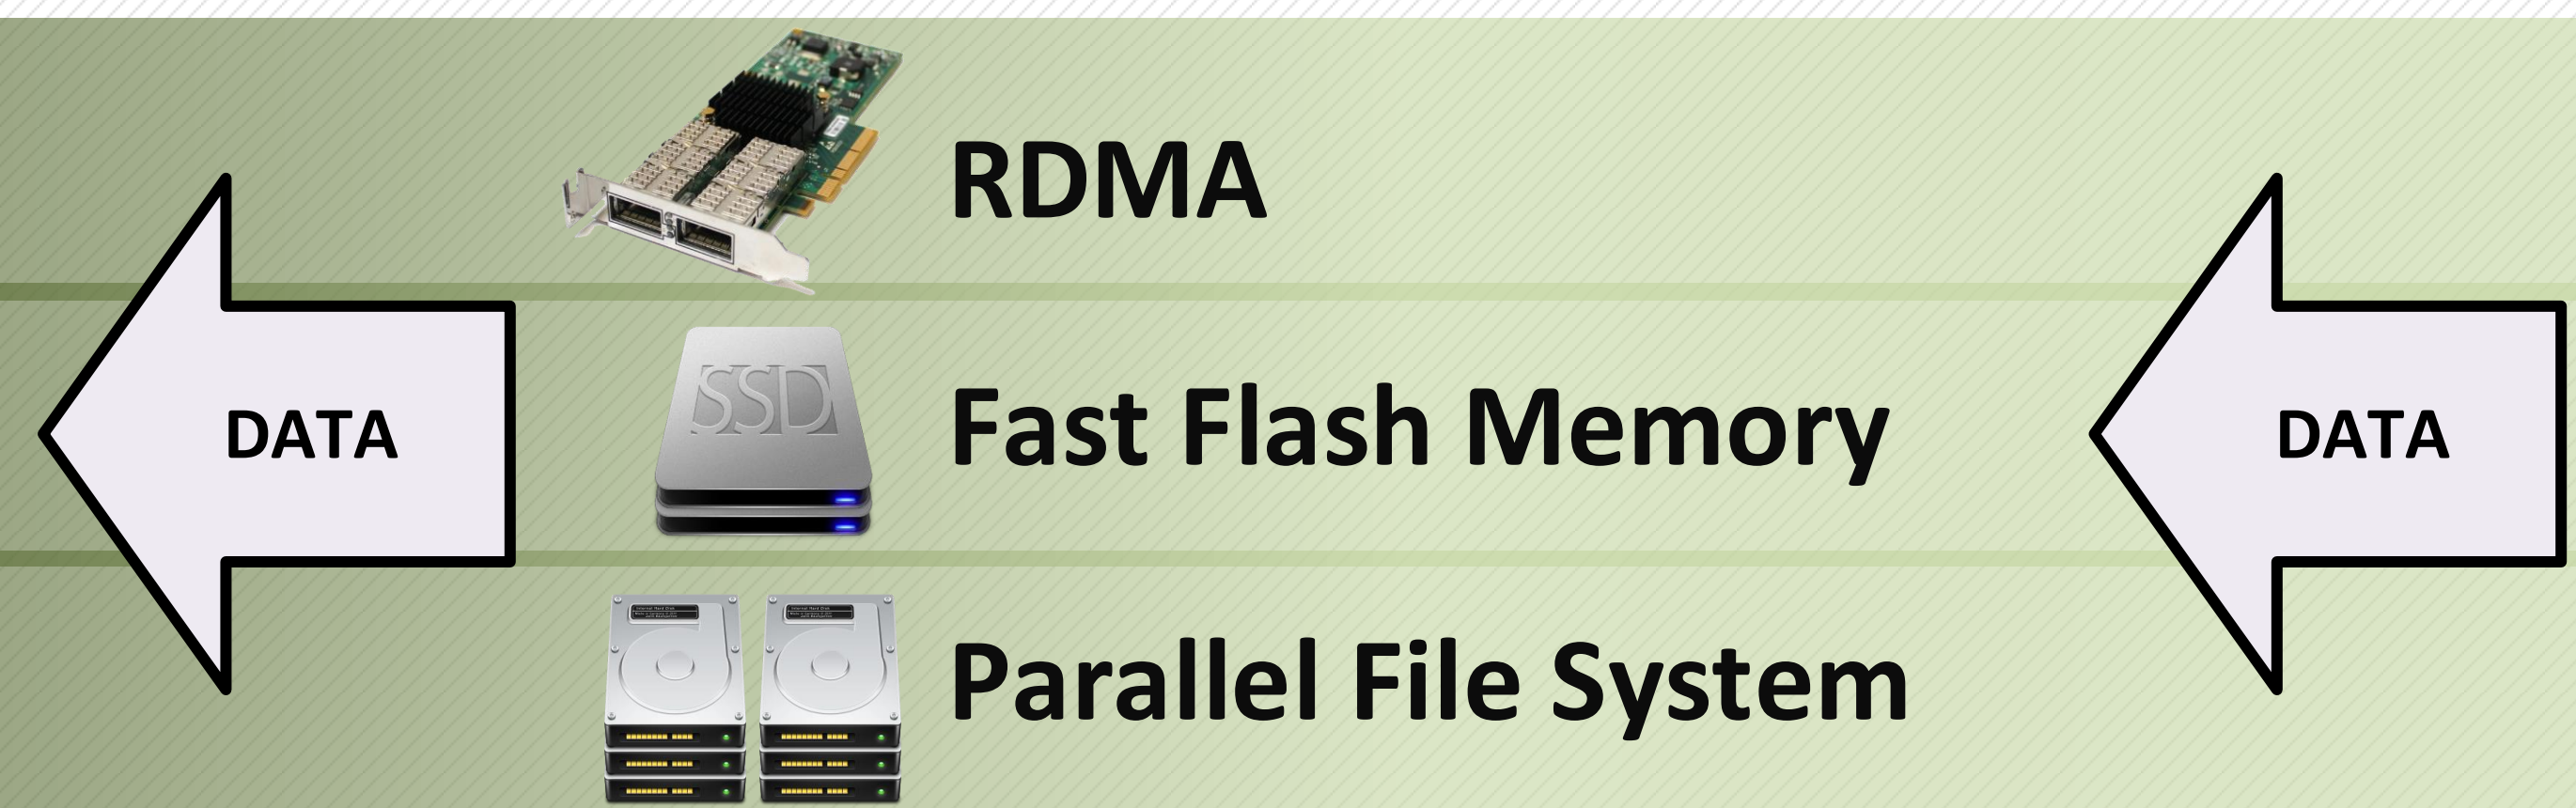
Synchronized:
 Sets of analysis tools operate in a block-synchronous manner.

Data Virtualizer

Restarts simulations for missing output steps.
 How to minimize computing resources?
Horizontal Prefetching: longer simulations for constant-trajectory analysis tools

Keeps track of analysis tools and their access patterns.
 How to optimize the response time?
Vertical Prefetching: Increase the number of parallel re-simulation to match analysis tool's BW

Manages the cache hierarchy.
 What is the best caching strategy in this context?



On-Demand Simulations

Modelling Simulations
 Simulations can be restarted at specific restart steps. Results are written in output steps.

Performance Model
 Time to simulate n output steps starting at time i:

$$T(i, n) = \alpha_{sim} + (i \bmod \Delta r) \cdot \tau_{sim} + n \cdot \tau_{sim}$$

SDaVi profiles simulation instances and updates the models' parameters at runtime.

How to interface a simulator?
 A set of simulator-specific functions (i.e., convert filenames in integer keys) can be provided as a LUA configuration file.

Experiments

