# XtremLab: Characterization of Volunteer Computing Resources

P. Malécot, D. Kondo, G.Fedak, F.Cappello (INRIA/LRI)

**Abstract :**
**Volunteer computing (VC) systems use the idle computing power of many volunteered desktop PC's to support large-scale computation and storage. For over a decade, VC systems have been the largest and the most powerful distributed computing systems, offering a plethora of computing power at a fraction of the cost of supercomputers. The volunteer desktops participating in VC projects are volatile and heterogeneous, but there is little detailed information about their volatility and heterogeneity. Yet this characterization is essential for the simulation and modelling of such systems. XtremLab is a project whose goal is to obtain detailed picture of the VC landscape. To this end, we will deploy resource sensors on the volunteer desktops to measure the availability of CPU, disk space, memory, and network bandwidth. Our immediate goal is to provide resource availability analyze. The resulting resource measurement data and characterization will be useful for a broad range of research areas, including distributed and peer-to-peer computing, and fault tolerance.**

## Volunteer Computing Challenge and Opportunities:

• VC systems (SETI@Home, distributed.net) are large and powerfull distributed systems, offering an abundance of computing power at a fraction of the cost of dedicated, custom-built supercomputers.
•Many applications from a wide range of scientific domains, including computationWe conduct our measurements via the BOINC desktop grid software infrastructure by submitting actual tasks to the system. These tasks perform computation and periodically
•report their computation rates. This method allows us to measure excatly what actual compute power a real compute-boucn application would be able to exploital biology, climate prediction, particle physics, and astronomy
• VC resources have two major characterisctics : they are shared and volatile. As a result, any number of factors can cause the task of a VC application to not complete : keyboard activity, the execution of other user applications, machine reboots, or hardware failures.
• Will help to design efficient VC system for a broader range of application : DAG, real-time, data intense, communicating applications

## Expected Results:

• Public Traces Repository. The measurements of resources activity will be useful to enable accurate simulation and modelling of VC systems. For example, the traces could be used either to directly drive simulation experiments or to create generative probability models of resource availability, which in turn can be used by simulators to explore a wide range of hypothetical scenarios.
• Desktop Grid Resources Characterization
• Prediction Model of Resources Activity

## Experimental Setup

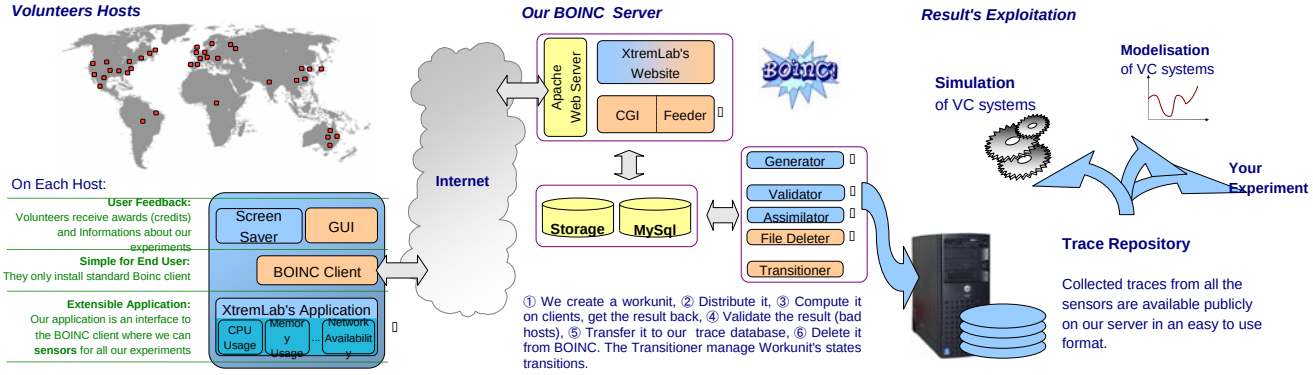


① We create a workunit, ② Distribute it, ③ Compute it on clients, get the result back, ④ Validate the result (bad hosts), ⑤ Transfer it to our trace database, ⑥ Delete it from BOINC. The Transitioner manage Workunit's states transitions.

## Resource Availability Characterization

• We conduct our measurements via the open source BOINC desktop grid software infrastructure by submitting actual tasks to the system.
• These tasks are executed only with the host is idle, as determined by the user's preferences and controlled the BOINC client.
• These tasks perform computation and periodically record their computation rates to file.
• These files are collected and assembled to create a continuous time series of CPU availability for each participating host.
• Our active trace method allows us to measure exactly what actual compute power a real, compute-bound application would be able to exploit.
• Compared to other passive measurement techniques, our method is not as susceptible to OS idiosyncracies and takes into account keyboard and mouse activity, which directly impact application execution.
• Using the resulting trace data, one can construct application performance models or use the traces to drive simulations, for example. See the graph of application work rate as an example.

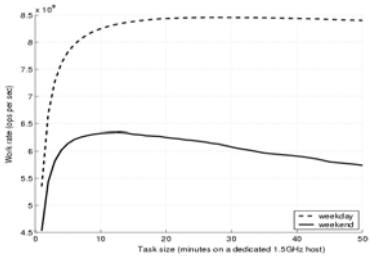


Model of application work rate for an enterprise desktop grid

# http://XtremLab.lri.fr

XtremLab Team Contact
Paul Malécot malecot@lri.fr
Gilles Fedak fedak@lri.fr
Derrick Kondo dkondo@lri.fr
Franck Cappello fci@lri.fr

Address:
INRIA Grand-Large
LRI, Building 490
University Paris XI
91405 Orsay Cedex
France