


# Ousting Ivory Tower Research: Towards a Web Framework for Providing Experiments as a Service

Tim Gollub, Benno Stein, Steven Burrows

## Motivation

- ▶ Comparability of experimental results is fundamental for empirical research.
- ▶ A lack of comparability has been attested for results in scientific publications.
- ▶ A review of SIGIR 2011 full papers shows underutilization of public data and software (right).
- ▶ Comparisons are significantly simplified if data and software assets are published together with papers.
- ▶ In our view, the most convenient technique to enhance comparability is “experiments as a service”.
- ▶ We are developing a platform for “experiments as a service” for information retrieval research.



## Design Goals

### 1. Local Instantiation

- ▶ Enables public research on private data.
- ▶ Enables comparisons with private software.

### 2. Unique Resource Identifiers

- ▶ Enables linkage of experimental results in papers with the respective experiment service.
- ▶ Enables reproduction of results on the basis of the resource identifier (digital preservation).

### 3. Multivalued Configuration

- ▶ Enables the specification of whole experiment series.

### 4. System Independence

- ▶ Enables a widespread usage of the platform.
- ▶ Enables the deployment of any experiment software without internal modifications.

### 5. Distributed Execution

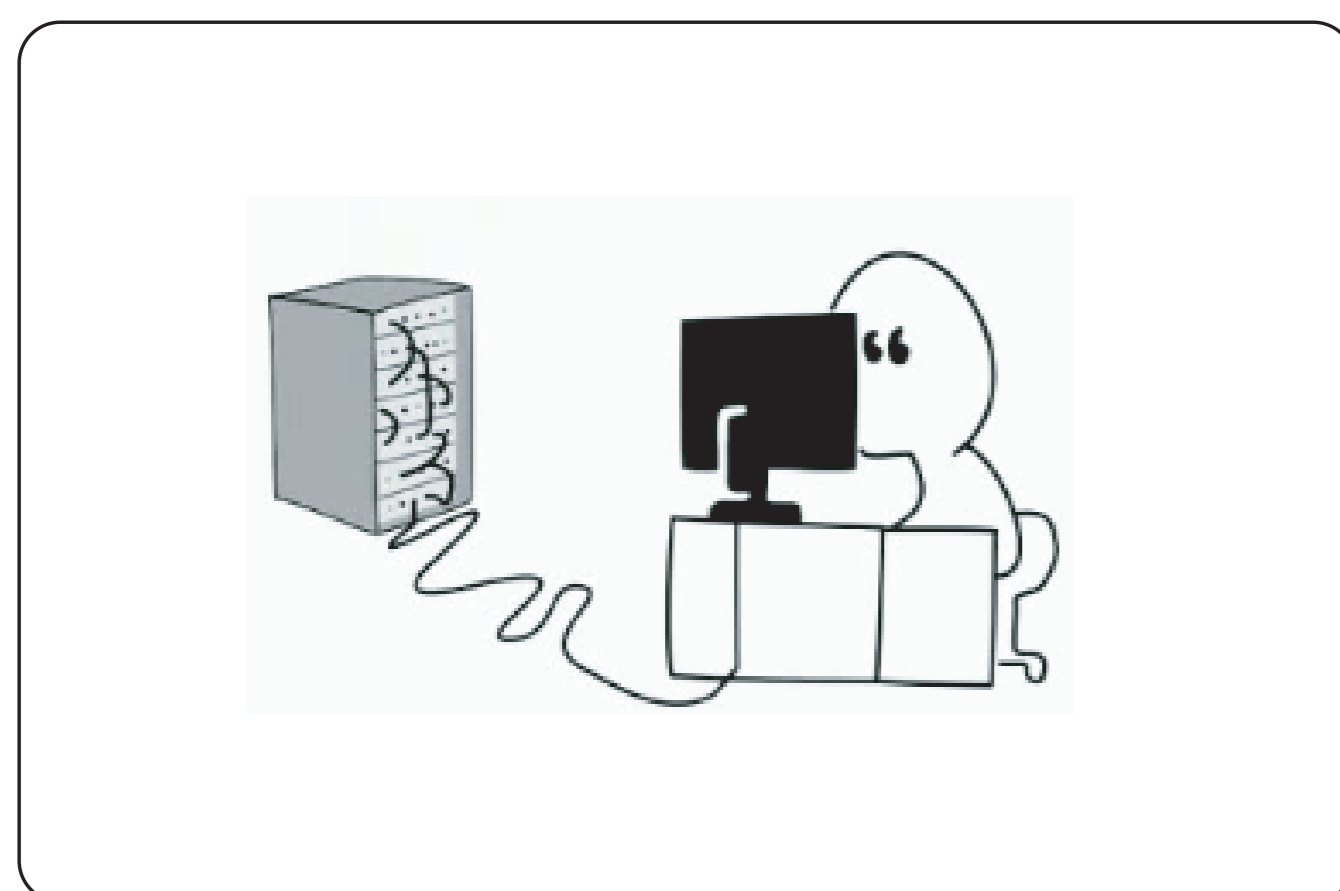
- ▶ Enables efficient computation of pending experiments.

### 6. Result Storage

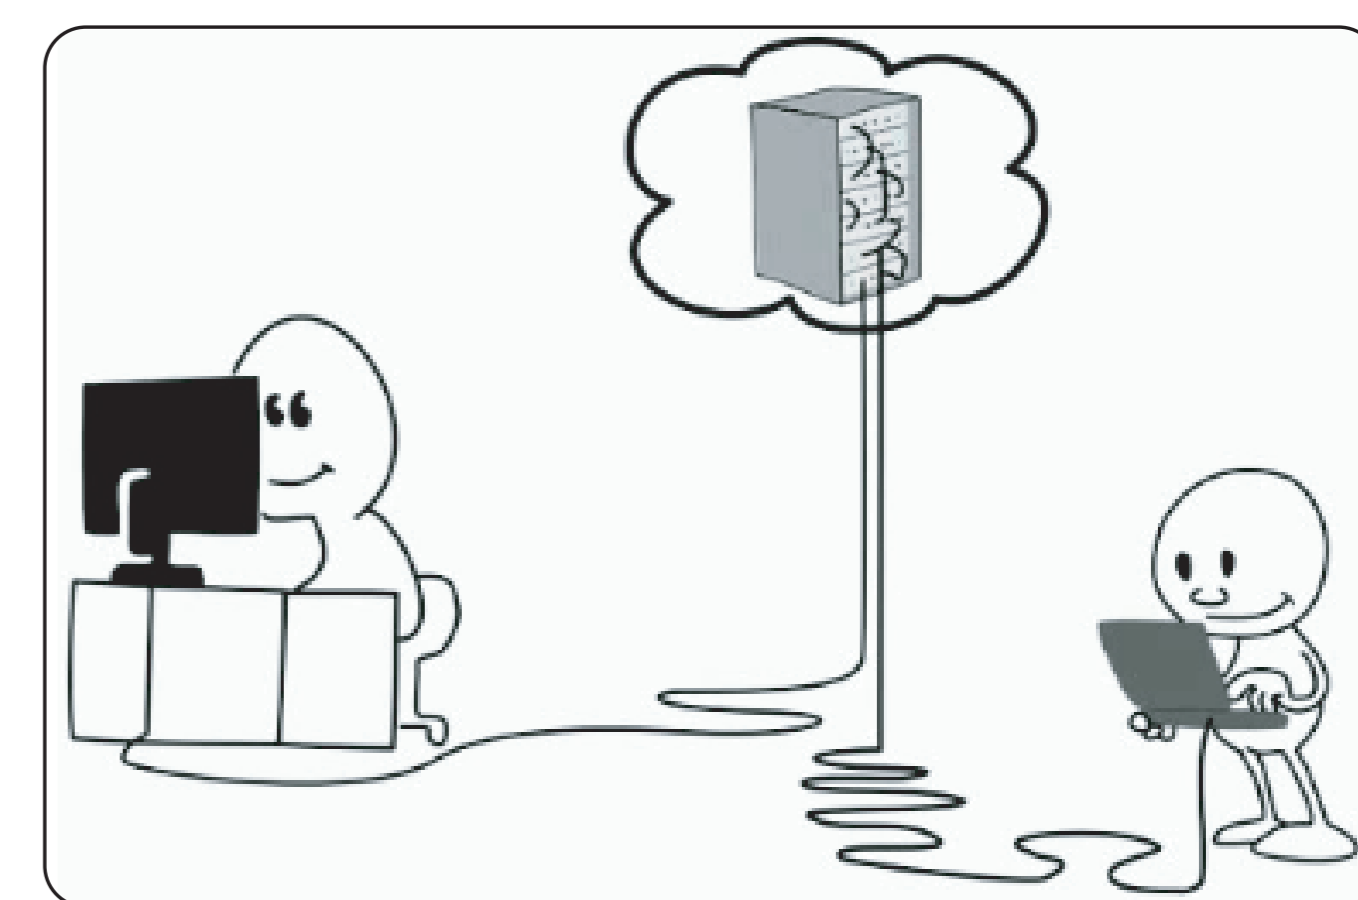
- ▶ Enables retrieval and maintenance of raw experiment results.

- localhost:2306/programs/examples/MyProgram?p1=42&p2=Method1&p2=Method2
- 
- tira@node1:~\$ ./myprogram.sh -p1 42 -p2 "method1"
- tira@node2:~\$ ./myprogram.sh -p1 42 -p2 "method2"
- | Parameter 1 | Parameter 2 | Output Directory                 | Performance |
|-------------|-------------|----------------------------------|-------------|
| 42          | Method1     | <a href="#">output-directory</a> | 0.89        |
| 42          | Method2     | <a href="#">output-directory</a> | 0.71        |

## Peer to Peer Collaboration



1. Instantiate a local instance of the experimentation platform.
2. Deploy your experiments to the local instance.
3. Grant public access to your local platform.
4. Allow other researchers to deploy algorithms to your platform.
5. Apply for an evaluation track at TREC, CLEF, etc.
6. Become the gatekeeper of your stream of research.



## Upcoming Publications

1. Tim Gollub, Benno Stein, Steven Burrows, and Dennis Hoppe. TIRA: Configuring, Executing, and Disseminating Information Retrieval Experiments. In Proceedings of the Ninth International Workshop on Text-based Information Retrieval at DEXA, Vienna, Austria, September 2012.
2. Tim Gollub, Steven Burrows, and Benno Stein. First Experiences with TIRA for Reproducible Evaluation in Information Retrieval. In Proceedings of the First SIGIR Workshop on Open Source Information Retrieval, Portland, Oregon, August 2012.