

StarExec: a Cross-Community Infrastructure for Logic Solving*

Aaron Stump¹, Geoff Sutcliffe², and Cesare Tinelli¹

¹ Department of Computer Science, The University of Iowa

² Department of Computer Science, University of Miami

Abstract. We introduce StarExec, a public web-based service built to facilitate the experimental evaluation of *logic solvers*, broadly understood as automated tools based on formal reasoning. Examples of such tools include theorem provers, SAT and SMT solvers, constraint solvers, model checkers, and software verifiers. The service, running on a compute cluster with 380 processors and several terabytes of disk space, is designed to provide a single piece of storage and computing infrastructure to logic solving communities and their members. It aims at reducing duplication of effort and resources as well as enabling individual researchers or groups with no access to comparable computational infrastructures. StarExec allows community organizers to store, manage and make available benchmark libraries; competition organizers to run logic solver competitions; and individual community members to run comparative evaluations of logic solvers on public or private benchmark problems. The StarExec web site provides facilities to upload and organize benchmarks, browse and query the stored benchmark libraries, run selected solvers on selected benchmarks, and view and analyze execution results and statistics.

1 Introduction

Ongoing breakthroughs in a number of fields depend on continuing advances in the development of high-performance automated reasoning tools such as SAT solvers, SMT solvers, theorem provers, model finders, constraint solvers, rewrite systems, model checkers, and so on, which we will generically refer to as (*logic*) *solvers*. Typically, application problems are translated into possibly large and complex formal descriptions (logical formulas, rewrite rules, transition systems, ...) for these tools to reason about. Different tradeoffs between linguistic expressiveness and the difficulty of the original problems have led to the adoption of different reasoning approaches and the use of different logical formalisms to encode those problems. Different research communities have developed their own research infrastructure to spur innovation and ease the adoption of their solver technology. This infrastructure includes standard input/output formats for logic solvers, e.g., [5, 11, 2], libraries of benchmark problems, e.g., [11, 9, 6, 4], solver execution services, e.g., [10, 1, 12], and solver competitions [7, 3, 8]. By and

* Work made possible in large part by the support of the National Science Foundation through grants 0957438, 1058748, and 1058925.

large, so far these different infrastructures have been developed separately in the various logic communities, at significant and largely duplicated cost in development effort, equipment and support.

StarExec is a solver execution and benchmark library service aimed at facilitating the experimental evaluation of automated reasoning tools. It is designed to provide a single piece of storage and computing infrastructure to all logic solving communities, with the twofold goal of reducing the duplication of effort and resources while also enabling communities and individuals that could not afford developing their own infrastructure. The service allows community organizers to store, manage and make available benchmark libraries, competition organizers to run competitions, and individual community members to run comparative evaluations of logic solvers on benchmark problems. These capabilities are accessible through a web browser interface at the StarExec web site which provides facilities to upload and organize benchmarks and solvers, browse and query the stored benchmark libraries, run user-selected solvers on user-selected benchmarks, and view and analyze execution results and statistics.

This paper gives an overview of StarExec, briefly describing its main design, components, functionality, and usage, and discusses its current status and further development plans. For more details, we refer the reader to the online documentation on the StarExec web site: <http://www.starexec.org>.

2 Main Concepts

StarExec is built as a service for logic solving *communities*, groups of logic solver developers and users such as the SAT, SMT, theorem proving, confluence, model checking, and software verification communities. In StarExec, a community is a user group administered by one or more designated *community leaders*, users with special privileges and responsibilities such as accepting new users or uploading and managing benchmarks and solvers. The StarExec system is built on top of a few basic concepts, *spaces*, *users*, *solvers*, *benchmarks*, *job pairs*, *jobs*, and *processors*, described below.

Spaces. A space is a collection of solvers, benchmarks, jobs, users (collectively referred to as *space items*) and other spaces (referred to also as *subspaces*). Spaces are similar to folders in a file system and are the means by which items within StarExec are organized. Spaces are structured hierarchically into a *tree* of subspaces. Hence each space other than the root space belongs to exactly one superspace. In contrast, solvers, benchmarks, jobs, and users can belong to more than one space. The direct subspaces of the root space are *community spaces*. Each logic community has one community space which contains, directly or within subspaces, all benchmarks, solvers, users and jobs in that community. Spaces can be *public* or *private*, *locked* or *unlocked*. A public space is visible to all StarExec users; a private one is visible only to the users contained in it. If a space is locked no one can copy items *out* of it. This is useful to store copy-restricted benchmarks and solvers while still allowing other users to include them in job pairs.

Users. A user is a registered and verified person who belongs to at least one community and is therefore endorsed within the system by a community leader. Every user has visibility over a set of (sub)spaces and a set of *permissions* for each of these spaces controlling the kind of permissible operations. Users also have *leadership* over some

spaces. A leader of a space has the full set of permissions over that space. A community leader is simply the leader of a community space.

Solvers. Solvers are programs that solve logic problems. Specifically, they are Linux programs that take a text file on standard input and may produce text on standard output. Solvers can consist of group of files and folders—as opposed to a single executable. Each solver is associated with one or more *configurations*, executable scripts that execute the solver with specific input flags and settings. Configurations and processor scripts (see below) can be written in any of the scripting languages (Bash, Perl, Python, and so on) available by default in Linux distributions.

Benchmarks. A benchmark is a single text file that is fed to a solver as input. Currently there is no support for multi-file benchmarks; however, StarExec supports benchmark dependencies, the referencing of other benchmark files (e.g., TPTP axioms files [11]) within a benchmark. Each benchmarks have an associated *type*. A type consists of a (system-wide) unique name and an associated benchmark processor.

Jobs and Job Pairs. A job pair consists of a solver configuration and a benchmark to be run by that configuration's solver. Job pairs are atomic execution units. After execution they are associated to various pieces of information such as results, CPU time, runtime, etc. A job is simply a collection of job pairs to be executed.

Processors. Processors are scripts that take textual input and produce textual output at certain stages in the job execution pipeline. They are currently of two kinds: benchmark processors and post-processors. The former are run on benchmarks being uploaded to a space. They can be used to perform various checks and validations on a benchmark and to extract meta-data to be stored in a central database. Post-processors are run on the output produced by the execution of a job pair. They can be used to extract specific data from this output which too are stored in a database for later access.

3 Functionality and Usage

StarExec was designed as a web service accessible through a web browser interface. Currently, most of the web service functionality is also available through a downloadable client application, called StarExecCommand. That program provides a command line shell and a large number of commands mirroring the operations that can be performed via a web browser.³ In the following, we focus on the web browser interface.

User Accounts Using StarExec requires first obtaining a user account. Since every registered user belongs to at least one of the StarExec communities, every new account request requires joining one of them (more communities can be joined later). Such requests are sent to all the leaders of the chosen community each of whom has the power to approve them or not. Communities themselves can be created only by the StarExec administrators. Registered users can be promoted to leaders of a community by existing leaders of that community or by the StarExec administrators. Registered users must log in to StarExec before being able to use it. External observers can view all public information on the StarExec website by logging in as the special user *guest*.

³ The interested reader is referred to the StarExecCommand manual, included in the tool's distribution, for the list of available commands and their functionality.

Spaces and Communities Using an interface similar to that of file explorers in computer desktop GUIs, users can browse the StarExec space hierarchy—or rather the subset of it visible to them—view the contents of any spaces they belong to, and inspect details of users, benchmarks, solvers, and jobs in those spaces. As an example, details about benchmarks include name, description, owner, upload date, file size, type, and any user-defined attributes automatically extracted from the benchmark when it was uploaded. Users can also perform a number of stateful operations on a space depending on the set of permissions they have on it. These operations may include creating and deleting subspaces, uploading, copying, linking and deleting benchmarks and solvers, creating and removing jobs, copying or removing users, making other users leaders, changing the default permissions, and locking or unlocking the space.

In general, users can copy any visible item from an unlocked space to any space for which they have permission to add such kind of item. While copy operations are done uniformly by a drag-and-drop action, their semantics depends on the copied item. Copying a user U from a space A to a space B amounts to giving certain permission on B to U . Copying a job J from A to B only creates a link in B to J . Copying a benchmark or a solver on the other hand can be done either way, as an actual copy or just as a link. All users are by default *leaders* of the spaces they create. The leaders of a space have full permissions on that space, which consist in the ability to add/remove items and subspaces as well as make other users to space leaders. A space leader can also decide which set of permissions from the above to grant to regular (i.e., non-leader) users in the space. For security and integrity reasons, no users can remove space leaders (including themselves) from a space or demote them to regular users.⁴

Since communities are spaces, users can operate on them as discussed above. However, they can also browse the communities visible to them, request to join or leave one of those communities, and download a community's benchmark processors. Community leaders (i.e., leaders or a community space) can, in addition, set several community-wide defaults such as default timeouts or post-processors for job pairs. They can also upload post-processors, define benchmark types and upload their associated processor.

Compute Cluster StarExec compute nodes are partitioned into a number of execution queues. Each job is submitted to one of these queues. There is a general queue available to all users. The other queues serve special functions, such as solver competitions, or are reserved for exclusive use by a community. Only community leaders can request the creation and the reservation of a queue. A dedicated page on the web server allows all users to inspect any queue to see pending job pairs for that queue, and inspect any node in a queue to see currently executing job pairs on that node.

Solvers StarExec users can add a solver to a space (in which they have the permission to add solvers) either by copying/linking the solver from another space or by uploading it. Solvers must be uploaded as tar archives or compressed tar or zip archives. The archives can consist of any collection of folders and files but must contain at least one configuration script—used to establish proper settings and input flags for the solver and then launch the solver. All configuration scripts must reside in a designated top-

⁴ This can be done only by StarExec administrators.

level directory (`bin`) in the archive and have a name starting with a designated prefix (`starexec_run_`) so that they can be easily identified by the StarExec system.

During job executions, StarExec will run a configuration script in an environment with a number of predefined environment variables. Those variables contain such information as the absolute path to the benchmark input file, the absolute path to a designated output directory for the script, and limits on wall clock time, CPU time, memory, and disk space. Any files written to the output directory during the execution of the script are saved by StarExec for later download by the user. Files written elsewhere are removed after each job run. If a job pair exceeds any of its time or space limits it is terminated.

As with any space item, users can inspect details of any solver visible to them. In particular, they can view the list and the content of each configuration script. If permitted by the containing space owner, they can also download the entire solver and its configurations as a compressed archive.

Benchmarks Users can add benchmarks to a space either by copying/linking them from another space or by uploading an archive file containing benchmarks all of the same type. The user has the option of (i) recreating the archive's directory structure as a space structure in the destination space and place each benchmark in the (sub)space corresponding to the benchmark's source (sub)directory, or (ii) simply placing all the files contained in the archive directly in the destination space. With the first option the new subspaces take the name of the corresponding subdirectory and all get the same set of permissions, specified by the user before uploading the archive. All the uploaded files are run through the benchmark processor associated to their type, and are added to their destination space only if accepted by the processor. An upload status page summarizes the results of the upload as it progresses and lists all the discarded files, if any.

Benchmark types are global across all communities. They can be created only by community leaders who must also provide a benchmark processor. A benchmark processor, which currently can be one a single file script, is expected to print for its input benchmark a sequence of lines of the form *key=value* where *key* is an attribute name and *value* its value. These attribute-value pairs are stored in StarExec's database and will be shown later when inspecting the benchmarks. With a couple of exceptions, all attribute names are user-defined. A predefined attribute, and the only mandatory one, is `starexec-valid`. The processor is required to print a pair with that key and the value `true` or `false` depending on whether the benchmark was accepted by the processor or not. This attribute is used by the system only to check whether to store or discard the benchmark. Other predefined, but optional, attributes allow the processor to specify dependencies on other benchmarks or an expected result for the benchmark.

Jobs Users can create and immediately execute a job in a space. At creation time they set execution parameters specifying timeouts, post-processors, and execution queues. They can choose to have job pairs executed in depth-first or in round-robin fashion. In the first case, StarExec will execute all job pairs in one subspace before moving on to the next; in the second, all job pairs in all subspaces will make progress in the execution concurrently. Users have different ways to generate a job's job pairs starting from the *root space*, the space in which the job is created. For instance, they can instruct the system to find all subspaces containing solvers and benchmarks, and execute all possible combinations of those benchmarks and available configurations for those solvers.

Alternatively, users can manually select which benchmarks and solvers to execute. In that case, they also have different options for how to pair solvers with benchmarks.

After a job has been set up, its job pairs are created automatically and sent to the execution queue as specified. Once created, the job can be monitored by looking at its details page which gets updated in real time with information on how many pairs have been completed, have been solved, have failed and so on. Each job pair in turn can be inspected for information on its status, final runtime, result, etc. After the job completes, its page will also provide various statistics in graphical form, with scatter plots and cactus plots. All of this job data can be downloaded in CSV format, with one line per job pair, for off-line analysis, process and visualization. The full job's output file created by the chosen post-processor can be downloaded as well.

4 Infrastructure and Technologies

The StarExec hardware infrastructure is located in a dedicated, state-of-the-art server hosting facility at the University of Iowa. The software infrastructure relies on several freely available software applications and libraries. Almost all of the software developed in-house for this service is written in Java, with the rest consisting mostly of shell scripts. We plan to open source the entirety of this software in the near future.

Hardware The StarExec service runs on a Red Hat Fedora compute cluster consisting of 3 head nodes and 190 rack-mounted compute nodes. Each node has two 4-core 2.4GHz Intel processors, 256GB of RAM, and a 1TB hard drive. Local disk space is used only for caching purposes during job execution. Solvers, benchmarks and all other persistent data and meta-data are stored centrally in a dual NetApp network-attached storage system with a capacity of 23TB.

Software tools The StarExec web service is built on top of the Apache Tomcat web server and servlet container. Web pages, based on the HTML5 and CSS3 standards, are generated dynamically using the Java Server Pages (JSP) platform and its Standard Tag Library (JSTL). All client-side scripting is done in Javascript and uses jQuery for most operations—such as animate DOM objects or send AJAX requests to the web server. Data is exchanged between web browsers and the server via AJAX requests and is written in JSON format. The StarExec back-end relies on the Oracle Grid Engine to distribute jobs to compute nodes and manage execution queues and workload. Most back-end persistent data (on spaces and their items) and all meta-data generated by processor scripts is stored in a MySQL database. Benchmarks, solvers and various system logs are stored directly in the networked file system.

5 Current Status and Future Development

At the time of this writing, StarExec has been used by two public events: the Confluence Competition (CoCo) 2013, which ran in June, 2013; and the SMT Evaluation (SMT-EVAL) 2013, which ran over several months in 2013. Each event had an execution queue giving it exclusive access to a subset of the nodes. Otherwise, they had quite different requirements for StarExec. CoCo 2013 ran during a meeting of the Confluence

Workshop, and thus it was very important that the CoCo organizers could initiate the job and monitor its results live. They used StarExecCommand for this. On the other hand, their workload was small, just 509 job pairs. In contrast, SMT-EVAL 2013 had a workload of 1,663,478 job pairs, split over four jobs. This meant that operating at scale was an important criterion for success for SMT-EVAL.

For the FLoC Olympic Games of Summer 2014, many additional events have expressed interest or plans to run on StarExec. We anticipate that future developments of the system will be driven by the needs of such events. For example, in some tracks of the Termination Competition, termination proofs produced by termination checkers for rewriting systems become inputs to tools which can certify such proofs. We anticipate adding support for pipelines of tools, to support the Termination Competition. The challenge in adding new features to meet the needs of individual events or communities is to try to cast such features in a more general way, that could be useful for other communities. The ability to create pipelines of tools is one good example of a feature which, while needed for a specific event, is likely to be of interest for others.

Acknowledgements We thank several people who have contributed in various capacities to the StarExec project so far. The following people were involved in the development of StarExec's software infrastructure at various stages of the project: E. Burns, T. Elvers, T. Jensen, W. Kaiser, B. McCune, M. Nassar, CJ Palmer, V. Sardeshmukh, S. Stark, and R. Zhang. Computer system support and assistance in designing and building the hardware infrastructure was provided by H. Brown, D. Holstad, J. Tisdale, and JJ Ulrich. Several people from user communities and from the StarExec Advisor Board provided useful feedback and input (please see the StarExec website for a full list).

References

- [1] SMT-Exec. <http://www.smtexec.org>.
- [2] The QDIMACS Format. <http://www.qbflib.org/qdimacs.html>.
- [3] C. Barrett, M. Deters, L. de Moura, A. Oliveras, and A. Stump. 6 Years of SMT-COMP. *Journal of Automated Reasoning*, 50(3):243–277, 2012.
- [4] C. Barrett, A. Stump, and C. Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). <http://www.SMT-LIB.org>, 2010.
- [5] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB Standard: Version 2.0. In A. Gupta and D. Kroening, editors, *Proceedings of the 8th International Workshop on Satisfiability Modulo Theories*, 2010.
- [6] H. Hoos and T. Stütze. SATLIB: An Online Resource for Research on SAT. In I. Gent, H. van Maaren, and T. Walsh, editors, *Proceedings of the 3rd Workshop on the Satisfiability Problem*, 2001. <http://www.satlib.org/>.
- [7] D. Le Berre and L. Simon, editors. *Special Issue on the SAT 2005 Competitions and Evaluations*, volume 2, 2006.
- [8] R. Nieuwenhuis. Special Issue: The CADE ATP System Competition. *AI Communications*, 15(2-3), 2002.
- [9] T. Rath, J. Otten, and C. Kreitz. The ILTP Problem Library for Intuitionistic Logic - Release v1.1. *Journal of Automated Reasoning*, 38(1-2):261–271, 2007.
- [10] L. Simon and P. Chatalic. SatEx: A Web-based Framework for SAT Experimentation. In H. Kautz and B. Selman, editors, *Proceedings of the LICS 2001 Workshop on Theory and*

Applications of Satisfiability Testing, number 9 in Electronic Notes in Discrete Mathematics, pages 129–149. Elsevier, 2001.

- [11] G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure. The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.
- [12] G. Sutcliffe. The TPTP World - Infrastructure for Automated Reasoning. In E. Clarke and A. Voronkov, editors, *Proceedings of the 16th International Conference on Logic for Programming Artificial Intelligence and Reasoning*, number 6355 in Lecture Notes in Artificial Intelligence, pages 1–12. Springer-Verlag, 2010.