Ⓔ

SCEC Broadband Platform: System Architecture and Software Implementation

by Philip J. Maechling, Fabio Silva, Scott Callaghan, and Thomas H. Jordan

Online Material: Text-based (BBP) input and output files and figures of goodness of fit; seismogram comparisons and source model illustrating the BBP validation of the Loma Prieta earthquake.

INTRODUCTION

The Southern California Earthquake Center (SCEC) Broadband Platform (BBP) is an open-source software distribution that contains physics-based ground-motion models capable of calculating earthquake ground motions at frequencies above 10 Hz across regional distances. In addition, the BBP contains software tools for evaluating ground-motion models and comparing simulation results to observed ground-motion recordings and against ground-motion prediction equations (GMPEs). The BBP also includes software utilities that help users run large numbers of ground-motion simulations and manage the simulation results. Several of the ground-motion simulation methods included in the BBP are described in related articles in this special focus issue. In this article, we describe how the BBP system architecture and software implementation support the scientific and engineering processes needed to assess ground-motion models for use in engineering applications.

BBP SYSTEM AND SOFTWARE REQUIREMENTS

The BBP design and development was initiated after a number of promising broadband ground-motion simulation methods were developed (e.g., Atkinson and Boore, 2003; Boore and Atkinson, 2008; Graves and Pitarka, 2010; Mai *et al.*, 2010; Schmedes *et al.*, 2010). Researchers were interested in comparing ground-motion simulation results between methods and in comparing simulation results to observations and to GMPEs. Such comparisons can be difficult because alternative methods often require different input parameters or specify input parameters in dissimilar formats. Fair comparisons between methods are only possible once multiple methods can input equivalent problem definitions and output comparable results. Based on previous experience evaluating simulation methods, seismologists and earthquake engineers recognized a need for a software tool to support the evaluation of the newly emerging broadband simulation methods. Drawing on discussions with both seismologists and engineers (e.g., Jordan and Maechling, 2003; Maechling *et al.*, 2009; Bielak *et al.*, 2010; Somerville *et al.*, 2011), we identified the following list of essential capabilities of the required tool. The BBP should do the following:

- implement multiple ground-motion simulation methods capable of outputting ground-motion seismograms;
- enable multiple ground-motion methods to run exactly the same problem and produce directly comparable results;
- produce repeatable results;
- provide seismogram postprocessing methods to support analysis of simulation results;
- be extensible, so that ground-motion simulation methods and postprocessing software utilities can be easily added or modified;
- be constructed using open-source software tools, such as the GNU compilers, and run in a Linux operating system computing environment; and
- be released as versioned open-source software so users can examine the source code and results can be attributed to a specific version of the platform.

The goals of the Southwestern United States Ground Motion Characterization (SWUS) and Next Generation Attenuation for Central and Eastern North America (NGA-E) projects described in this special focus issue added additional BBP requirements. For SWUS and NGA-E, it was important that the BBP be capable of the following:

- separate common input parameters, such as the source description and velocity model used by all methods, from method-specific parameters used by individual methods;
- allow each method to define region-specific configuration parameters and provide tools to determine which configuration parameters work best for a given region;
- support comparison of simulated ground-motion amplitudes to observed ground-motion amplitudes;
- support comparison of simulated ground-motion amplitudes to ground-motion amplitude ranges given by existing GMPEs;
- support large suites of scenario ground-motion simulations as needed in evaluation processes by running many

simulations quickly using high-performance computers; and

- support collaborative scientific and engineering ground-motion simulation evaluation processes.

In the remainder of this article, we describe how we designed and implemented the BBP to meet these user requirements.

BROADBAND PLATFORM CONCEPTS AND TERMINOLOGY

A set of concepts and related terminology has developed around the BBP. These terms, once defined, make communications about the platform easier. To begin, we describe the BBP as a platform to communicate the fact that the BBP is more than a single scientific application program. In SCEC terminology, a "computational platform" is a well-integrated collection of validated scientific software and geophysical models that can perform a useful research calculation. The BBP implements a component-based software architecture. Within the BBP, a "component" is a scientific computer code or program that has been integrated into the platform, accepts BBP format inputs, and produces BBP format outputs. A "method" is a collection of one or more components that can perform a ground-motion simulation. "Workflows," or "pipelines," are equivalent concepts that refer to a sequence of calculations that may include both ground-motion simulations and validation processing, in which the output of one component is used as input to the next processing stage. A "simulation" is the basic unit of work for the BBP. A simulation represents an invocation of the platform using a specific source description, site list, velocity model, and method that produces one or more output seismograms. Practical use of the BBP typically requires running many simulations. A "scenario" is a general description of an earthquake as observed by a particular set of stations and for a specified velocity model. A scenario defines the earthquake magnitude, fault dimensions (length, width, depth, and segmentation), and fault geometry (strike, dip, and average rake). A scenario does not define the details of the rupture history (slip distribution, slip time function, and rupture speed) and does not define hypocenter location. Earthquake source definitions, given as a part of a scenario, are elaborated into distinct source realizations. Source realizations may differ by initial stress distribution, slip distribution, slip time function, hypocenter location, and rupture velocity. Source realizations are created in the platform by changing source parameters, a numerical seed parameter, or both, in the source definition file. In BBP usage, "verification" involves comparison against expected results, whereas "validation" involves comparison against observations. Figure 1 shows an example of a BBP validation workflow to illustrate the concepts we are presenting. Using these concepts, we can more easily describe the capabilities of the current BBP system.

DATA FLOW PROCESSING MODEL WITH FILE-BASED DATA EXCHANGE FORMATS

The BBP was constructed by integrating a collection of independently developed scientific software programs (e.g., Zeng *et al.*, 1994; Atkinson *et al.*, 2009; Graves and Pitarka, 2010, Mai *et al.*, 2010; Schmedes *et al.*, 2010). The availability of these existing open-source scientific codes led to a modularized BBP system architecture that separates BBP processing into multiple small stages. By assembling the BBP from multiple software programs, we avoid a monolithic software program that is difficult to understand and modify. However, with this modular approach, multiple software programs must be used to implement a basic BBP ground-motion calculation, and most of the independently developed scientific codes used in the BBP

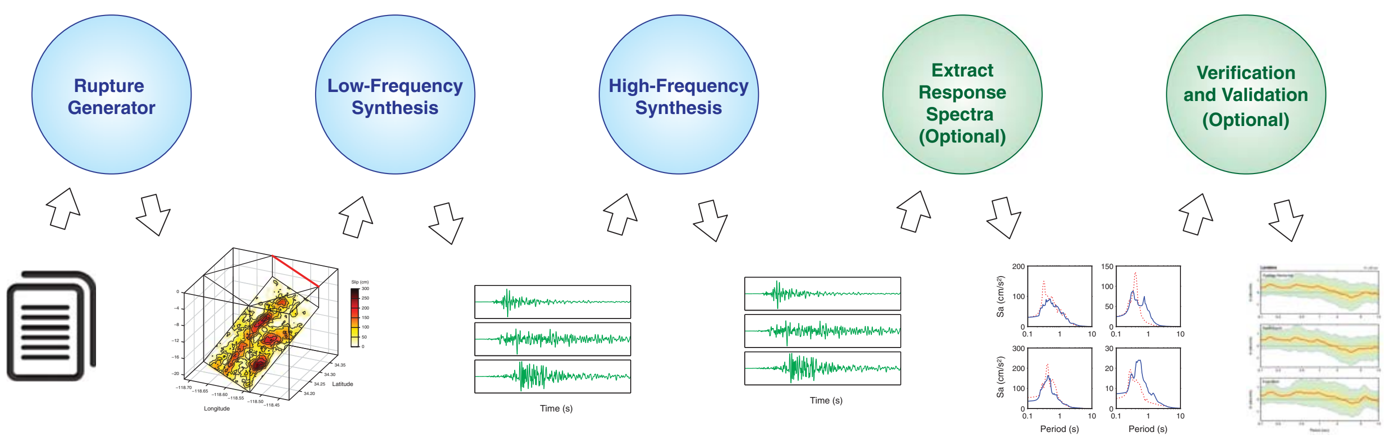

▲ **Figure 1.** Overview of a Broadband Platform (BBP) simulation workflow showing how a ground-motion simulation is implemented in three processing stages (blue circles), followed by two stages of postprocessing (green circles). The ground-motion simulation method shown is implemented using three components: a rupture generator, a low-frequency ground-motion simulation, and a high-frequency ground-motion simulation. In the first postprocessing stage, ground-motion time series are converted to peak pseudospectral acceleration (PSA) at different periods. In the second postprocessing stage, BBP validation processing compares observed ground motions against ground-motion simulation results.

were not developed with interoperability in mind. Our basic approach for making the BBP codes interoperable is to define a BBP calculation as a specific sequence of programs in which files output by one processing stage are used as inputs by subsequent processing stages. BBP developers describe this as a dataflow architecture with file-based data dependencies. BBP users may refer to this as a workflow or pipeline processing approach.

As mentioned above, the BBP uses file-based data exchange between processing stages. File-based communication between programs is slower than several other possible approaches, but it is easily implemented and reduces the need to modify the components involved. The BBP defines standard text-based file formats as interfaces between components. BBP components are configured to read files in BBP format and to output results in BBP file format. Internally, components can use their own input and output formats and may use method-specific input files such as Green's functions (GFs), but they must convert their internal formats to BBP formats at component interfaces. This format conversion approach works as long as native file formats contain the information required by the BBP standard format. Ⓔ Examples of standard BBP file formats are available in the electronic supplement to this article. A list of standard BBP file formats is provided below:

- *Source (SRC):* A source file contains simple rectangular fault geometry, an earthquake magnitude, and a hypocenter location. There are no time parameters, such as origin time or rise time, in the SRC file source description.
- *Site list (STL):* A site list contains a list of one or more sites for which ground motions will be calculated. Both the site list and the SRC file contain geographical locations, and the relationship between the SRC hypocenter location and the site locations determines the distance from the event to the sites of interest.
- *Velocity model (VMOD):* A velocity model file contains a list of velocity layers, their thickness, and V_P, V_S, density, Q_P, Q_S for each layer, assumed to be in the geographical region of interest. For methods that use GFs, the input 1D velocity model must match the 1D velocity model used to generate the GFs.
- *Standard rupture format (SRF):* A rupture generator component converts a simple SRC file into a SRF file based on seismological rules implemented in the rupture generator. The SRF format describes the earthquake source in a time-series format.
- *BBP seismograms:* The BBP outputs ground-motion seismograms in a three-component time-series format. BBP seismogram timestamps are given relative to event origin time, not arrival time or absolute time, as found in most observational data formats. The BBP seismogram format is used for multiple ground-motion types, including acceleration, velocity, and displacement. BBP seismogram format is used as the standard input and output format for many BBP components and utilities.

In addition to these BBP exchange formats, there are two more standard file formats used within the BBP. These two text formats are important to the operation of the BBP platform, but they are not, strictly speaking, data exchange formats.

- *XML workflow:* The BBP defines each simulation in an XML workflow format that defines the components that will be run, and the input files to be used in each stage. The BBP XML format defines a directed acyclic graph, a common workflow definition form.
- *Observational ground-motion time-series data:* BBP validation processing requires management of observational data sets. The BBP can process observed ground-motion seismograms in Pacific Earthquake Engineering Research Center seismogram format, or alternatively, in BBP seismogram format.

The BBP implements its multistage computing capabilities using two important software constructs. First, the BBP defines an XML workflow specification format to describe the input files and computer programs needed to run a ground-motion simulation and any required postprocessing stages. Second, the BBP provides a software framework that interprets these XML workflow specifications and executes the given list of programs, using the given input files, in the correct order. With these two constructs, the BBP can flexibly support multistage ground-motion simulation calculations, making it easy to add or remove computational stages without changes to the BBP infrastructure. BBP XML files are important metadata for BBP calculations because they define the exact inputs and programs used to calculate a particular result. Given a working copy of the BBP software and the required input files, a BBP XML workflow specification can be used to rerun a specific BBP ground-motion simulation. The number and type of processing stages in a particular BBP simulation depends on the specific research calculation defined in the XML file. Figure 2 shows the data file inputs and processing components defined in a BBP XML workflow specification needed to implement the BBP workflow in Figure 1. Ⓔ An example of a BBP XML workflow specification is available in the electronic supplement to this article.

SCIENTIFIC SOFTWARE COMPONENTS

The BBP performs its processing by running a series of BBP components. Each component accepts file-based inputs in one of the BBP standard formats and outputs the results in one of the BBP standard formats. A summary of current BBP ground-motion simulation methods, showing the components used to implement each method and two common postprocessing components, is provided in Table 1.

BBP components are implemented, at their outermost layer, as a Python language software program. An example BBP component is shown in Figure 3. Some BBP components, such as postprocessing utilities, are written completely in Python. For many more-complex BBP components, a BBP Python language software program calls a software executable written in a compiled programming language, such as C or FORTRAN. In this case, we call the Python language program a "Python wrapper," and the BBP component is composed of the Python wrapper and

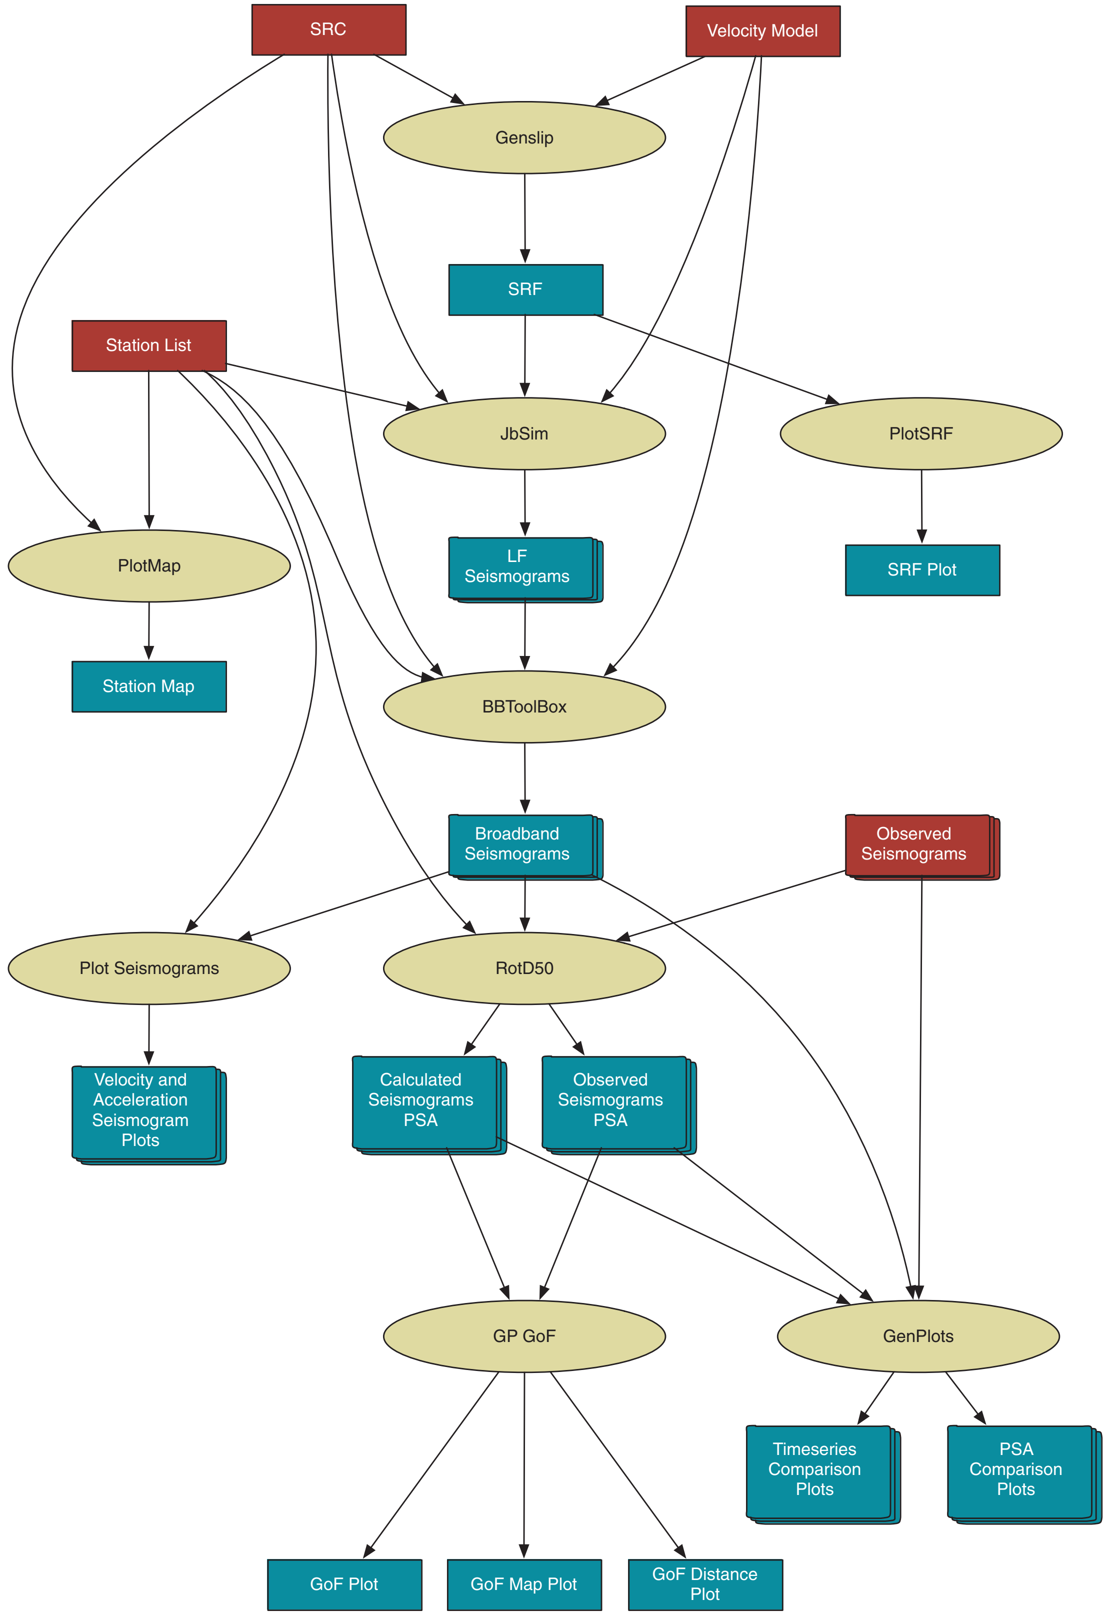

▲ **Figure 2.** This workflow diagram shows how a BBL XML workflow specification defines a BBP validation simulation based on the SDSU method. Input files needed for a validation simulation are shown as red boxes, BBP components are shown as yellow ovals, and BBP output files are shown as blue boxes. The SDSU method uses the rupture generator and low-frequency components of the Graves and Pitarka (GP) method but also includes its own high-frequency component that generates broadband seismograms. Output validation data products produced by the BBP are produced using software that is shared by all methods.

Table 1
Broadband Platform (BBP) Ground Motion Simulation and Postprocessing Components

<table>
<tr><th>BBP Ground Motion Method Name</th><th>Rupture Generator</th><th>0–1 Hz Motions</th><th>1–100 Hz Motions</th><th>Time-Series Postprocessing</th><th>Goodness-of-Fit Postprocessing</th></tr>
<tr><td>GP</td><td>Gen_Slip</td><td>JB_Sim</td><td>HF_Sim</td><td>RotD50</td><td>Bias Plot</td></tr>
<tr><td>SDSU</td><td>Gen_Slip</td><td>JB_Sim</td><td>BB_Toolbox</td><td></td><td></td></tr>
<tr><td>UCSB</td><td>UCRMG</td><td colspan="2">Syn1D</td><td></td><td></td></tr>
<tr><td>EXSIM</td><td></td><td>EXSIM</td><td></td><td></td><td></td></tr>
<tr><td>CSM</td><td></td><td>Simula</td><td></td><td></td><td></td></tr>
</table>

Table columns describe BBP ground-motion method names and three ground-motion simulation stages and two common postprocessing stages. Table rows show how selected BBP simulation methods are implemented in one or more processing stages within the BBP and how the common postprocessing programs are used to analyze results from all simulation methods

the compiled language program together. Python wrappers are an important design technique used to build the platform. By constructing each component as a Python language program, the overall BBP framework calls only Python programs, hiding the use of other programming languages and executables. Wrappers enable us to integrate programs in a variety of languages into the platform with a minimum of modifications. To reduce modifications to scientific programs, component wrappers often perform file format conversion processing, converting BBP format inputs to the native input format of the scientific codes, then converting the native output format of the scientific code to BBP format before the component processing completes. In this way, minimum modifications to the original scientific codes are required.

To accrue the benefits of this software architectural approach, we accept a number of disadvantages. Reformatting input files requires additional processing time, increases the disk-based input/output requirements of the platform, and increases the data storage requirements by introducing duplicate files in alternative formats. We currently accept these inefficiencies, but recognize that BBP performance can be improved in each of these areas by minimizing duplicate files and processing in future versions.

Within the platform, we distinguish between BBP components used to implement ground-motion simulation methods and BBP components used to implement common postprocessing stages. Examples of BBP components used to implement ground-motion methods include rupture generators, low-frequency deterministic waveform processing, and high-frequency stochastic components. As described earlier, a method may be implemented by combining components from different sources. For example, the Graves and Pitarka (GP) low-frequency ground-motion simulation component is used in both the GP method (Graves and Pitarka, 2010), when combined with the GP stochastic

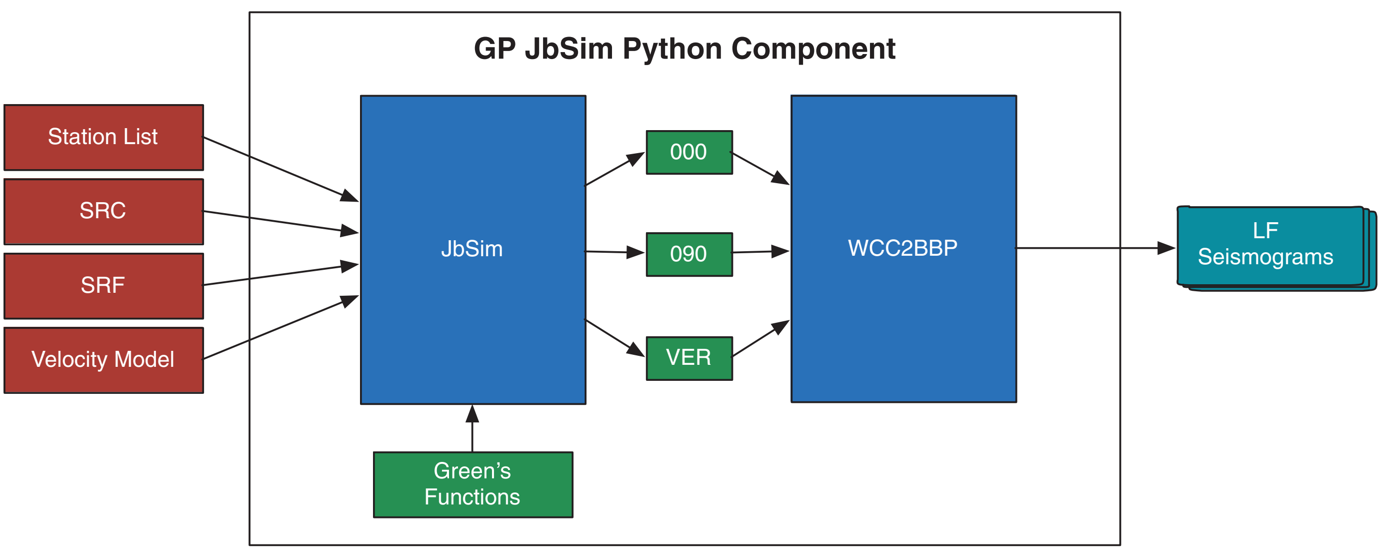

▲ **Figure 3.** Example of how the BBP component-based architecture is implemented by wrapping scientific codes, such as JbSim and WCC2BBP, into a Python language component, called GP JbSim. Only Python language components are called with BBP simulation workflows, hiding the use of the scientific codes written in other programming languages. BBP components input and output BBP format files (left, red rectangles; right, teal blue rectangles), but they may use internal file formats (center, green rectangles) to perform their processing.

high-frequency component, and in the San Diego State University (SDSU) method (Mai *et al.*, 2010), when combined with the SDSU stochastic high-frequency component.

All BBP methods produce output seismograms. In some BBP uses, such as scenario simulations, seismograms are the desired end product. However, if additional postprocessing is desired, the BBP provides a set of common postprocessing programs, and the same postprocessing programs are used regardless of the ground-motion simulation methods involved. Currently, the BBP provides common postprocessing components that can integrate and differentiate time series, components that can calculate amplitude values such as peak acceleration, spectral acceleration at various periods, and specialized ground-motion amplitudes such as RotD50.

Several BBP methods make use of random numbers as they model stochastic processes. Use of random numbers can lead to difficulty providing reproducible results. To support both the use of random numbers and reproducible results, the BBP makes use of numerical seed values defined in input SRC files to initialize a random number generator used in BBP calculations. As long as the same numerical seed value is used, the random number generator used by the BBP will produce the same sequence of random numbers. Currently within the platform, numerical seed values are used to (1) randomize the slip distribution in the rupture generators, (2) seed the scattering wavelets calculation in the SDSU high-frequency component, (3) provide random numbers for the GP high-frequency component, and (4) initialize some aspects of extended finite-fault simulation (EXSIM) processing. Details of how numerical seeds values are used are method dependent. From the platform's perspective, numerical seed values are used as inputs to methods to provide both reproducibility of results and variability in results. If the platform runs a method twice using the same seed value, the platform produces the same results. If the platform runs a method twice using different seed values, the platform produces different results.

BBP DATA PRODUCTS

Users can view the BBP as a software black box that inputs configuration files, consumes computer-processing time, and outputs file-based data products. The file-based data products produced by the BBP vary, depending on the type of simulation the user performs. When a user runs a BBP simulation, the BBP assigns a unique simulation ID (or SimID) to the simulation and creates a data directory using that SimID to organize the inputs, temporary files, and outputs produced by that simulation. Output data products for a simulation are found in the output data directory when the simulation run completes. Organizing results into simulation-specific directories separates results from different simulations and ensures that results from one simulation do not overwrite the results from a previous simulation.

We can categorize output BBP data products into two types: standard data products and comparison data products. Standard BBP format output data products can be produced for both validation and scenario simulations, and some standard output products are optional and user selectable. Standard BBP output data products include seismogram files in BBP format, seismogram plots in PNG format, SRF slip plots, spectral response plots, and fault and station maps. Comparison data products are produced when the platform is used in validation mode. In this case, the platform will produce the standard data products, as well as data files and plots that compare simulation results against observed data. Comparison data products include seismogram comparison plots, goodness-of-fit (GoF) bias plots, and GMPE comparison plots. Examples of BBP data products for scenario simulations are shown in Figure 4. Figure 5 displays BBP data products for validation simulations, and Figure 6 shows GoF data products that the BBP can produce for validation simulations. Ⓔ Additional examples of BBP output data products are available in the electronic supplement to this article.

BROADBAND DISTRIBUTION AND INSTALLATION

SCEC distributes the BBP as a collection of Python language scripts, C and FORTRAN language source codes, configuration files, GF data files, test cases, and observational data files. Most of the scientific codes used within the BBP must be compiled before the platform can be used. The BBP distribution separates the Python codes and compiled programs into separate directories. The platform distribution includes a top-level makefile that invokes a series of method-specific makefiles that build all required executable programs when run on a properly configured Linux system. The target development and runtime environment for the BBP is a 64-bit Linux computing environment with GNU and Intel compilers and tools. The BBP platform uses an open-source software policy to provide transparency into the scientific codes, ensuring that users can review and test any of the methods or software utilities within the platform. When SCEC releases a new BBP distribution, users need to retrieve the new distribution from the BBP website (http://scec.usc.edu/scecpedia/Broadband_Platform; last accessed November 2014), install the platform on their Linux computer, configure the installation properly, build the executables, and test the new installation before using it for research purposes.

Currently, new releases of the BBP are distributed in multiple separate TAR files that we call "packages." We have separated the BBP into multiple packages because the size of all packages combined exceeds 8 GB and because not every BBP user needs all the packages. We assign version numbers to SCEC software using a date-based versioning scheme in a Year.Month.Revision-Number format. So, BBP v13.6.0 is a version of the BBP software released in June 2013. BBP v13.6.1 is the first revision of the v13.6.0 release. Each BBP distribution package (source code package, velocity model package, and validation package) has its own separate version number. This allows us to update individual packages without having to rerelease the entire set of files every time a modification is made. The BBP source file package version number identifies the official BBP version, such as v13.9.0 or v14.3.0. In some cases, we associate a version of

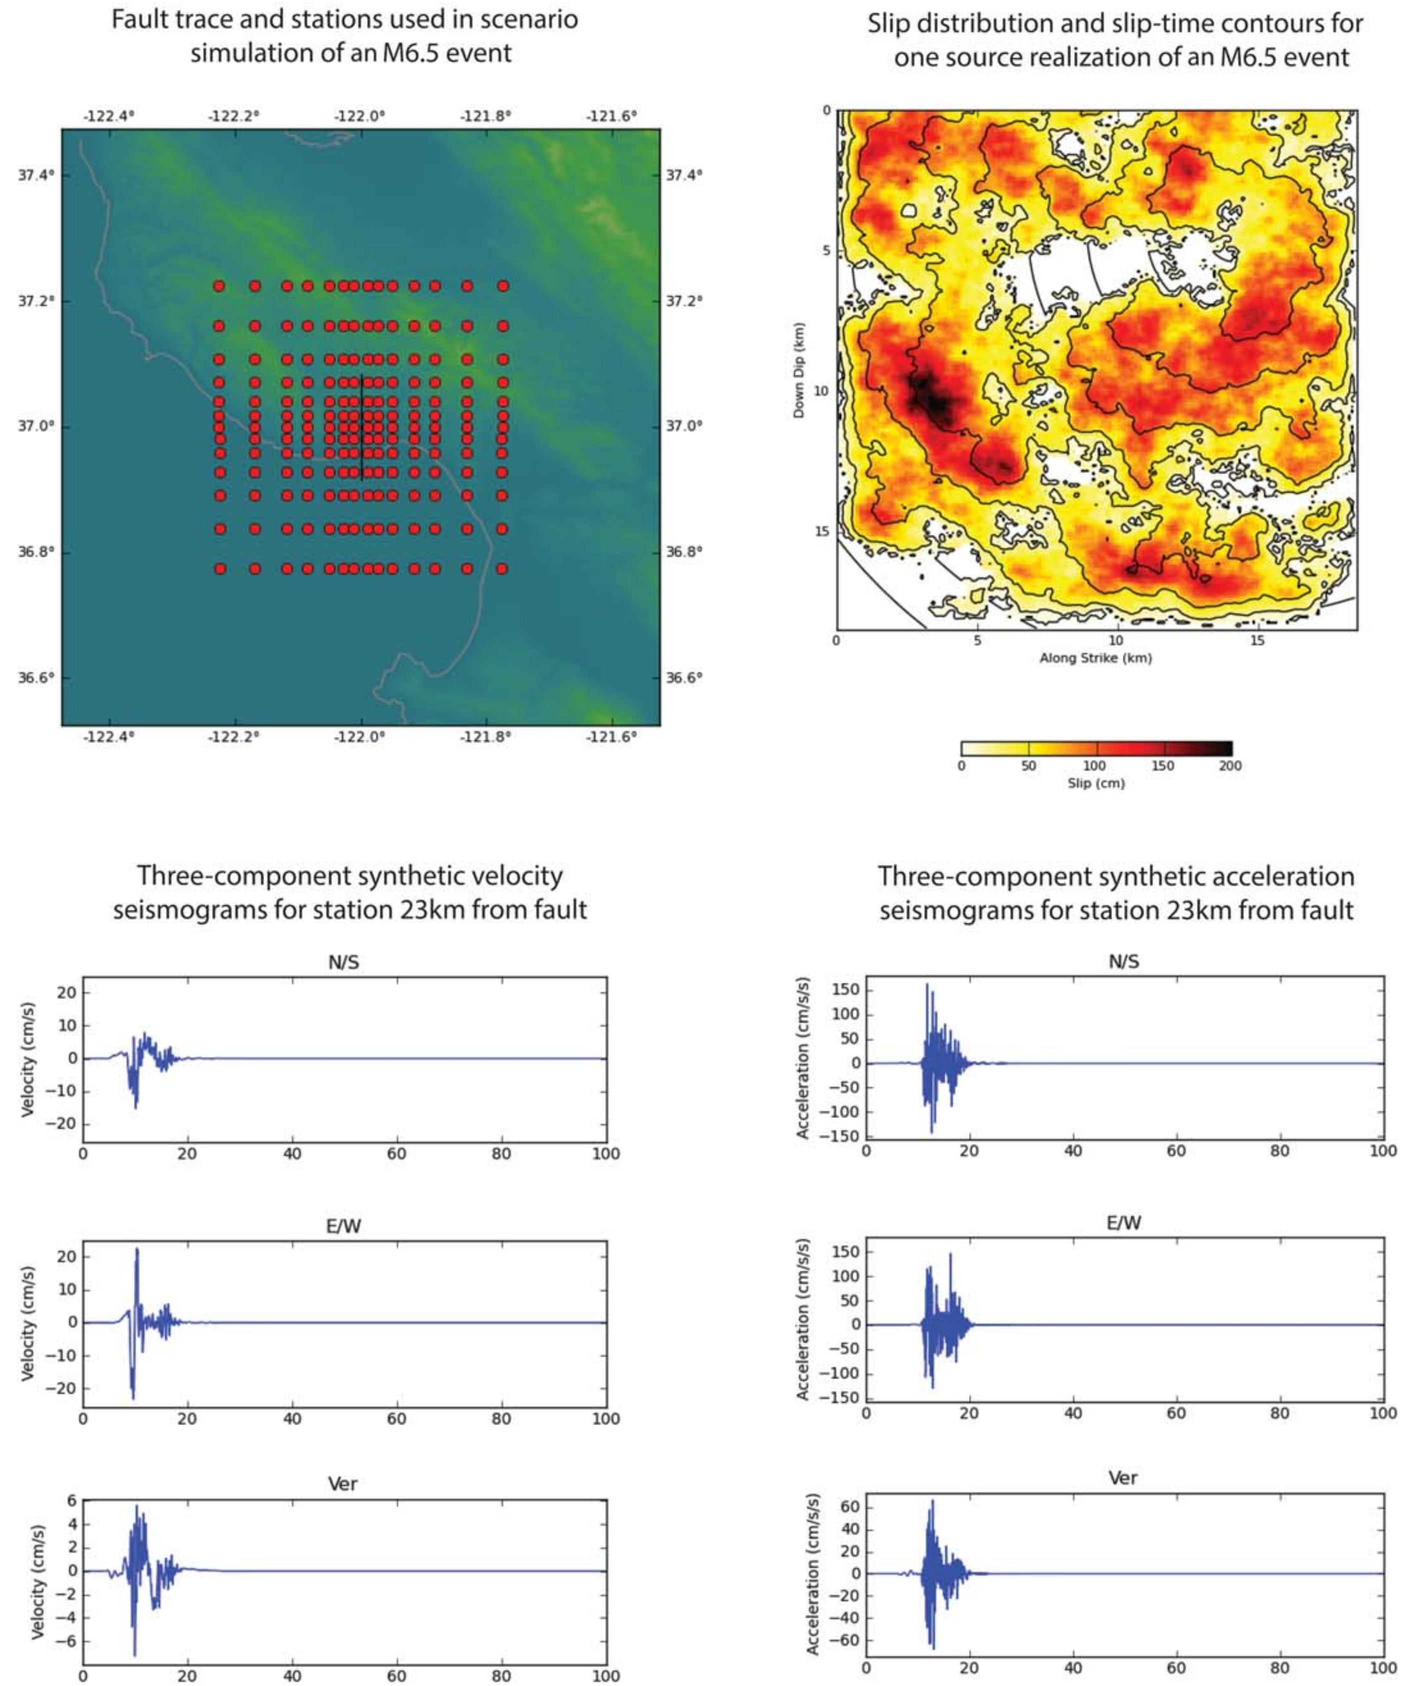

▲ **Figure 4.** Examples of output data products produced by a BBP scenario simulation. Output BBP data products include a map showing a projection of the fault plane on the surface and the location of stations (top left), a standard rupture format slip plot (top right), and ground-motion time-series data as velocity (bottom left) and as acceleration (bottom right).

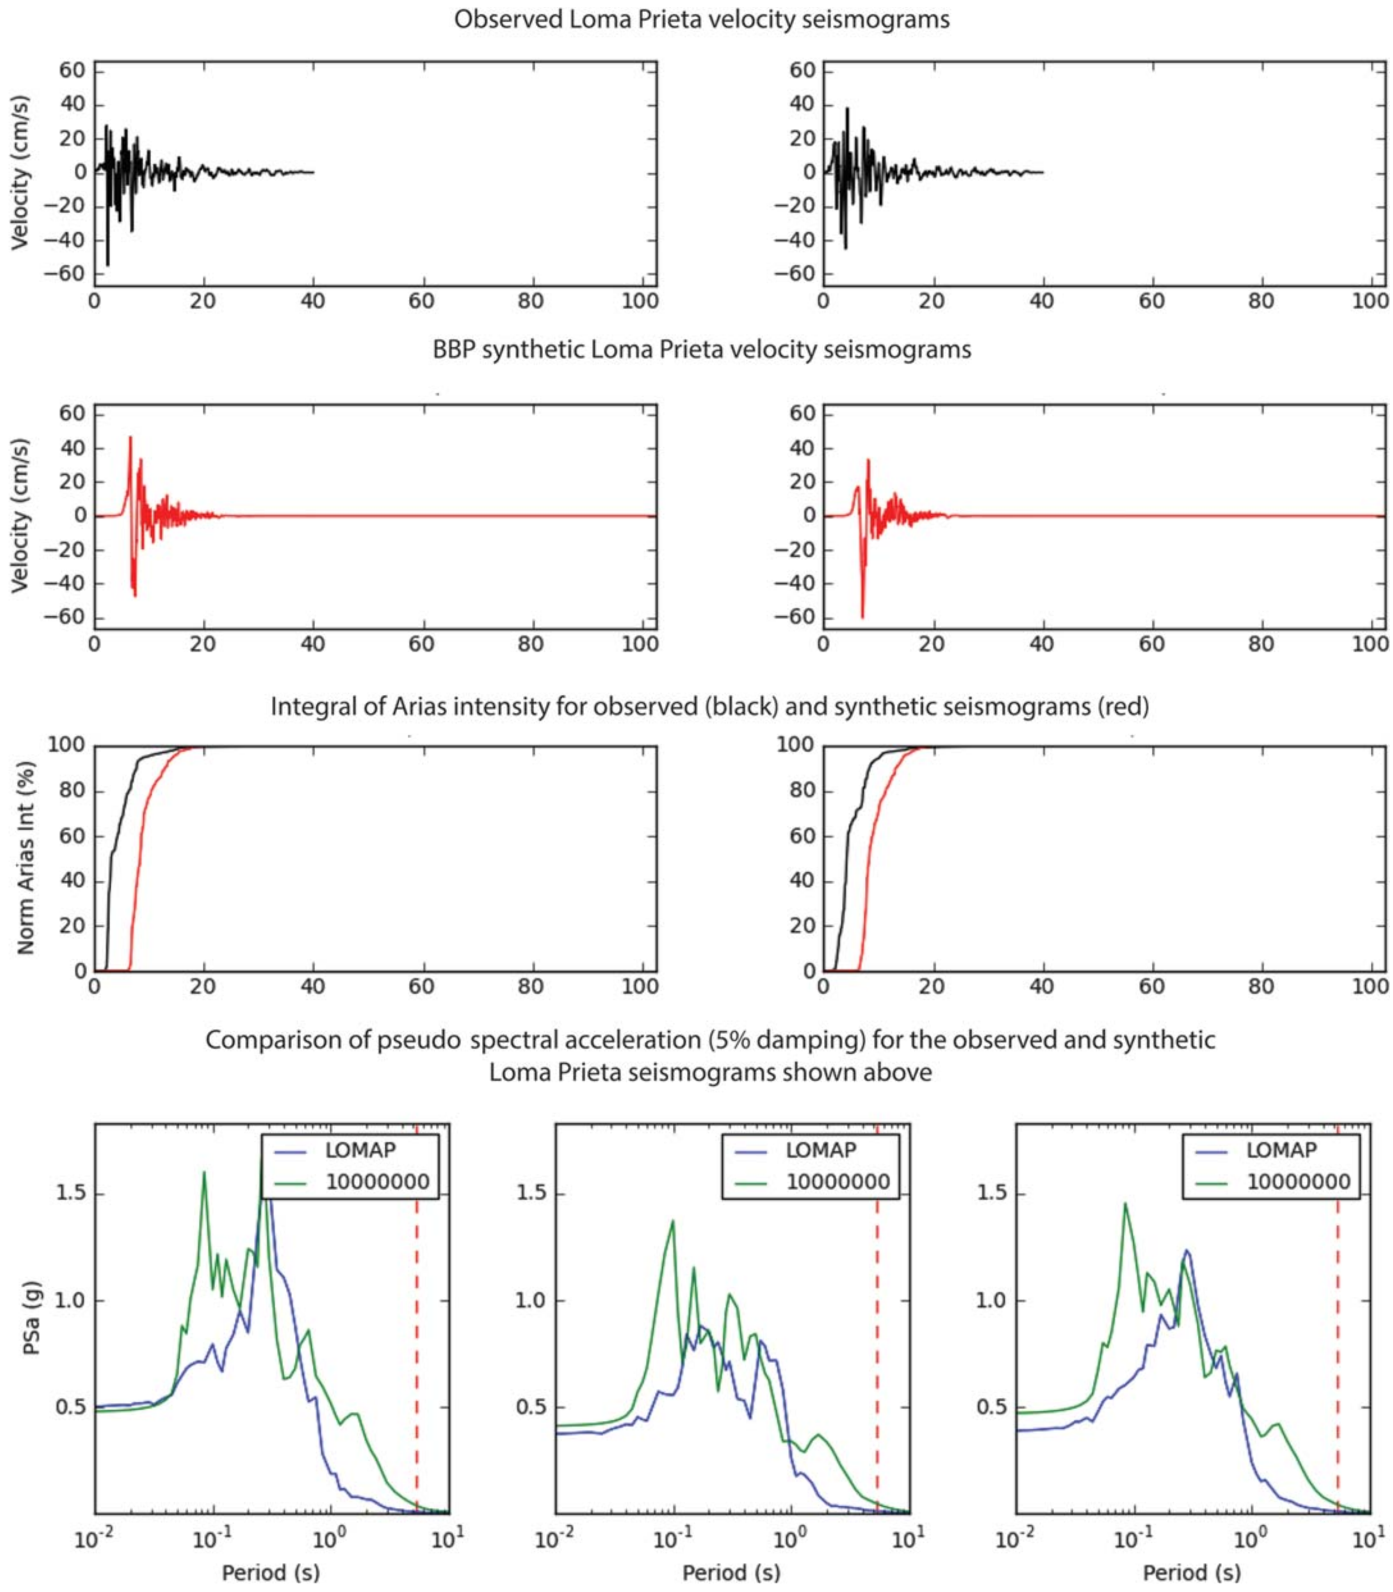

▲ **Figure 5.** Examples of output data products produced by a BBP validation simulation. When the BBP is used to simulate well-observed historical earthquakes, the output BBP data products can include plots showing comparisons between observed and simulated ground-motion time series and the integral of Arias intensity as a function of time plots for both the simulated and observed seismograms (top). The BBP can also produce PSA (5% damping) comparison plots that compare ground-motion simulation time-series PSA response against the pseudospectral response of observed historical earthquake recordings (bottom). The vertical red lines in these plots indicate the maximum long-period response of the instruments that recorded the observed seismograms.

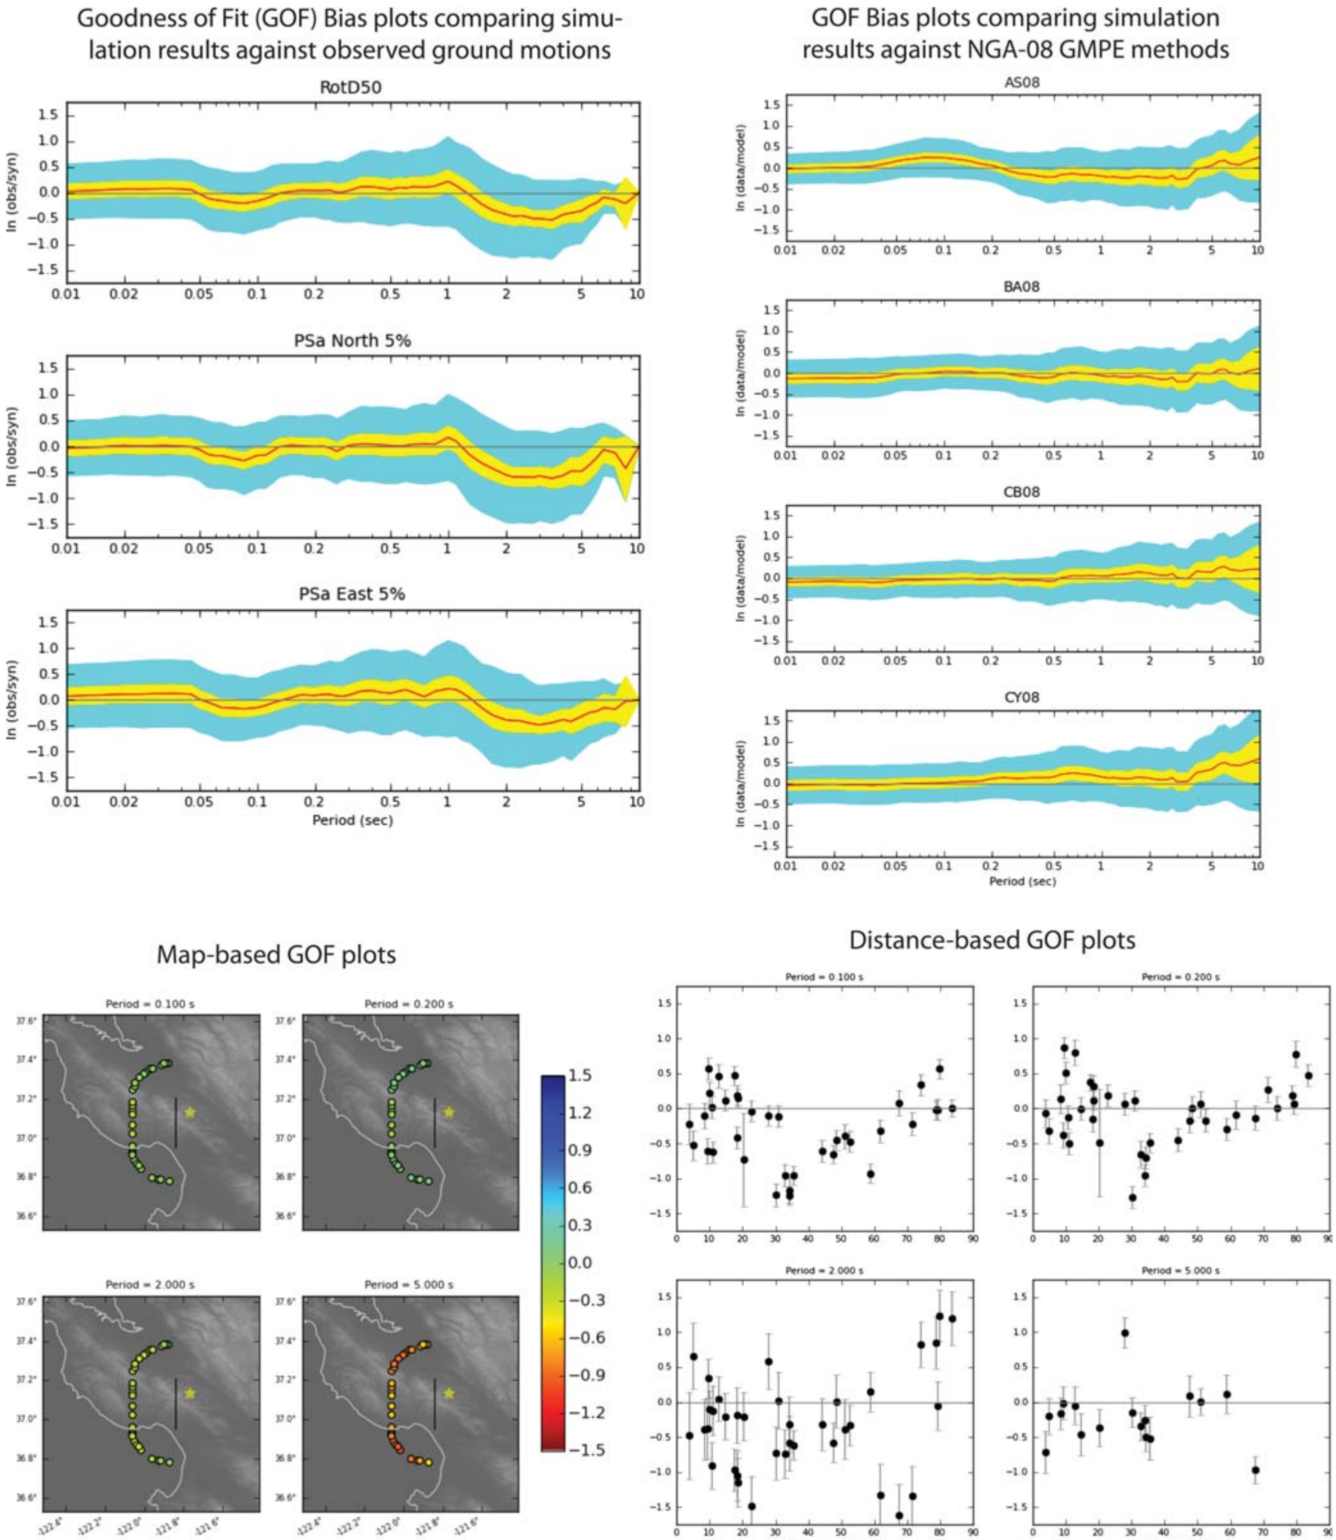

▲ **Figure 6.** Examples of goodness-of-fit (GoF) plots generated by the BBP, comparing ground-motion simulation results against observed earthquake ground-motion recordings and ground-motion prediction equation (GMPE) models. The BBP can generate GoF by period, distance, and location and can plot results from the 2008 Next Generation Attenuation (NGA-08) GMPE models (Abrahamson and Silva, 2008; Boore and Atkinson, 2008; Campbell and Bozorgnia, 2008; Chiou and Youngs, 2008) against observed and simulated ground-motion time series. Bias plots compare the spectral response of the BBP simulations from short-to-long periods (0.01–10 s) against multiple ground-motion observations (top left) and against NGA-08 GMPE models (top right). Map-based GoF plots (bottom left) and distance-based GoF plots (bottom right) help evaluate directional amplitude bias at multiple periods.

the Broadband Platform with a validation and review exercises. For example, the BBP v13.6.1 release was used in the 2013 SWUS validation exercise, and the BBP v14.3.0 release was used in the NGA-East validations exercise (Dreger *et al.*, 2014).

The BBP source code package contains the main Broadband Platform files, including the scientific codes, method-specific data and configuration files, plotting utilities, and testing codes used to ensure the Broadband Platform is properly installed and generating expected results. Along with the source code packages, users need to install at least one velocity model package. The BBP v14.3 distribution includes one velocity model package for each of the supported regions, which currently include the Los Angeles Basin, Mojave, northern California, western Japan, central Japan, eastern United States, and eastern Canada. In addition to the standard BBP format velocity model file defined for the region, each of these packages includes GFs for the GP (Graves and Pitarka, 2010) and the UCSB (Schmedes *et al.*, 2010) methods, region-specific files required by other methods, such as EXSIM (Motazedian and Atkinson, 2005; Atkinson *et al.*, 2009) and composite source model (CSM) (Zeng *et al.*, 1994), and a configuration file that lists all region-specific parameters for each of the scientific methods. Because velocity model packages are large, around 1–2 GB each, each region is distributed separately, allowing users to download and install only the packages they need. BBP users need to install validation packages only if they want to run validation simulations. Validation packages contain carefully selected observational data recorded for the earthquake being studied. Each BBP validation package is distributed separately so users only need to download the validation events of interest. The full procedure for installing a BBP distribution is described in the BBP users manual, available on the SCEC BBP website.

Once the required set of BBP files are installed, the BBP platform uses two types of automated software test suites (unit tests and acceptance tests) to ensure the platform is working correctly before is it used in a research calculation. Once the unit test suite passes, it confirms that the BBP has been installed correctly and that individual components are producing the same results on the deployment computer as they produced on the development systems at SCEC. When the BBP acceptance tests pass, it shows that, at a minimum, the platform has run one end-to-end earthquake simulation using each method, and each method produces ground-motion amplitudes that match the development system results within acceptable tolerance established by scientific and engineering users at the time the code was released.

In general, we recommend researchers use the most recent public release of the BBP software because the ground-motion models and computational infrastructure continue to improve rapidly. In some cases, however, users may want to reproduce results from versions of the platform that have been reviewed by an evaluation panel. SCEC maintains an archive of previously released versions of the BBP platform to support such usage.

RUNNING THE BBP

Once the platform is installed and running correctly, as established using the unit tests and acceptance tests, a BBP user can configure and run a ground-motion simulation through a text-based interface that presents the user with a series of questions and creates a BBP workflow based on the user's responses. Once created, the user can run the workflow immediately or save the workflow as an XML file, which can then be used to run the calculation repeatedly without respecifying the workflow. The BBP software is designed to separate the command line user interface, used to construct the workflow, from the workflow engine that runs the processing sequence defined in the XML file. The example command line (below) shows how a user can run a BBP simulation that will execute the processing sequence defined in a specific BBP XML workflow file. `$ run_bbp.py -x northpalmsprings-gp-0000.xml`

Currently, all the programs in the BBP are single-threaded, serial programs that run on 64–bit × 86 computers. The platform has been run successfully on computing clusters, such as the USC high performance computing cluster, using unmodified XML files. When BBP calculations are run on a computer cluster, the cluster is treated as a collection of stand-alone Linux computers. Modern computer clusters often have multiple nodes, the nodes contain multiple processors, and the processors contain multiple cores. On such a cluster, many BBP simulations can be run simultaneously, typically one BBP simulation per core. Use of a cluster can speed up a BBP calculation that involves many simulations, but it does not speed up individual BBP simulations, because each simulation is implemented as a serial process.

BBP DEVELOPMENT AREAS AND PLANNED IMPROVEMENTS

From the software development perspective, close collaboration with scientific and engineering users has greatly improved the capabilities of the BBP, making the software development more efficient, and helping to prioritize the software development tasks. So, we expect that future BBP development will be driven by the needs of the emerging BBP user community. In the following sections, we describe several planned BBP improvements based on recent discussions with scientific and engineering groups.

With new ground-motion models under development and existing models continuing to improve, we expect to integrate new ground-motion simulation components into the BBP, including alternative rupture generators and new ground-motion simulation methods that perform all processing stages. Computational methods that include nonlinear site response processing stages are also under development. We also expect to add additional validation events into the platform. Adding a new validation event to the BBP typically involves developing scientific consensus on how to parameterize the source description, how to define the 1D velocity model, and how to select ground-motion observations to be used in GOF evaluations.

Engineering users would like simulation methods within the BBP to model multi-segment ruptures, which are often needed to simulate very large earthquakes. Scientists and engineers continue to develop and refine evaluation procedures for engineering use of ground motion simulations, so we expect additional post-processing and GoF methods will be required in the future. Several existing BBP methods depend on precalculated GFs, defined for specific 1D velocity models, and BBP simulations are currently limited to using these predefined, provided, 1D models. We plan to remove this limitation by integrating code into the BBP that can calculate GFs for arbitrary 1D models. At some point in the future, the BBP may be modified so it can make use of 3D low-frequency seismograms, rather than the current 1D low-frequency seismograms. However, 3D ground-motion simulations require high-performance computers (HPC). So, initial BBP support for 3D simulations is likely to be based on importing externally calculated 3D seismograms rather than on integrating 3D HPC simulation codes into the platform.

We have identified several software engineering improvements that we believe will improve the BBP user experience. Anticipated BBP improvements include the following:

- refine and improve the standard BBP data exchange file formats to improve their usability;
- document the individual components within the platform so the components are easier to use as stand-alone programs;
- develop a more flexible interface for constructing workflows that improves upon the current command line interface;
- distribute the platform in ways that make it easier for users to install and run the BBP in their own computing environment;
- help users contribute new methods by clearly defining the required input parameters and providing example results;
- establish a persistent, searchable, data archive of BBP releases and associated evaluation results; and
- provide researchers with complete descriptions of standard reference ground-motion simulation problems that can be used to evaluate new ground-motion simulation methods.

CONCLUSIONS

The SWUS and NGA-E ground-motion simulation evaluation processes show that the BBP successfully enables nonseismologists to run multiple, complex simulation methods, using a specific software version for each code, and to process results from all methods using a common suite of postprocessing algorithms for the generation of statistics and plots. As a result of a careful software integration and evaluation process, the BBP software now allows users to run simulations without detailed knowledge of how a particular method is implemented, making it possible for nonscientists and third parties to run ground-motion simulation models and produce useful results. The BBP also provides a working implementation of ground-motion simulation evaluation processes defined by seismologists and earthquake engineers to qualify ground-motion simulations for engineering use. By providing these tools and capabilities, the BBP has developed into a system capable of supporting the scientific and engineering ground-motion simulation evaluation processes needed to widen engineering use of ground-motion simulations.

ACKNOWLEDGMENTS

The authors would like to thank the Broadband Platform Validation Group, particularly group leaders Christine Goulet and Paul Somerville, for valuable direction and feedback on Broadband Platform software development. This research was supported by the Southern California Earthquake Center (SCEC). SCEC is funded by National Science Foundation (NSF) Cooperative Agreement EAR-1033462 and U.S. Geological Survey Cooperative Agreement G12AC20038. This research was supported by the NSF Software Infrastructure for Sustained Innovation program under award SI2-SSI: A Sustainable Community Software Framework for Petascale Earthquake Modeling (OCI-1148493). Major funding for development of, and improvements to, the Broadband Platform was provided by the Pacific Gas and Electric Company. Computation for the work described in this article was supported by the University of Southern California's Center for High-Performance Computing. Additional computation described in this article used computing resources provided by NSF XSEDE allocation titled SCEC Petascale Research: An Earthquake System Science Approach to Physics-Based Seismic Hazard Research (TG-MCA03S012). This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF Grant Number ACI-1053575. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin and the San Diego Supercomputer Center (SDSC) at the University of California, San Diego, for providing high-performance computers and data management resources that have contributed to the research results reported within this article. The SCEC contribution number for this article is 1981.

REFERENCES

Abrahamson, N. A., and W. J. Silva (2008). Summary of the Abrahamson & Silva NGA ground-motion relations, *Earthq. Spectra* **24,** no. 1, 67–97.

Atkinson, G. M., and D. M. Boore (2003). Empirical ground-motion relations for subduction-zone earthquakes and their application to Cascadia and other regions, *Bull. Seismol. Soc. Am.* **93,** no. 4, 1703–1729.

Atkinson, G. M., D. M. Boore, K. Assatourians, K. Campbell, and D. Motazedian (2009). A guide to differences between stochastic point-source and stochastic finite-fault simulations, *Bull. Seismol. Soc. Am.* **99,** 3192–3201.

Bielak, J., R. W. Graves, K. B. Olsen, R. Taborda, L. Ramírez-Guzmán, S. M. Day, G. P. Ely, D. Roten, T. H. Jordan, P. J. Maechling, J. Urbanic, Y. Cui, and G. Juve (2010). The ShakeOut earthquake scenario: Verification of three simulation sets, *Geophys. J. Int.* **180,** no. 1, 375–404, doi: 10.1111/j.1365-246X.2009.04417.x.

Boore, D. M., and G. M. Atkinson (2008). Ground-motion prediction equations for the average horizontal component of PGA, PGV, and 5%-damped PSA at spectral periods between 0.01 s and 10.0 s, *Earthq. Spectra* **24,** 99–138.

Campbell, K. W., and Y. Bozorgnia (2008). NGA ground motion model for the geometric mean horizontal component of PGA, PGV, PGD

and 5% damped linear elastic response spectra for periods ranging from 0.01 to 10s, *Earthq. Spectra* **24,** no. 1, 139–171.
Chiou, B. S.-J., and R. R. Youngs (2008). An NGA model for the average horizontal component of peak ground motion and response spectra, *Earthq. Spectra* **24,** no. 1, 173–215.
Dreger, D. S., G. C. Beroza, S. M. Day, C. A. Goulet, T. H. Jordan, P. A. Spudich, and J. P. Stewart (2014). Validation of the SCEC Broadband Platform V14.3 simulation methods using pseudospectral acceleration, *Seismol. Res. Lett.* **86,** no. 1, doi: 10.1785/0220140118.
Graves, R. W., and A. Pitarka (2010). Broadband ground-motion simulation using a hybrid approach, *Bull. Seismol. Soc. Am.* **100,** no. 5A, 2095–2123, doi: 10.1785/0120100057.
Maechling, P., E. Deelman, and Y. Cui (2009). Implementing software acceptance tests as scientific workflows, in H. R. Arabnia (Editor), *Proc. of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2009)*, 2 Volumes, Las Vegas, Nevada, 13–17 July 2009, CSREA Press, 317–323, ISBN: 1-60132-123-6.
Mai, P. M., W. Imperatori, and K. B. Olsen (2010). Hybrid broadband ground motion simulations: Combining long-period deterministic synthetics with high frequency multiple *S*-to-*S* back-scattering, *Bull. Seismol. Soc. Am.* **100,** no. 5A, 2124–2142, doi: 10.1785/0120080194.
Motazedian, D., and G. M. Atkinson (2005). Stochastic finite-fault modeling based on a dynamic corner frequency, *Bull. Seismol. Soc. Am.* **95,** 995–1010.
Jordan, T. H., and P. Maechling (2003). The SCEC community modeling environment: An information infrastructure for system-level earthquake science, *Seismol. Res. Lett.* **74,** 324–328.
Schmedes, J., R. J. Archuleta, and D. Lavallée (2010). Correlation of earthquake source parameters inferred from dynamic rupture simulations, *J. Geophys. Res.* **115,** no. B03304, doi: 10.1029/2009JB006689.
Somerville, P. G., S. Callaghan, P. Maechling, R. W. Graves, N. Collins, K. B. Olsen, W. Imperatori, M. Jones, R. Archuleta, J. Schmedes, and T. H. Jordan (2011). The SCEC broadband ground motion simulation platform, *Seismol. Res. Lett.* **82,** no. 2, 275, doi: 10.1785/gssrl.82.2.273.
Zeng, Y., J. G. Anderson, and G. Yu (1994). A composite source model for computing realistic synthetic strong ground motions, *Geophys. Res. Lett.* **21,** no. 8, 725–728, doi: 10.1029/94GL00367.

Philip J. Maechling
Fabio Silva
Scott Callaghan
Thomas H. Jordan
Southern California Earthquake Center
University of Southern California
Bldg. ZHS Room 156
3651 Trousdale Parkway, Suite 169
Los Angeles, California 90089-0742 U.S.A.
maechlin@usc.edu

Published Online 17 December 2014