リナックス計算状態スクラップブックシステム

1．背景

Saving and restoring intermediate computation state is a useful technique. A well know example is the hibernation feature on notebook computers, which saves all the intermediate state of a user's application programs to hard disk so that the power can be turned off without losing work in progress. Another example is that some applications quickly startup by restoring previously saved computation state. Users of scientific computing save the intermediate state of long running computations so that they can be resumed on different hardware without losing already completed parts of the computation. Users of databases and fault tolerant systems save checkpoints that make graceful recovery from hardware are software faults possible. Even though all these examples use basically the same idea, each requires separate code to save and restore the computation state.

2．目的

The purpose of this project is to build a general system for saving and restoring the intermediate computation state of the Linux operating system. Because so many software applications and programming systems have been ported to Linux, this one system has potential to support many different types of uses. The overall system is called SBUML, which stands for Scrapbook for User-Mode Linux, because it builds on top of the established User-Mode Linux project, and because it provides a Computation Scrapbook that can save multiple Linux computation states persistently for various creative and useful applications.

One of the main motivating uses for SBUML is to support programming tools. Software visualization and debugging tools are much harder to use than they could be because it is difficult to set up the state to be observed or debugged. Running a program can destroy state in which users are interested, and users must then use ad hoc and awkward methods to regenerate the program state. SBUML can help solve this problem and make existing Linux programming tools more powerful and easier to use. It also provides an infrastructure for building new advanced model checking and debugging tools.

An important part of this project is to establish SBUML as an open source project so that others can contribute to testing and debugging it. Use by others will help prove that it is a reliable and general system. This will make it more attractive to other researchers and developers who wish to experiment with new applications for systems that save and restore intermediate computation state.

3．開発の内容

SBUML currently runs well on installations of RedHat Linux versions 8.0 and 9.0 and can save and restore snapshots of the complete User Mode Linux runtime state, including file systems and

all processes, both kernel and user state. The snapshots can then be archived or migrated to different host hardware and then restored.

All of SBUML's functionality is accessible from shell commands, and these commands also make up its Application Programming Interface (API). Below the main commands are listed in **bold** text with parameters in *italic* text.

- "**sbumlboot** *machine-name base-name*" creates a new virtual machine with the name *machine-name* by booting from a standard base hard disk image named *base-name*.
- "**sbumlhalt** *machine-name*" halts the machine and deletes its resources. The optional parameter **–all** deletes all machines currently running.
- "**sbumlvnc** *machine-name ip-suffix*" changes the IP address of the machine to class-C-subnet.*ipsuffix*, and then attaches vncviewer to the machine, which gives the machine a graphical user interface.
- "**sbumlfreeze** *machine-name*" freezes the machine so its state does not change.
- "**sbumlcontinue** *machine-name*" unfreezes the machine so that it continues execution.
- "**sbumlsave** *machine-name snapshot-name*" will save an image of the machine named *machine-name* by creating a snapshot named *snapshot-name.* The optional parameter **–f** will freeze the machine before saving. The optional parameter **–c** will continue the machine after saving.
- "**sbumlrestore** *machine-name snapshot-name*" will restore a snapshot named *snapshot-name* by creating a machine named *machine-name.* The optional parameter **–c** will continue the machine after restoring. The optional parameter **–v** *ip-suffix* will run the **sbumlvnc** command automatically after restoring and continuing.
- "**sbumlmakedelta** *snapshot-name*" compresses the snapshot compared with the snapshots ancestor, so that only the changes since the ancestor are saved.
- "**sbumlexpanddelta** *snapshot-name*" expands a delta-compressed snapshot back to its full information content.

The above commands form SBUML's core functionality and API. They make it easy to save Linux computation states as ordinary file directories. These snapshot directories can be copied, transferred, organized and deleted with ordinary Linux commands. The delta compression commands can compress snapshots to be only a few megabytes in size, and sometimes less than a megabyte. This makes it possible to archive large numbers of snapshots in today's typical multi-gigabyte disks.

Another benefit of delta compression is that it can make snapshots small enough to download quickly over the Internet. To explore this benefit, SBUML includes the following Web Page Snapshot functionality:

- The **sbumlrestore** and **sbumlmakedelta** commands were modified to take an extra **–sd** *location* parameter, which tells these commands to automatically search in *location* for the snapshot. The location can be specified as a web address such as http://sbuml.source.net/snapshots/, in which case the matching snapshots will be downloaded to the local machine before processing.
- **"sbuml--mozilla-helper-application"** command is called by a web browser when a

SBUML Web Page Snapshot has been selected by the user. It simply transfers the information from the web page link to the **sbumlrestore** command so that it can download and restore the snapshot.

- **"sbuml--prepare-snapshots-for-posting** *url"* command is used to prepare a directory of snapshots for posting to a website that has the address *url*. After preparation, the directory can be copied to the web site and web pages can link to special files in the directory that look like this: link-*snapshot-name.*sbuml.

This functionality makes it possible to distribute complete Linux computation states on web pages. All users have to do is click on a link, and the User Mode Linux machine is automatically downloaded and restored on their local machines.

A website was developed at http://sbuml.sourceforge.net to distribute information about SBUML and its potential applications. In particular, the website has demonstrations of several simple Web Page Snapshot uses:

- **Software demonstrations**: Researchers and software developers can post fully configured demonstrations of their software on website. The software can be written in any language that has been ported to Linux, and users do not have to be burdened with software installation or program initialization.
- **Programming Language documentation**: Manual pages about programming language constructs can contain links to fully configured Linux machines setup to demonstrate the particular language construct. The machine can also include debuggers that allow users to explore in detail how the construct can be used in a concrete example.
- **End-user application documentation**: Documentation for the features of end-user application such as word processors and spreadsheets can contain links to snapshots of application configurations that make it easy for users to practice using specific application features.

4．従来の技術（または機能）との相違

Four distinguishing characteristics of SBUML are that it (1) saves state at the operating system level, (2) is designed to save multiple persistent copies of state, (3) saves thread state, and (4) is an open source project. VMWare is a commercial product that can save thread persistent state at the operating system level, however it is not open source and therefore not as flexible for exploring new creative uses. SBUML is different from most other systems that save and restore intermediate computation state because these systems designed to save one temporary snapshot. In contrast, SBUML is designed to save multiple persistent snapshots in a Computation Scrapbook.

5．期待される効果

SBUML is already complete enough for practical use. It can be used to distribute pre-configured already-running software inside a User-Mode Linux virtual machine. Many different types of users can take advantage of this. Researchers can use it to distribute demonstrations of their research prototypes. Teachers can use SBUML to prepare lectures that involve software

simulations. Students can use SBUML to download and review the simulations without worrying about installation and setup. Commercial software publishers can use it to set up and distribute demonstrations that show specific features of their product.

SBUML is also expected to stimulate research on systems that save and restore computation state. It can be difficult to add such functionality to a system because when a program is running some of its state is maintained by the operating system. SBUML solves this problem by capturing the state of the operating system too. Now researchers can save and restore the state of their research prototypes by simply porting them to Linux and running them inside of SBUML.

6．普及（または活用）の見通し

Since User-Mode Linux is one of the most active project of the SourceForge.net open source website, it is expected that SBUML too will be widely used. It is already part of active research projects at the University of Tokyo and the University of Tsukuba, and there have been inquires from universities in Europe and America about using it.

7. 開発者名（所属、ｅ－ｍａｉｌアドレス）

Richard Potter
potter@is.s.u-tokyo.ac.jp