


# DPM® SP



addendum



**PEAVEY®**

## INTRODUCTION

Hi and welcome to the DPM SP Plus addendum. This addendum is meant to provide you—the user—with any late-breaking developments. In this case that means features that have been added to the DPM SP Plus. In addition to all the really cool features described in the manual, the SP Plus now has the following features:

- Effects: Reverb, DDL (Digital Delay Line), Chorus and EQ
- Akai S1000 Compatibility

These features allow you to use the SP Plus in ways you never imagined possible.

## EFFECTS SETUPS AND PRESET SELECTION

The effects setup feature is located under the **Preset** button. To reach the effects setup feature press the **Preset** button, then press the **+ / INC** button until you reach the following screen:

```

Edit FX setup: P# 1
      <press exec>

```

Press the **EXEC** button to access the effects setup screens.

## ABOUT THE EFFECTS SETUP

An effects setup stores all settings for all effects. Effects setups in the SP Plus are stored as part of each preset. In “normal” operation—that is, in Omni or Poly mode—changing the active preset, whether done manually via the front panel or using MIDI program change messages, will cause the new preset effects setup to be loaded automatically, making it the active or “current” setup. Preset setups are edited via a set of screens which appears as a submenu under the Preset button. When presets are copied, the effects setups they contain are copied along with them. Effects setups are stored to and recalled from disk along with the presets they are part of and become available immediately upon being loaded into the SP Plus—no special action on the part of the user is required.

## EFFECTS OUTPUT ASSIGNMENT

Effects are applied to either the Main or Sub outputs. All sounds playing through the selected pair of outputs receive identical effects processing. At present there is no capability for specifying effects send levels independently for different sounds. However, sounds may be individually selected to receive effects processing by assigning only those sounds to the “wet” output pair using map or preset output routing controls.

A screen under the **Global** button allows the user to select which pair of outputs will have effects processing applied to it. To reach the effects output assignment press the **Global** button, then press the **+ / INC** button until you reach the following screen:

```

FX are applied to:
      main outputs

```

The chosen setting is retained when power is turned off and will remain unchanged until the user changes it again. Note that this setting is not part of any effects setup, is not stored to disk, and is not affected by preset changes. (At present, changes to this setting do not actually take effect until the next time the SP Plus is turned on. This may be changed in a future update).

## **EFFECTS SETUPS AND PRESET SELECTION IN MULTI MODE**

Multi mode works just a bit differently with respect to effects setup selection. In Multi mode, there are as many as sixteen active presets, each of which may have a unique effects setup. However, the SP Plus is capable of accommodating only a single effects setup at any one time. Therefore, the Multi setup includes an “effects MIDI channel” setting. When in Multi mode, the effects setup will be obtained from the preset which is active on the designated effects MIDI channel. Effects processing will be the same for all sounds regardless of the MIDI channel they are played on (assuming they are assigned to the “wet” output pair). Program changes sent to the effects MIDI channel will cause the active preset on that channel to change in the usual way, and this will also result in effects setup changes. Note that if this MIDI channel is disabled in the Multi setup, the effects setup of the preset selected for this channel will still be used, but program change messages sent on this channel will be ignored, and the effects setup will therefore not change in response to them.

A side effect of Multi mode as it affects the editing of effects setups is worth mentioning, since it may lead to some confusion if it is not understood. Generally, in all preset editing screens, including those in the effects edit submenu, the **+/INC** and **-/DEC** buttons can be used to select the preset to be edited. However, while in Multi mode, the effects editing screens will only display the preset which is designated by the Multi setup for the effects channel—the **+/INC** and **-/DEC** buttons cannot then be used in these screens to change the display to a different preset. This is done both as a convenience and to prevent inadvertent editing of an effects setup which is not currently the active one. It does not affect MIDI playback in any way—if a MIDI program change is received on the effects channel, the effects edit screens WILL switch to the new preset. If editing of other preset effects setups is desired, this can be defeated by simply switching the SP Plus out of Multi mode, or by selecting a different preset for the Multi setup's effects MIDI channel.

## **EFFECTS SIGNAL FLOW AND MIX LEVELS**

Four different effects are provided: a five-band graphic equalizer, a digital delay line (DDL), an eight-voice chorus and a reverb. They are all available simultaneously. Although only the EQ has a true stereo input—it actually consists of two separate EQ sections—each of the effects produces a stereo output which is maintained as the audio moves down the effects chain. The effects chain configuration is fixed—sounds pass in series through the EQ first, then the DDL, then the chorus and finally the reverb. However, by means of the mix controls provided on the output of each effect, the four effects may be used individually or in any combination.

Each effect module is provided with both wet and dry output level mix controls. The wet mix level control adjusts the amount of processed or effect signal which is fed to the module output. This signal contains only the result of the effect processing, and no component of the original input signal. The dry mix level simply controls the amount of the unprocessed input signal which is passed through to the output of the module. It is the dry mix level which allows the stereo outputs of preceding effects modules to pass through each additional module with stereo separation maintained. The adjusted dry and wet outputs are mixed together to produce the signal which is fed to the input of the next effect module in the chain (or in the case of the reverb, to the designated audio outputs). Thus, the input to each effect module consists of a mix of dry and wet signals from all preceding modules. There is no separate control for the level of signal input to each effect, but this can be simulated by adjusting both the wet and dry mix level settings of the preceding effect up or down together as needed.

Here are some practical examples of typical mix level settings: If all of the wet mix levels are set to zero and all of the dry mix levels to 99 (the maximum mix level), the result is a full bypass of all effects—only the unprocessed signal gets passed from one module to the next. (However, there are easier ways to obtain effects bypass—see below). When using the EQ, it will usually be preferable to turn the EQ wet mix level to 99 and the dry mix level

to zero, so that only the EQ'd signal is passed through. A typical DDL setup—for slap echo, for example—might have both wet and dry mix levels set to the same value. For just a touch of reverb, leave the reverb dry mix level at 99 and set the wet mix level to 10-15. For reverb only, without other effects, set the EQ, DDL and chorus wet mix levels to zero. If the chorus seems to be distorting, then turning down both the wet and dry mix levels of the DDL, which feeds into the chorus, may stop the overloading.

## EFFECTS BYPASS OPTIONS

```
Preset FX ctrl P# 0
effects are:active
```

Effects bypass may be achieved in two different ways:

A global effects bypass switch is accessible under the **Global** button. It offers an easy means of “getting all of the effects out of the way” in order to hear unprocessed sounds. When global effects bypass is turned on, effects processing is disabled for all presets and under all conditions. This switch setting is retained when power is turned off and will remain unchanged until the user changes it again, so it is especially useful in cases where no on-board effects processing is ever desired. It is not part of any effects setup, is not stored to disk, and is not affected by preset changes, nor does it alter preset effects setups in any way. Note that editing of effects setups is blocked when global bypass is active—effects edit screens are locked out, with a message explaining the reason why. This is done as a safety measure, to prevent effects setups from being edited at a time when the results of the edits cannot be heard.

Effects bypass can also be done on a preset-by-preset basis. The effects setup within each preset contains a bypass switch whose setting takes effect only when that setup becomes the current setup—that is, when the preset is selected. This switch can be used to temporarily or selectively disable effects processing for individual presets while leaving the other effects settings in those presets intact, and without affecting other presets. Unlike the global effects bypass switch, this switch is a part of the effects setup for each preset and does get stored to and recalled from disk along with the preset. Since it is a part of the effects setup, it is accessed via a screen in the effects edit submenu under the Preset button. Be careful here—preset-level effects bypass does not block editing of other settings in an effects setup. If the effects edit controls seem to be having no effect, check to make sure that the effects bypass switch is not on in the current preset.

Certain combinations of playback and MIDI modes can also result in temporary effects bypass. This is discussed in the next section.

## HOW PLAYBACK MODE SETTINGS AFFECT EFFECTS OPERATION

Because an effects setup is part of a preset, the selection of MIDI mode and playback mode have some effect on the way the effects system behaves.

Operation is simplest when the playback mode is set to “per MIDI rev mode.” Under this setting, the selected MIDI mode (Omni, Poly or Multi) is always active. The discussion up to this point has assumed this setting for playback mode.

There are four playback mode settings which cause playback to directly access the currently selected preset, map, tone or wave—that is, the one which appears in the display when a corresponding editing screen is called up. These playback mode settings also override the MIDI mode selection and force the SP Plus into Omni mode. When playback under the current preset is selected, effects operation is as described above for Omni mode—it behaves “normally,”

with the effects setup in the selected preset becoming the active one. However, playback under the current map, tone or wave causes effects to be bypassed. This is partly for the simple reason that in these modes, which bypass presets entirely, there is no way for the SP Plus to determine which preset to use for the effects setup—no preset is currently active. It also serves to minimize sonic confusion when attempting to edit a map, tone or wave, while saving the user the trouble of switching the effects into bypass manually. Under these conditions, as when global effects bypass is active, and for the same reason, it is not possible to gain entry into the effects setup editing screens.

Finally, there is the “wildcard” playback mode setting “what’s on display.” When this setting is chosen, the actual playback mode will switch automatically among the five modes described above depending upon which screen is on the display. Thus, while under Map, Tone or Wave buttons, playback switches to the displayed map, tone or wave and effects are automatically bypassed. Under the Preset button, the displayed preset and its effects setup are active, but under forced Omni mode. Under Disk, Global or Utilities buttons, operation reverts to “per MIDI rev mode”—that is, “normal” operation, with Omni, Poly or Multi mode active as selected. If the effects seem to cut in and out of their own accord as different screens are called up, it’s probably because the playback mode is set to “what’s on display.”

## **THE DEFAULT EFFECTS SETUP**

When a new preset is created on the SP Plus (using the “Create preset” function), it includes a default effects setup. This effects setup creates a moderate reverb patch in which only a small amount of reverb is heard. It contains usable settings for each of the other effects as well, but the effects mix levels are set to leave the DDL and chorus out completely. These can be added in by simply turning up their wet output mix levels, and their other settings can then be adjusted as desired. The EQ mix levels are set for full equalized output and none of the pre-EQ audio, but the boost/cut controls in each EQ band are all set to zero. Thus, the EQ initially does nothing to the sound passing through it, but it is already part of the setup and will immediately affect the sound if its boost/cut controls are adjusted.

## **EXISTING SP DISKS/PRESETS**

The SP Plus is fully compatible with existing SP presets and sound disks. Presets created for the SP prior to the release of the SP Plus do not contain effects setups. Such presets will be automatically and transparently updated when detected by the SP Plus. This detection occurs **ONLY** when a preset gets selected as the current preset from which to obtain the effects setup. If the SP Plus tries to read the effects setup information from the preset and finds that the preset does not contain an effects setup, it instead implants the default effects setup into the preset, and then uses this as the current setup. The new effects setup can then be edited as desired. If the contents of SP Plus memory are now saved to disk, this newly-created effects setup is stored as part of the updated preset and will be recalled when the disk is again loaded into the SP Plus. Note that SP presets which are imported into the SP Plus via MIDI SysEx messages are also updated automatically upon being selected.

As a result of the automatic updating feature and the settings in the default effects setup, which is a reverb patch, all SP presets which are played from the SP Plus will come through a reverb patch by default. If this is not wanted, a simple method of getting around it is to activate the global effects bypass option, described earlier. However, this will suppress effects processing on all presets, including SP Plus presets for which effects processing may still be wanted. Another method is to use the preset effects bypass switch which is a part of each effects setup. The effects

setup copy utility, which is described in the next section, makes it easy to apply this method to multiple presets. After loading a bank of SP presets from disk, the preset effects bypass switch can be turned on in just one of the presets—or any other desired effects setup can be created in that preset—and then this effects setup can be copied to all of the other presets in the bank using this copy utility, updating them in the process. It is important to remember that the automatic updating feature does not operate on all presets as soon as they are loaded from disk, but **ONLY** on those presets which actually get selected for playback. The effects setup copy utility is therefore the most effective way to ensure that all presets in a bank get updated and that the default effects setup will not crop up where it is not wanted.

SP Plus disks which include presets containing effects setups can still be loaded into the SP without problems. The SP can make no use of the effects setup information and will simply ignore it. Also, although this is not guaranteed indefinitely to be so, the SP will generally retain intact the effects setup in an SP Plus preset even though it cannot use it, and if this preset is saved to disk again by the SP, it will usually still contain its effects setup. (It would be wise not to depend upon this, however).

### THE EFFECTS SETUP COPY UTILITY

```
Copy FX setup to:
P# 0 Mega Kit
```

In order to allow effects setups to be copied from one preset to another, an effects setup copy utility has been provided. This utility appears in the effects editing submenu under the **Preset** button and is very simple to use. It simply copies the current effects setup to the indicated preset when the **Exec** button is pressed. When importing SP sound banks into the SP Plus, it will often be useful to duplicate a new effects setup in all currently loaded presets. The selection of “all presets” as a destination for the effects setup copy appears at the end of the list of individual presets and allows this duplication to be done in a single operation. Note that the copy utility has the effect of updating SP presets which did not previously contain effects setups and can be used to prevent unwanted automatic updating from occurring.

### THE INDIVIDUAL EFFECTS

The types of effects provided by the SP Plus are very familiar to most musicians. This section focuses on one effect at a time and describes the aspects of each which are unique in the SP Plus, or which may not be obvious from the editing screens.

### THE FIVE-BAND GRAPHIC EQUALIZER

```
EQ boost/cut: P# 0
freq: 100Hz 0dB
```

Actually there isn't very much unique about this effect. The single set of boost/cut controls affects both the left and right EQ sections together instead of having a separate set of boost/cut controls for each section, as is common in outboard analog equalizers. Another departure from the outboard version of this effect is that the output mix controls of the SP Plus equalizer make it possible to send any desired amount of the non-equalized input signal through along with the equalized signal. The frequency bands for the five band graphic EQ are located at 100 Hz, 330 Hz, 1000 Hz, 3300 Hz and 10 kHz. The maximum boost setting is +12 dB and the maximum cut is -24 dB.

## THE DIGITAL DELAY LINE (DDL)

```
DDL levels:      P#  0
L:75 C:99 R:75 Fb:00
```

The DDL provides three delay outputs or taps which are fixed-panned hard left, center and hard right. Each tap can be individually programmed for any delay between zero and 278 msec. In addition to the output taps there is also a feedback tap which does not produce directly audible output but whose output is fed back to the input of the DDL to produce repeating echoes. All three output taps and the feedback tap have individual level adjustments as well. In the case of the feedback tap, the level control affects the number of times a sound will repeat before its echoes die away.

```
DDL times (ms) P#  0
L180 C090 R270 Fb270
```

The design of the SP Plus DDL requires that the delay time setting of the feedback tap always be as large as or larger than that of any of the output taps—that is, the feedback tap must come last in line. The DDL delay time editing screen takes care of this by automatically adjusting the feedback tap delay time when necessary to accommodate changes to one of the output taps, in addition to preventing the feedback tap delay from being set to a time which is smaller than one or more of the output tap delay times.

## THE EIGHT-VOICE CHORUS

The simple editing screens of the SP Plus chorus mask some rather complex action. As with driving an automobile, you don't really need to understand it in order to use it, but if you want some knowledge of its inner workings—read on!

The chorus consists of a small DDL with four pairs of output taps—thus, eight voices. The delay time values of these output taps change continuously under the control of delay time sweep generators. Each output tap pair has its own sweep generator, so that the sonic “motion” created by the sweep can become very complex when all four pairs are swept at different rates—which is usually the case. Within each output tap pair, one of the taps is panned hard left and the other hard right, resulting in interesting stereo effects when the two taps are operated at different delay times—which is usually the case. Additionally, although the two taps share a sweep generator in common, they are swept out of phase with one another, which dramatically expands the stereo effect. The chorus editing screens have been designed to make this complexity manageable.

```
Chorus delay:  P#  0
base:10.0  spread:50
```

Delay times for all eight taps are set by just two controls. The base delay setting sets the “starting” position for all taps. The delay spread control does exactly what it sounds like it does. When set to zero, all taps have the same delay time. At higher settings, it spreads the tap delay times out from one another in both directions to increasing degrees, in a left-right alternation, so that the effect produced by each tap is slightly different from that of every other tap. This results in a richer, more diffuse chorusing effect.

Note that the effect of the delay spread control varies somewhat depending upon the base delay value. Base delay settings which are close to either end of the available delay range leave less space for the delay spread control to work with. It still works, but the delay taps are not spread as far apart for a given spread setting as they would be if the base delay were set closer to the middle of the range. This is not necessarily a bad thing or something to be avoided, but simply something which it may be useful to be aware of. The spread “squeeze” is most significant when base delay is set into the range (roughly 5 msec and under) in which the effect becomes flanging rather than chorusing.

```
Chorus speed: P# 0
base:30 spread:50
```

Sweep speed controls are set up similarly. The base sweep speed setting is the main speed control for all four sweep generators. When the sweep speed spread control is set to zero, the sweep speed is the same for all four output tap pairs, while raising the sweep speed spread setting causes the sweep generators to run at increasingly different speeds from one another. This also tends to create a more diffuse sound by ensuring that the sweep generators do not run in lockstep.

```
Chorus ctrl: P# 0
depth:50 fdbk:-50
```

The depth setting controls the amount of sweep on all taps. The most important thing to know about this setting is that it is affected by the base delay setting and, to a lesser extent, by the delay spread setting. Base delay settings which are close to either end of the delay range leave less room for movement of the delay taps by the sweep generators, with the result that a given depth setting translates into less actual sweep than it would if the base delay were set closer to the middle of the range. This depth “squeeze” is even greater if a large delay spread setting is used.

```
Chorus levels: P# 0
tap1:+25 tap234:+25
```

Output tap levels are also controlled by two settings. One setting controls the output level of tap pair 1, while the other controls the output level of tap pairs 2-4 in tandem. The main reason for this division is that it provides a simple and flexible way of cutting tap pairs 2-4 out of the picture. This is sometimes useful as a way of “tightening up” the effect by deliberately decreasing its complexity, as for example when using the chorus with the very short delays (roughly 5 msec and under) in which the effect becomes flanging rather than chorusing.

Unlike the DDL, which has a dedicated feedback tap that is separate from the output taps, the signal used for feedback in the chorus is simply the sum of all of the output taps—after they have been affected by their respective output level controls. The chorus feedback setting controls the amount of this signal which is fed back to the chorus input. It should be noted that high settings on both the feedback and tap output level controls can result in overloading and even sudden oscillation of the chorus (turn the volume down when playing with these settings!). When all four output tap pairs are in use, their cumulative effect causes these thresholds to be reached at even lower settings.

The chorus output tap and feedback level controls can all be set to values in the  $\pm 99$  range. Negative settings produce a phase inversion in the resulting signal. One very nice way to use this is to set up the chorus in the flanging delay range and use negative feedback with positive tap output, or vice versa. This produces a flanging effect with a distinctive sound which is markedly different from that obtained when feedback is in phase with the output.



## THE REVERB

With the reverb you'll find that different types of reverb sound best with particular types of signals. For example, "smooth" reverb works well with percussive sounds, whereas "medium" reverb seems well-suited to piano sounds.

```
Reverb config: P# 0  
Type:Hall Size:Enorm
```

The reverb allows you to select from three reverb types, hall, plate and room, each with its own characteristics. The size parameter allows you to choose from small, medium, large, huge and smooth, each altering the apparent size of the reverberant space.

```
Reverb adjust: P# 0  
Time: 5.5 Damp:8kHz
```

The reverb time sets the length of the reverb "tail" (how long it takes for the reverb to fade away) from 0 to 8 seconds in .1 second increments and from 8 to 30 seconds in 1 second increments. Using the Damp parameter you can set the high-frequency absorption of the room by using a low-pass filter to simulate a more acoustically dead environment. Select from cutoff frequencies of 250 Hz, 500 Hz, 1 kHz, 2 kHz, 4 kHz, and 8 kHz. Of course, if no damping is desired you can turn this parameter off.

## **USING THE SP/SP PLUS WITH THE AKAI® S1000 SOUND LIBRARY**

The Peavey DPM-SP and SP Plus now feature the ability to access the extensive sound libraries of the Akai S1000 sampler. This addendum to the operating manual describes these functions in detail as of firmware version 2.00.

### **OVERVIEW**

The S1000 disk access function operates on SCSI disks. It does not support reading of floppy disk volumes in the S1000 format. However, the supported SCSI devices include removable cartridge hard disks, CD-ROMs and almost any other SCSI disk storage device which would normally work with an S1000, in addition to plain old hard disks. The operation and capabilities of this function are currently identical on the SP and the SP Plus. For simplicity, therefore, the discussion of the function will use the name "SP" to refer to both the SP and the SP Plus from this point onward.

The S1000 disk access function allows the SP to access both samples and programs in an S1000 volume. The contents of the disk being accessed are presented to the user for selection in a hierarchical fashion that experienced S1000 users will feel right at home with. The conversion of not only sample files but also program parameters, including key mapping, is fairly thorough. Because of differences between the program architecture of the S1000 and that of the SP, the conversion of programs cannot always be perfect. However, most programs will play immediately as soon as they have been loaded into the SP, with little or no additional tweaking by the user. Even when some additional programming is required to get the imported programs to play as desired, the conversion routines will have already taken care of most of the really tedious work automatically.

One more thing—the S1000 disk access function is a read-only function. It does not have the ability to write data out to disk in an S1000 format. Volumes imported using this function are fully converted to SP files and can be saved to SCSI or floppy in the SP file format.

### **BASIC OPERATION**

The S1000 disk access function appears in the Disk menu. It presents itself in a manner akin to the program and sample loading functions of the S1000 itself. The SP lacks the large display of the S1000 and cannot put a list of several files on screen at one time, but it reads disk directories more efficiently than the S1000 and therefore the process of selecting files to load can actually be much quicker than on the S1000, especially when working with a large and relatively slow device such as a CD-ROM drive. A logical series of screens accessed by a small number of keystrokes provides for making selections at progressively lower levels in the file hierarchy until a specific item is selected and loaded. When reentering these screens, the SP remembers your most recent selection at each level.

```
LOAD Akai data:SCSI0
      <press exec>
```

In the first screen, the SCSI device ID of the drive containing the S1000 volumes is selected. Hitting the **Exec** button causes the SP to access the selected drive and move to the second screen, in which the volume list for a specific partition on that drive is presented.

```
LOAD Akai data:SCSI5
P:A vol:VOLUME 001
```

(Partitioning is a technique used by the S1000 to simplify the handling of large disk drives by dividing their storage space into chunks of 60 megabytes or less—the SP does not partition the hard disks it uses). The SP reads in a partition's entire volume list at one time, allowing the list to be scrolled through very rapidly on screen without pauses for additional disk access. The volume names as well as the single-letter designation of the selected partition appear just as they would on the S1000. However, unlike the S1000, the SP avoids the needless distraction of displaying inactive volumes. While in this screen, in addition to selecting a volume from the volume list, it is also possible to select a different partition. Doing so will cause the SP to access the same drive again, read in the volume list for the newly selected partition and return to the same screen. (The partition chain on the S1000 disk must be traced from the beginning each time a new partition is selected, so this access takes progressively longer as partitions higher in the chain are selected).

Once the desired volume has been found, hitting the **Exec** button causes the SP to read in the file directory for that volume and advance to the third screen.

```
LOAD Akai data:SCSI5
Program:ALL
```

The list of files which can be loaded by the SP is presented in this screen. The first choice is always "program:ALL." After this the individual programs in the volume appear by name. Following the last individual program file, the choice "sample:ALL" appears, in turn followed by individual sample file names. Loading all programs is the SP equivalent of loading an entire S1000 volume.

If either "program:ALL" or "sample:ALL" is selected, hitting the **Exec** button will cause the SP to completely clear its memory and begin the load operation. Each incoming program or sample is assigned by the SP to the lowest-numbered preset or wave location available at the time.

If an individual program or sample is selected, hitting the **Exec** button will call up a fourth screen in which it is possible to select the preset or wave location into which the program or sample will be loaded. Hitting **Exec** one more time will cause this load operation to begin. When loading an individual S1000 program or sample, the SP will *not* clear its memory first. This allows programs and samples from different S1000 volumes to be combined in the SP. In the case of programs, only unoccupied preset locations will be offered for loading into. However, a single sample can be loaded into any wave location, including one which already contains a wave—the existing wave will be deleted first.

Until loading has actually begun, the screen sequence can be escaped at any point by hitting any function button. Because of the way the SP captures directory information from an S1000 disk, once such an escape is taken it is necessary to begin again with the first screen. However, for convenience, all of the selections which were last made in each screen—SCSI ID number, disk partition, volume, specific file and load destination—are remembered by the SP and will be used as the default selections upon reentering the screen sequence. Three or four quick keystrokes will get you right back to where you were.

## ABOUT LOADING PROGRAMS

In the S1000 disk format, programs contain direct references or links to samples, without which the program is meaningless. Thus, whenever the SP loads an S1000 program, it will also load the samples which are used by that program. The display will show this clearly while it is happening. The program load procedure is good at detecting when some or all of the samples needed by a particular program have already been loaded into the SP, and will link the incoming program to these samples wherever possible instead of loading duplicate copies of the same

samples. Note that this identification is based on the names of the samples, which therefore should not be changed needlessly.

After the samples associated with each program have been loaded, the SP will scan through them looking for any stereo sample pairs which may have been loaded and need to be merged together. See the section on stereo samples below for details on why, how and when this is done.

When loading all programs from a volume, the SP simply repeats the one-program load procedure for each program it finds in the volume, and the display will show alternating program and sample loads. Since the programs in a volume must make use of the set of samples which is also in that volume, it will often happen that most or all of the samples in the volume get loaded along with the first program. Programs loading after the first one will get linked to the samples already loaded and may not cause any additional samples to be loaded along with them. Another point worth noting is that even in an all-programs load, it is possible that some of the samples in a volume will not get loaded. This would occur if there are samples in the volume which are actually not used by any of the programs in the volume.

For each program which is loaded, the SP creates a single preset and a unique set of maps and tones to tie the new preset to its waves. The preset is given the name of the program from which it is created. Anywhere from one to four maps is created, depending upon the contents of the S1000 program. These maps are given the same name as the preset, with "m1," "m2," etc., appended to the end of the name to allow the maps to be more easily distinguished from one another. Each S1000 sample loaded as a "wave" in the SP receives the original name of the sample. A unique SP tone is created for each wave used by the new preset. Each new tone is given the name of the wave it applies to, with the two-digit number of the new preset added to the end of the tone name. A tone created for an S1000 program is used only by the preset for which it is created—the SP does not attempt to "reuse" existing tones for later S1000 program loads. As a result, when several programs are loaded from the same S1000 volume, each wave may end up with several tones pointing to it—one for each program that was loaded. Appending the preset number to the tone name prevents these tones from all ending up with the same name, which would not be a problem for the SP but could be very confusing for the user if editing of the tones was later attempted.

## **ABOUT LOADING SAMPLES**

The S1000 disk access function can also be used to load samples only into the SP. This can be done one sample at a time, or all of the samples in a volume can be loaded in a single operation. When loading only samples from an S1000 volume, only waves are created in the SP—no tones, maps or presets are created.

As mentioned earlier, an all-samples load will cause all previous memory contents to be cleared—not just waves, but tones, maps and presets as well—while a one-sample load will simply add to existing memory contents or possibly replace one existing wave. The other main difference between the two modes of sample loading is that, like one-program and all-program loads, the all-samples load includes a scan for stereo sample pairs needing to be merged into stereo waves, whereas the one-sample load does not do this. The one-sample load, if directed at a sample which is actually part of a stereo pair, will load only that sample—it will not pull in the other member of the pair. This allows the option of loading only one member of a stereo pair when desired.

S1000 samples are stored on disk as individual files and are loaded one at a time by the SP. When loading an S1000 program or doing an all-samples load from an S1000 volume, the SP will not attempt to figure out in advance the total amount of wave memory which will be needed and whether all of the samples to be loaded will fit into the available space. Instead, as each sample comes up for loading, it simply checks to make sure there is room for that sample, and then loads it if there is room for it. Note that if a sample cannot be loaded because it is larger than the available space, the load operation will continue anyway, since there may be smaller samples waiting in line which

will fit—or, in the case of an all-programs load, additional programs which will can be loaded without using any additional sample memory. Thus the SP always loads as many of the desired samples (and programs) as it possibly can. If samples had to be skipped at any time during a load operation, a message to that effect will appear on the display when the operation is complete.

## **HANDLING OF STEREO SAMPLES**

The S1000 does not actually “know about” stereo sounds. Its programming architecture provides the means whereby two independent samples can be set up to play together such that if they happen to be left and right sides of the same stereo sound, then the original stereo sound is recreated, but the same programming technique can be used with any two samples, related or not, and is not provided specifically for stereo playback. Stereo sample pairs in the S1000 customarily have names which are the same except for the last two characters, which contain “-L” or “-R” to identify them as a stereo pair, but this is not a rule imposed by the S1000, nor is it necessary for S1000 stereo playback, and the S1000 does not treat them any differently because of their names.

In the SP, on the other hand, waves are inherently either mono or stereo. A stereo SP wave is a single entity with both left and right sides embedded in it, which is stored to disk, appears in memory and in all other respects is handled as a single wave, EXCEPT when it comes to playback, at which time it is played through a phase-locked pair of oscillators, both of which are under common control of a single set of envelopes, etc. Therefore the S1000 disk access function includes provisions for detecting and linking samples which are halves of a stereo pair, and merging them into single SP stereo waves.

The merge scan occurs automatically following most S1000 file load operations—one-sample loads are the only exception. Whenever it runs, it scans the entire wave directory attempting to match pairs of waves together according to their names, looking for “samplename-L” and “samplename-R” as described above. Upon identifying such a pair, it also checks the sample lengths to make sure they are equal. If they match, the SP considers them a pair and merges them, using the same procedures that are used by the SP stereo merge function in the wave edit menu. The R sample is merged onto the L sample and disappears from view, while the “-L” at the end of the sample name is replaced with “-S.”

Depending upon how the samples were loaded, the merge may go very quickly or may take some time. Normally the SP will load stereo sample pairs in a simple L-R-L-R order, because that is the order in which it will naturally go looking for them. This is optimal for the merge procedure, which then can get its job done with mere “hand waving” and doesn’t need to shuffle the contents of wave memory around. If the L and R samples get loaded in some other order, the merge may take quite a while longer, because then the SP needs to arrange the sample pairs into the L-R-L-R order before they can be merged. The order of loading is not always under the control of the SP, but depends upon the S1000 programs. One case in which this out-of-order load can happen is with library volumes which use large stereo samples. These volumes often provide “shrink” mono versions of their programs which use only the L sample from each pair in order to allow the program to fit into less sample memory. If the “shrink” program appears first in the volume directory, the SP will load it first, pulling in all of the L samples with it in a group. If the full-size program is loaded next, the SP will then pull in all of the R samples in a group, since it will detect that the L samples have already been loaded, and the merge procedure will then have quite a bit of work to do to get the samples arranged as required for merging. This is harmless, but it may challenge your patience! If you catch the S1000 disk access function spending a lot of time doing its merge scans, you should suspect that something of this sort is going on. If you know that a volume contains more than one version of a program try to load the larger, stereo version of the program first if you can tell which one it is, or skip loading the “shrink” version of the same program if you don’t need it.

Another thing to know about the stereo merge scan is that it doesn't confine itself only to the samples loaded along with a particular program, but as mentioned above, works its way through the entire wave directory merging everything it can. This means that samples can get "spontaneously" merged even if they were previously loaded as single samples—the only necessary conditions are that both samples are present in SP memory, and that a merge scan occurs. An example of this is the same case used to illustrate the last paragraph. The "shrink" program is linked to all of the L waves that were loaded with it, but after the full-size version of the program is loaded, all of those L waves will get merged to their R counterparts, and the "shrink" program will end up being linked to the S waves which take the place of the L waves (and in fact the two programs will now be the same size in terms of wave memory usage). On the other hand, if the "shrink" program had been created using R samples instead of L samples, then it would end up not linked to any waves at all, because the R waves would be gone from the wave directory after the stereo merge joined them onto their L counterparts. This can be prevented when necessary by changing the names of the samples in question—for example, replacing the capital "L" or "R" at the end of the name with a lowercase "l" or "r."

Aside from the stereo merge scan itself, another important aspect of stereo sample handling is the matching of stereo sample pairs on an S1000 disk to already-merged stereo waves in SP memory. This becomes significant when multiple programs are loaded from the same volume. When the SP detects that a program being loaded contains a reference to a stereo sample pair, it seeks to match this sample pair not to a pair of waves in SP memory but to a single stereo wave created by merging the same pair of samples on a previous load operation. If such a wave is present in memory, then the incoming program is linked to it and the samples it refers to are not loaded redundantly. In order to make this connection, the SP uses the sample names exclusively, looking for the "-S" which takes the place of the original "-L" and "-R." Therefore, the new "-S" sample name should not be changed casually. As mentioned above, if a matching S sample cannot be found in memory, the SP will look for matching L or R samples, in case one of them has already been loaded without the other, and only the samples which cannot be found in memory will be loaded.

### **SOME NOTES ON PROGRAM/SAMPLE CONVERSION**

Apart from the merging of stereo sample pairs described above, translation of samples from S1000 format to SP format is fairly straightforward and complete. Samples are pretty much the same thing in all samplers, and the SP decodes and uses almost all of the information found in an S1000 sample file, including not just the sample itself and its name, but also its playback and loop points, its sample rate, and pitch and fine-tuning information as well. Only very esoteric users dealing with arcane sampling techniques are likely to ever notice a difference here—for example, the SP doesn't have a loop hold time parameter, and it supports only a single loop instead of the 8 independent loops provided on the S1000.

At the program level, however, there are big differences between the S1000 and the SP. Though in essence their goals are the same, the means whereby these two machines achieve these goals are distinct from each other in a number of important ways. In fact, programs can be created on either sampler which would be difficult or impossible for the other one to accurately emulate. In using the S1000 disk access function, however, this is rarely an issue. Most "real world" S1000 programs are well within the SP's own performance capabilities.

With this in mind, the approach which has been taken in the SP is to assume that the programs being converted have been created in a reasonably logical manner to yield typical, straightforward sampler setups of the most commonly used sorts. This is certainly true of most programs in professionally-produced sample libraries. The conversion function handles all of the essential and tedious data grinding, so that all of the necessary data components are imported from an S1000 program, and then the SP makes a selection from a very small set of choices as to how these components should be deployed in an SP preset. The resulting preset will almost always be immediately usable,

even if it differs from the original in some of its details. With all of the pieces ready at hand for manipulation by the user, and with no need to spend time dismantling complicated and unpredictable conversion results, any necessary programming will often amount to nothing more than changing a few of the global preset controls such as the preset mode, the layer switch control source, etc.

At the heart of the program conversion function is a procedure which scans the keygroups in an S1000 program and creates a zone in an SP map based on the information found in each S1000 keygroup. Actually, the S1000 keygroups are translated into as many as four SP maps, because each keygroup contains four velocity zones. Each SP map which gets built contains information from the same velocity zone level in all of the keygroups. These maps can then be layered, crossfaded or velocity- or otherwise switched in the SP preset. Not all of the velocity zone levels are actually used in every S1000 program, which is why the number of SP maps which gets created can vary. Special handling for stereo sample pairs is included here so that they get translated into an entry in a single SP map even though they occupy two different velocity zone levels in an S1000 keygroup. The SP preset is then configured according to the number of maps which were created. For one map, the preset plays layer A only, and the single map is selected for all map switch levels. If two maps were created, then the first one is assigned to all levels of layer A and the second to all levels of layer B, with a velocity-switch setup between the two layers. For three or four maps, a three-way or four-way velocity split between the maps is set up via layer A only.

As might be expected, the preset which results from this procedure will not always be exactly what is desired. An example which will often be encountered is that of a “chorus” program in which samples are layered against themselves with a slight amount of detuning between them to create a simulated chorus effect. As indicated above, the SP program conversion function will create two maps from such a program and will set them up in a velocity-switch preset. However, to put this preset into the desired state, only two settings need to be changed: the preset mode must be changed from “layer a/b switch” to “play both layers,” and the preset layer detune setting must be changed from 0 to a value such as +10. Or, if the original program was a velocity crossfade or simple layer program, simply changing the preset mode setting accordingly may do the trick.

We recognize that the current version of the S1000 program conversion function does have its limitations. Ongoing effort is being directed towards improvements to it, to extend the range of S1000 programs which can be faithfully converted by the SP with little or no assistance from the user. Periodic updates to the S1000 disk access function should be expected. Meanwhile, given the information above about how the program conversion function currently operates, the “diehards” among you may be able to take S1000 programs which the SP won’t scan satisfactorily back to the S1000 and modify them slightly so that the SP can convert them.

## **ADDITIONAL SPECIFICS**

A great deal of conversion work goes on “behind the scenes” in the S1000 disk access function. A full discussion of all of the details is beyond the scope of this manual, and to the extent that the conversion does its job smoothly, knowledge of the details is not necessary. This section focuses on a few additional differences between the S1000 and the SP which are most likely to be encountered in practice are mentioned, along with recommendations for dealing with some of them. The information here applies to SP / SP Plus firmware version 2.00. As noted above, some of the items discussed here may change in future firmware revisions.

**MIDI program numbers:** Each S1000 program includes a MIDI program number. The S1000 allows any number of programs to be assigned to a single MIDI program number, and when that program number is selected, the S1000 will play all of the programs assigned to that number simultaneously. The SP does not have this MIDI program layering capability—it can play only one preset at a time on any MIDI channel—and so does not make use of the S1000 program numbers. If only two programs are being layered, the SP can be made to simulate the S1000 response by

combining maps from two different presets within a single preset, assigning one to layer A and the other to layer B in that preset, and setting the preset mode to “play both layers.”

Certain settings in the S1000 allow overlapping ranges. This includes keygroup note ranges in a program as well as velocity zones within each keygroup. Furthermore the S1000 is capable of crossfading between the overlapping items within the overlap range. The SP does not support either of these features. Currently, it converts overlapping key ranges in an S1000 program to non-overlapping ones by truncating the overlap. In addition, the SP currently assumes that keygroups will appear in an S1000 program in order of increasing key range positions—i.e., moving up the keyboard. Keygroups which are “out of order” may appear misplaced in the resulting SP maps, or in extreme cases may not get translated at all. As mentioned above, editing such a program on the S1000 itself is one cure for such a problem. Simply splitting the keygroups among two or more separate S1000 programs to eliminate any key range overlaps will allow the SP to convert all of them accurately, at which point maps from the resulting presets can be layered together in a single SP preset (although the SP will still not crossfade the overlaps). As discussed earlier, velocity ranges within keygroups are currently not utilized directly by the SP conversion function, whose results are therefore not affected by the presence or absence of overlapping velocity ranges.

Unlike the S1000, which uses fairly conventional ADSR (Attack Decay Sustain Release) envelopes, the SP uses somewhat simpler envelopes which might be called ADR envelopes—the envelope does not have an actual sustain phase with an adjustable sustain level towards which the decay phase settles. The program conversion function makes use of the more compact parameter set of the SP envelopes to approximate the envelope settings found in an S1000 program. Envelope control in samplers is usually less critical than it would be in a non-sample based synthesizer because the samples often have a natural envelope built into them. Thus, even though the SP envelopes cannot do everything that an ADSR envelope is capable of, this is not usually a problem in converting S1000 programs, because the full capabilities of S1000 envelopes are not required in the first place. As a result, the envelopes created by the program conversion function are usually quite appropriate. The two envelopes of the SP are assignable, but by convention, ENV1 is assigned to the DCA, while ENV2 is assigned to the DCF—this closely parallels the fixed assignment of envelopes in the S1000, and envelope settings are translated accordingly from the two envelopes of the S1000 to those of the SP. Additionally, when the S1000 program applies some amount of envelope control to sample pitch, the resulting SP program will route ENV2 to the DCO in a corresponding proportion.

The S1000 provides two independent LFOs per program, applying to all voices played by that program. The SP handles LFO controls at the tone level—thus there may be many different LFO settings at play within a single preset—but each voice has only a single LFO. The SP will utilize the pitch LFO settings of an S1000 program to set up the LFO controls in each tone created for an SP preset, and will route the LFO to DCO pitch if the pitch LFO was active in the S1000 program.

With eight assignable mono outputs in addition to the main stereo output pair, the output structure of the S1000 is quite different from that of the SP. Currently the program conversion function does not attempt to make use of the output assignment settings in an S1000 program.