

Performance Evaluation

Integrated Systems
Advanced Technology Group

IOTracer 6.0 User Manual

Justin Bishop

April 6, 1995

IOTracer can be used to record information on file system, disk driver and resource manager requests by applications and the operating system in a Macintosh environment. It traces various I/O and I/O related operations including Read/Write, Open/Close, and Get/Load Resource Atraps. A trace post processor called "IOTracer Analyzer" takes the traces as input and reports information on traced operations. The format of IOTracer records is described in a separate document .

1. Introduction

IOTracer traces selected file manager and resource manager Atraps. A complete list of traps for which data is collected is shown in table 1. Using these traces, it is possible to gauge how a particular workload is using system, application and data files as well as how the workload utilizes the disk cache.

IOTracer is an INIT - CDEV combination. It will not work with systems earlier than System 7.0. IOTracer can be configured to dump its trace data array to disk while it is capturing data (dump and go mode) or to wait until data taking has terminated. When IOTracer is operating in dump and go mode and its data array reaches a set percentage of capacity, data collecting is turned off and the data array is written to disk. Data collection is then resumed. Because of this capability, system memory does not limit the maximum amount of trace data captured. The amount of free space on the disk holding IOTracer's output file (the current startup disk) gives the maximum output file size limit. When IOTracer is not working in dump and go mode, data will be collected until the data buffer is full and written to disk after data taking has been terminated. Figure 1 shows a picture of IOTracer's window.

2. Installation and Use

To install IOTracer, place it in the Control Panels folder and reboot the machine. (See the appendix for a step by step description of how to use IOTracer). IOTracer can then be launched from the Control Panels folder or from an alias.

Before starting data collection, the category of trap to trace must be selected. This is done using the two checkboxes labeled "File System Traps" and "Resource Traps" (see figure 1). One or both of the checkboxes must be selected before taking data. The specific traps belonging to the two categories (file system traps and resource traps) are detailed in section 5.

Two methods of starting and stopping data collection are provided. The second method does not require making the IOTracer's window the active window.

1. Click on IOTracer's "Start Taking Data" button to start data collection. Click on the "Stop Taking Data" button to stop data collection.
2. Data taking can be toggled on and off using an FKEY. An FKEY is a resource in the system resource fork (see section 3). The FKEY provided with IOTracer is activated by pressing the key combination Shift-Command-1. When the FKEY is used to trigger data taking, IOTracer's window will not display the "Taking Data" status. Instead, a SysBeep will occur to provide confirmation that data taking has started. When the FKEY is used only file system traps are recorded.




Figure 1. IOTracer

IOTracer's output file is named "IOTracer.out". If data taking is terminated using the "Stop Taking Data" button, the trace file will be placed in the active system folder. If the FKEY is used to stop data taking, the file will appear in the current directory at the time data taking is terminated.

Before the first data collection session, the trace buffer size can be changed from the default 2 megabyte buffer. The minimum buffer size is 1 Megabyte. To do this, click on the Preferences button and enter the size of the desired array in megabytes in the Buffer Size text field of the preferences dialog box (Figure 2). This box also allows control of the disk spooling mechanism by checking/unchecking the "Spool While Capturing" checkbox.

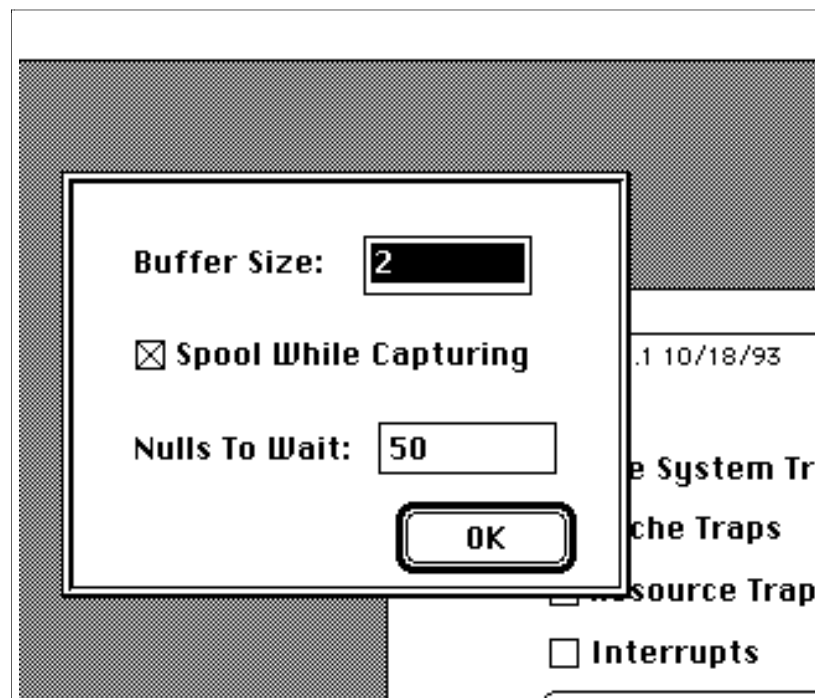


Figure 2. Preferences Dialog Box

When IOTracer is in spooling mode it uses null events sent to it by the Event Manager to check the number of records in its buffer. IOTracer is configured to check its buffer after a given number of null messages have been received. The default number is 50, but since different workloads may allow few null messages to be sent (i.e. do not return control to the system for long periods of time) and others may send many, this number can be changed using the preferences dialog box. Simply type the desired number of null messages to wait before checking the buffer into the Text Edit field labeled "Nulls To Wait". It is possible that IOTracer may not receive enough null messages in time to dump its buffer and so may lose data due to buffer overflow. If this occurs a flag will be set in the trace record header (detailed later in this document). After the first data taking session, the Preferences button will no longer appear on the CDEV.

If IOTracer is used with the spooling option, do not close IOTracer for an extended period of time while taking data. The INIT depends on the CDEV being up and receiving null messages for its file dumping mechanism. If the CDEV is closed, the INIT will not receive any "checkBuffer" messages and the trace buffer may overflow. This is not an issue for traces known to be smaller than the trace buffer, as the buffer will be emptied when the "Stop Taking Data" button is pressed.

3. FKEY Resource

An FKEY is a resource which must be included into the system resource file to add its functionality to the system. To install the FKEYs used with IOTracer, launch the provided application "Install IOTracer FKEY". Once the FKEY is installed, it is important to restart the machine before attempting to use them. This application will add one FKEY resource to the active system folder. The FKEY is used to toggle data collection and is triggered by the key combination Shift-Command-1. When the FKEY is used to control data collection the IOTracer window will not report the correct data collection mode and only file system traps are recorded.

<i>Trap</i>	<i>Category</i>
Read	File System
Write	File System
Open	File System
Close	File System
Create	File System
Delete	File System
Get File Info	File System
Create Resource File	File System + Resource
Open Resource File	File System + Resource
Close Resource File	File System + Resource
Load Resource	Resource
Get Resource	Resource
Get Named Resource	Resource
Get Individual Resource	Resource

Table 1. Traps and their Categories

4. Timers

IOTracer will search any NuBus slots for an MCP card. If one is found it will load the card's timer, setting it to give 25.6 microsecond resolution ticks, using this timer when collecting data. If no MCP card is found, IOTracer uses the Atrap \$A193 one microsecond resolution timer. IOTracer reports the type of timer used in the trace record header.

5. IOTracer Output Analysis

The analysis program *IOTracer Analysis* takes an IOTracer trace file as input and provides a simple analysis of the trace or a text output describing each trace record. The analysis package is described in a separate document.

6. Record Format

This section describes the IOTracer record format. This information is provided for those who wish to implement their own analysis programs. A detailed record format is given in the separate document *IOTracer Record Formats*. All records are 32 bytes.

Table 2 shows the header record format. The version field of traces taken by IOTracer 6.0 will contain 0x00000060. The timer can be one of two types, an MCP card timer or the Atrap timer. If the timer type field contains 0xFD, the time stamps are from an MCP card. If the timer type field contains 0xFB the Atrap timer has been used. The byte at offset 5 will be set to 1 if IOTracer was in dump and go mode (spooling). The four byte field at offset 12 will hold the timestamp read when data taking was started, the field at offset 16 will hold the timestamp at which data taking was turned off. The field at offset 24 from the header record contains the number of 32 byte records in this trace.

The remaining three fields only apply if the trace was taken using the dump and go option. The number of times the buffer was dumped to disk will be in offset 6. Offset 20 will contain 0x00000001 if a disk error occurred during the dumping of the data array. Offset 28 will contain the number of times the buffer overflowed (the number of times data was lost due to the lack of null messages). If this field is not 0, either increase the buffer size or decrease the number of null events to wait using the preferences dialog box the next time the trace is attempted.

<i>Offset</i>	<i>Size</i>	<i>Contents</i>
0	4	version
4	1	timer type
5	1	set to 1 if spooling
6	4	times buffer emptied
12	4	start timestamp
16	4	closing time stamp
20	4	set to 1 if disk error
24	4	number of records
28	4	buffer overflow count

Table 2. Format of Header Record

7. The Dump Record and the Go Record

If the "dump and go" option is used (spooling data to disk during capture), a 32 byte record will be inserted into the data array both before and after writing the data to the disk. Table 3 shows the format of these records. At the beginning of the spool operation a record with a four byte time stamp followed by 0x60 (Table 3) will be placed in the data array. Data capturing will then be turned off. after the spooling is completed a record with a time stamp followed by 0x61 will be placed in the array and data capturing will be turned on. These two time stamps can be used to determine how much time was spent spooling.

<i>Offset</i>	<i>Size</i>	<i>Contents</i>
0	4	time stamp
4	1	0x60 (starting) or 0x61 (finishing)

Table 3. Format of Dump and Go Records

8. Appendix A. Installing IOTracer

1. To install IOTracer, drag the IOTracer file into the Control Panels folder of the startup disk. The Control Panels folder is found in the System Folder.

2. If FKEY control is desired, double click on the application "Install IOTracer FKEY" to install the FKEY resources into the system folder of the startup disk. The installation program will install the resources and allow a choice of quitting and restarting or simply quitting the installation program (a restart is recommended). THE FKEYS MUST NOT BE USED BEFORE RESTARTING THE MACHINE (step 3). The FKEYs are not required to use IOTracer. IOTracer can also be started using the "Start" and "Stop" buttons.

3. Restart the machine.

4. Launch IOTracer from the Control Panels folder.

5. Select the category of traps to trace using the two checkboxes on IOTracer's window ("File System Traps" and "Resource Traps"). One or both may be selected. If the FKEY is used to start collecting data, only file system traps will be traced and checking the check box is not necessary.

6. Turn on data collection by either clicking on the "Start Taking Data" button or, if using the FKEY resource, hit the keys Shift-Command-1. When the FKEY is used IOTracer does not need to be the active window. The FKEY will also work when a modal dialog (such as one of the Standard File Package dialog boxes) is up. If the FKEY is used a SysBeep will occur to verify that data taking has started.

7. Execute the operation to be traced.

8. Turn off data taking by either clicking on the "Stop Taking Data" button or, if using the FKEY resource, hit the keys Shift-Command-1. The trace file will be named "IOTracer.out". If the button is used, it will be placed in the active system folder. If the FKEY is used, the file will be placed in the current directory at the time the FKEY is invoked.