CAGE Code 81205

Real-time CORBA Trade Study

Volume 2 – Basic IDL Scenario 1a

DOCUMENT NUMBER:	RELEASE/REVISION:	RELEASE/REVISION DATE:
D204-31159-2	**ORIG**	

CONTENT OWNER:

Phantom Works Advanced Information Systems (9-5430)

All future revisions to this document shall be approved by the content owner prior to release.

Document Information

Document Type Formal	**Original Release Date**	**Contract Number (if required)** F34601-96-C-0720 98-ESC-13
Preparing Organization (if different from owning organization)		**Hardware and Software Used** IBM PC Microsoft Word 97
Location of Software Files (optional)		
Boeing Web URL (optional) http://diicoe.web.boeing.com		
Notes and Limitations (optional)		

Signatures for original release

AUTHOR:	H. Rebecca Callison and Daniel G. Butler	9-5430	
	Sign and type: First Name MI Last Name	Org. Number	Date
APPROVAL:	Marilynn B. Goo	9-5430	
	Sign and type: First Name MI Last Name	Org. Number	Date
DOCUMENT RELEASE:			
		Org. Number	Date

Table of Contents

List of Figures

List of Tables

1. Introduction

This volume presents the detailed results for Basic IDL Scenario 1a, "Client and Server on Single SPARC", 70 ms frame time. Measurements were taken for both Call & Return (two-way) and One-way data transfers.

2. Call & Return (Two-way) Operations

Figure 1 summarizes the comparative performance of the three ORBs when the BasicIDL Call & Return methods execute with client, server, and background processes running in a single SPARC computer. Each of the lines in the graph captures the *average* operation time for messages of increasing size for transfers involving a particular data type. The Any series are labeled with average values in Figure 1. Average operation times for all other transfers are listed explicitly in the data table of Figure 2.

Scenario 1a: Client, Server on Single Solaris Host

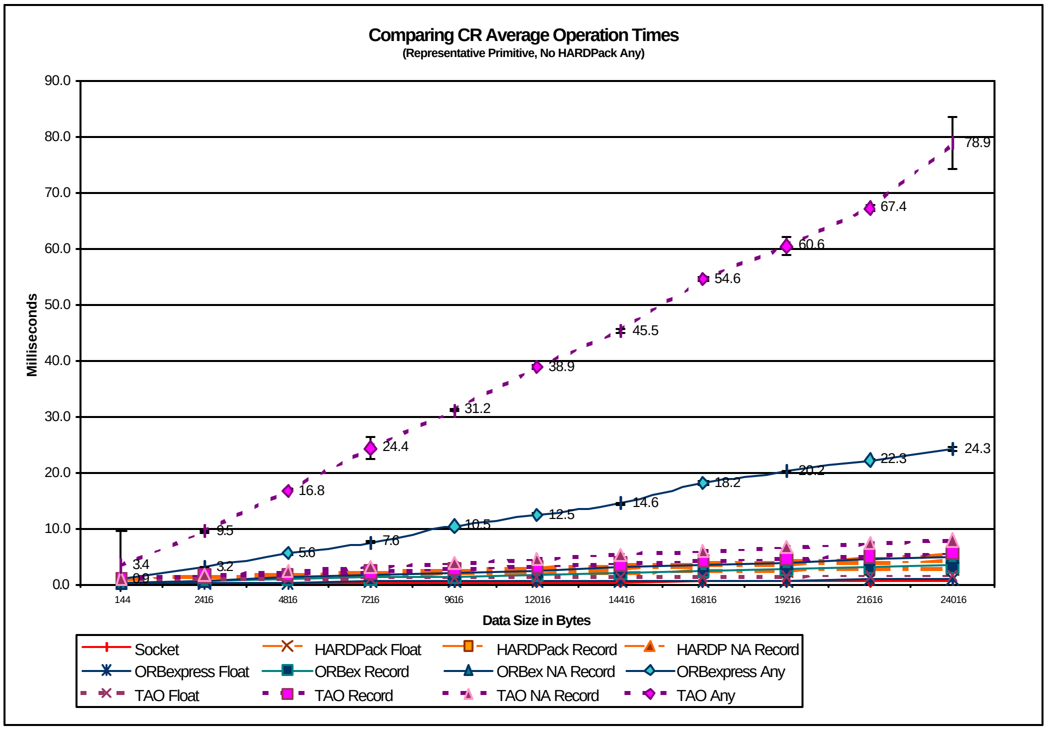

Figure 1. Call & Return Operations in a Single Solaris Host: Average

2.1 Summary: All Data Types

Performance information for five kinds of data transfers appears in Figure 1:

1. A basic TCP socket transfer of the designated data message size.
2. An ORB transfer of a C++ struct containing an array of float data. In general, we found that the performance of the ORBs when transferring arrays of primitive data was roughly comparable across the range of primitives, at least in comparison to some of the more complex data types like record, non-aligned record, or CORBA Any transfers. For these summary graphs, we arbitrarily selected the “Float” transfer for inclusion as *representative* of ORB behavior for primitive data types. We have tried to note, in the accompanying text of this report, any instances in which behavior across primitive data types was *not* roughly the same for a particular ORB.
3. An ORB transfer of a C++ struct containing an array of records in which the data items were neatly aligned on (32-bit) word boundaries.
4. An ORB transfer of a C++ struct containing an array of non-aligned records, records in which the data items were intentionally poorly aligned with respect to (32-bit) word boundaries.
5. An ORB transfer using the CORBA Any transfer method.

In most cases, data for all three ORBs appears on the graph for each ORB transfer type. The performance of HARDPack for the Any transfer method, however, was an order of magnitude slower than for the other ORBs. Including this data on a single graph skewed the presentation so drastically that it obscured differences in behavior in other areas. For this reason, HARDPack Any information is routinely omitted from the summary graphs.

Since all of the ORBs under evaluation use sockets to transfer data internally within the ORB, the socket performance represents a practical lower bound on the performance that can be achieved, helping us isolate the overhead added by the ORB. The socket performance we measured should not be construed as the best performance that can be achieved on basic sockets. We tuned our socket program just enough to get rid of obvious knees, peaks, and valleys for the program under test but did not explore the limits of socket performance. Our tuning may be typical of the level of effort that “real” programs might apply to the problem. It may even be slightly above average, since some programs may never consider the impact of socket tuning on system performance. But it certainly does not represent optimal socket performance, just representative.

Unless otherwise noted, any error bars in the graphs of this section depict the range of one standard deviation around the mean observed operation time. We use these bars to visually convey a minimal feeling for the temporal predictability of operations in the series. In Call & Return operations, however, larger standard deviations often arose from the cost of a single operation in the series, often the first. When this is the case, the standard deviation error bars exaggerate the amount of jitter that the ORB user can expect to observe over a routine series of operations.

The summary data in Figure 1 provides a few fairly obvious insights:

1. The CORBA Any transfer method is expensive and should be used with caution.
2. ORB*express* outperforms other ORBs on Any transfers by a significant margin. (We found this advantage to hold across all test scenarios.)
3. For other transfer methods, the ORB behaviors are fairly closely grouped, too closely for any conclusions to be drawn from this particular graph.

2.2 Records and Primitives

In Figure 2 we remove Any transfers from the graph, enabling a closer look at other transfer methods and data types. We see that in the single machine environment, ORB*express* has a significant advantage in terms of base overhead over the other ORBs. ORB*express* performance for the smallest message size in all data types converges on an average operation time of about .4 milliseconds, or about .15 milliseconds above the basic socket time of less than .25 millisecond. Both TAO and HARDPack begin at a lower limit of almost 1 millisecond, or four times the base overhead of the socket.

Scenario 1a: Client, Server on Single Solaris Host

Comparing CR Average Operation Times

(Representative Primitive, No Any)

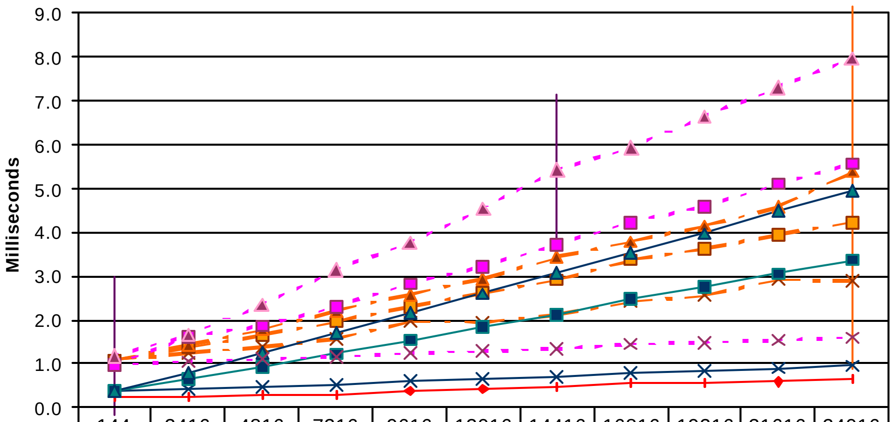

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
Socket	0.2295	0.2552	0.2868	0.3096	0.3787	0.4277	0.4528	0.5474	0.5643	0.5901	0.6415
HARDPack Float	1.0898	1.2432	1.3948	1.5625	1.9688	1.952	2.1142	2.4197	2.5606	2.9214	2.893
HARDPack Record	1.0869	1.3722	1.657	1.9569	2.3124	2.6102	2.9122	3.3726	3.6354	3.9394	4.2366
HARDP NA Record	1.0922	1.4252	1.7842	2.2135	2.5757	2.9488	3.4387	3.7943	4.1481	4.5795	5.3821
ORBexpress Float	0.3641	0.4166	0.4673	0.5216	0.6025	0.6544	0.7039	0.8091	0.8503	0.9028	0.9579
ORBex Record	0.374	0.6471	0.9417	1.2293	1.5489	1.8487	2.1292	2.4861	2.7627	3.0627	3.3696
ORBex NA Record	0.3835	0.7996	1.2439	1.7134	2.1619	2.6059	3.0913	3.5355	3.9776	4.4926	4.9569
TAO Float	0.9772	1.0533	1.1124	1.1659	1.2429	1.2915	1.3462	1.4526	1.4924	1.5458	1.6033
TAO Record	0.9869	1.6218	1.8693	2.3091	2.8332	3.23	3.7276	4.2357	4.5968	5.1049	5.5719
TAO NA Record	1.1762	1.6569	2.3505	3.1484	3.7713	4.5428	5.4355	5.9304	6.6506	7.31	7.9801

Figure 2. Call & Return Operations Without "Any"s: Average

Figure 3 presents the data a little differently so that trend lines can be calculated for operation time versus data size.

Scenario 1a: Client, Server on Single Solaris Host

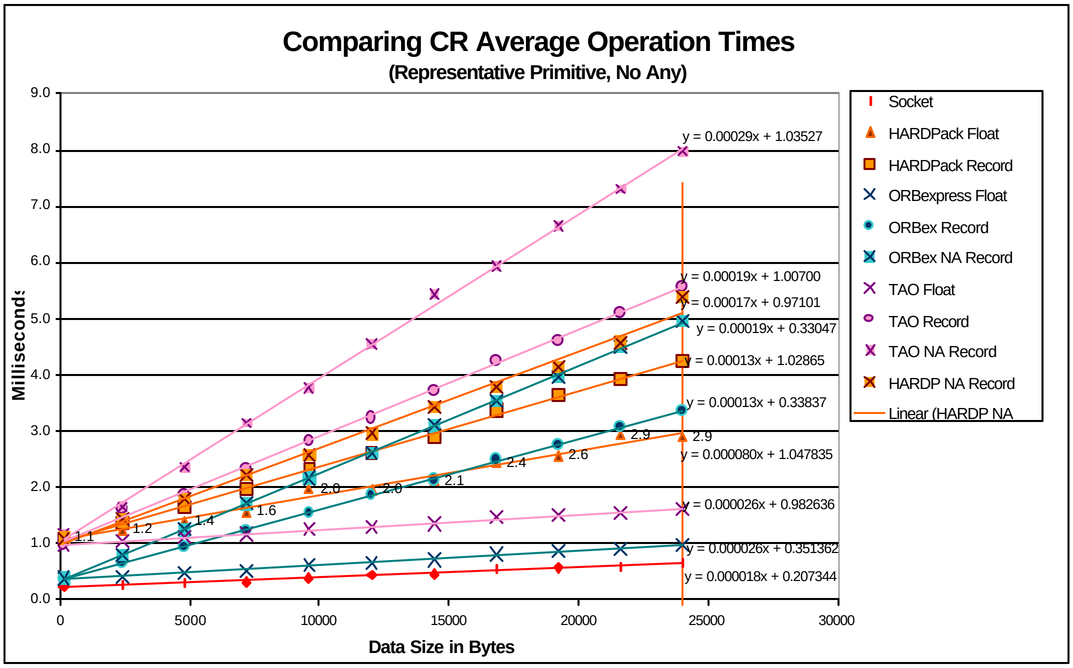

Figure 3. Trend Lines and Equations for CR Operations

The incremental cost of increasing data size in the transfers is approximately the same for ORB*express* and TAO for primitive data types. Table 1 contains the trend line equations computed by Microsoft Excel for the "primitive" (float) data sets displayed. In these equations, "x" represents the number of bytes of (application) data in each transfer, so the coefficient of x approximates the incremental cost (in milliseconds) of adding a byte of data to the message. The coefficients for ORB*express* and TAO indicate that the incremental data handling costs for these ORBs is essentially the same. In contrast, the cost of each incremental byte under HARDPack is roughly three times the cost for the other ORBs. Again, we reiterate that these trends apply only to transfers of primitive data inside a single SPARC host.

Table 1. Comparative Trends in CR Operations with Primitives

Middleware used	Trend line equations for "float" operations
Socket	y = 0.000018x + 0.207344
ORB*express*	y = 0.000026x + 0.351362
TAO	y = 0.000026x + 0.982636
HARDPack	y = 0.000080x + 1.047835

2.3 Aligned Records

The relationships change when the data is organized into records. As shown in Table 2 and Figure 2, raw performance information and incremental trends give HARDPack an advantage over TAO for all data sizes except the smallest. HARDPack roughly equals the incremental performance of ORB*express*, although with a much larger basic overhead. There is a caveat to the improving performance numbers for HARDPack, however: TAO and ORB*express* use IIOP in their handling of data, so these numbers represent performance based on the ORB standard for interoperability. HARDPack, by contrast, uses a proprietary protocol that may improve its marshalling performance for records by bypassing the standard. Further, the Basic Data Integrity tests described in other sections of this report showed that HARDPack was internally inconsistent in its handling of some data at the time of these tests. The integrity of these measurements is therefore suspect. They may or may not measure all the computation required to ensure the integrity of ORB transfers.

Table 2. Comparative Trends in CR Operations with (Aligned) Records

Middleware used	Trend line equations for "record" operations
Socket	y = 0.000018x + 0.207344
ORB*express*	y = 0.00013x + 0.33837
TAO	y = 0.00019x + 1.00700
HARDPack	y = 0.00013x + 1.02865

2.4 Non-aligned Records

The advantage of HARDPack over TAO in this environment persists when the records are not aligned on word boundaries. In this case, as characterized by the equations in Table 3, the incremental cost of increasing data size is lower for HARDPack than for ORB*express* as well. Unfortunately, the same caveats regarding protocol and integrity apply: Use HARDPack measurements cautiously unless the ORB environment is homogeneous and until the data integrity issues for HARDPack are resolved.

Table 3. Comparative Trends in CR Operations with Non-Aligned Records

Middleware used	Trend line equations for "NA record" operations
Socket	y = 0.000018x + 0.20734
ORB*express*	y = 0.00019x + 0.33047
TAO	y = 0.00029x + 1.03527
HARDPack	y = 0.00017x + 0.97101

2.5 Standard Deviations

Figure 4 plots standard deviations calculated for the data sets of the scenario. In studying these graphs, we are looking for data sets with unusual jitter and/or the highest number of or most excessive anomalies. Runs for both HARDPack and TAO show data sets with significantly large standard deviations. Because the HARDPack Any timings are much larger than other measured operation times, the unusually large standard deviation is not particularly surprising. With this dominating data set removed, as shown in Figure 5, other erratic behaviors are more easily observed.

Scenario 1a: Client, Server on Single Solaris Host

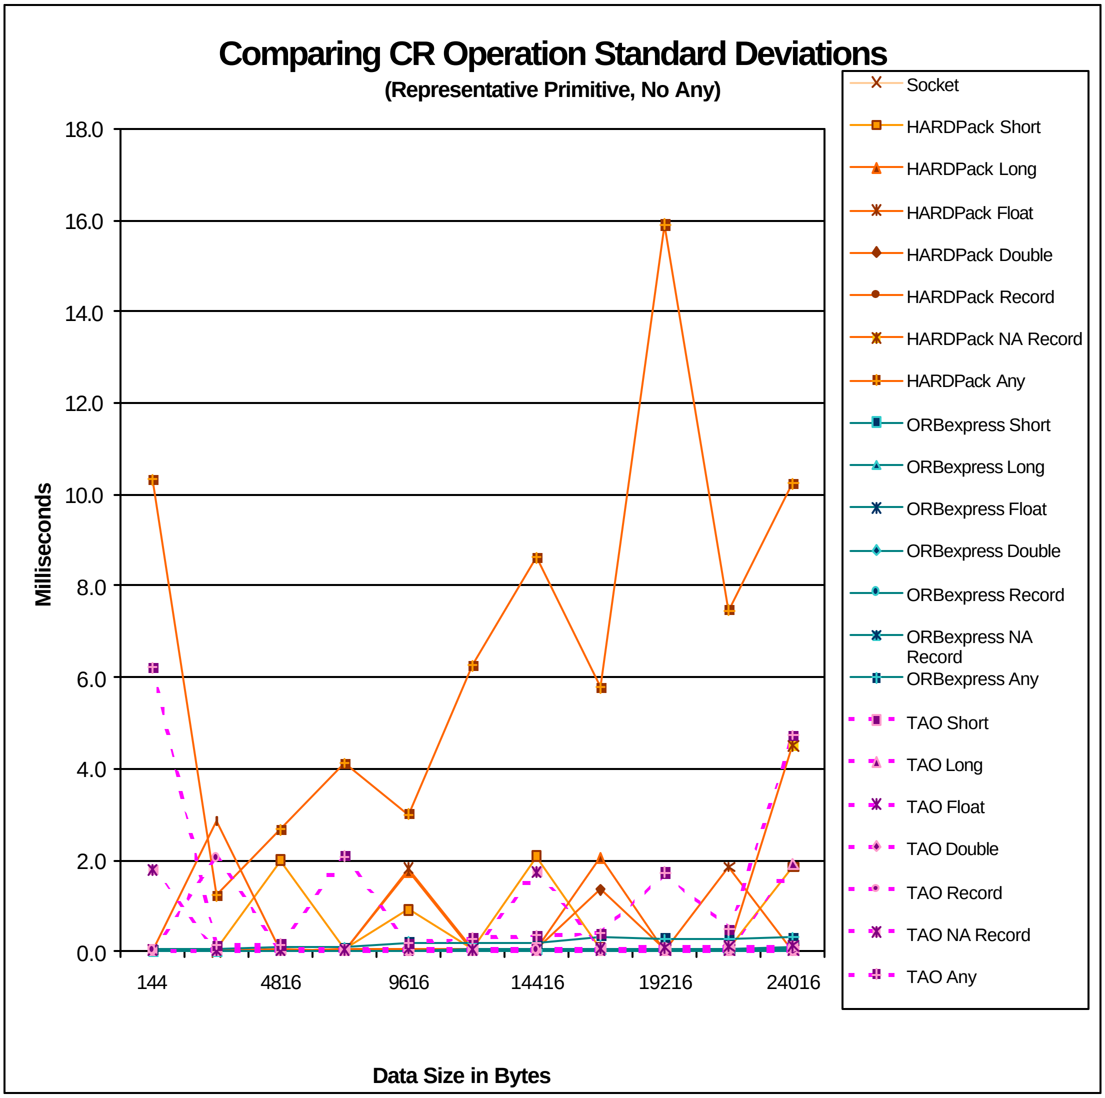

Figure 4. CR Operations in Single Solaris Host: Standard Deviations

Scenario 1a: Client, Server on a Single Solaris Host

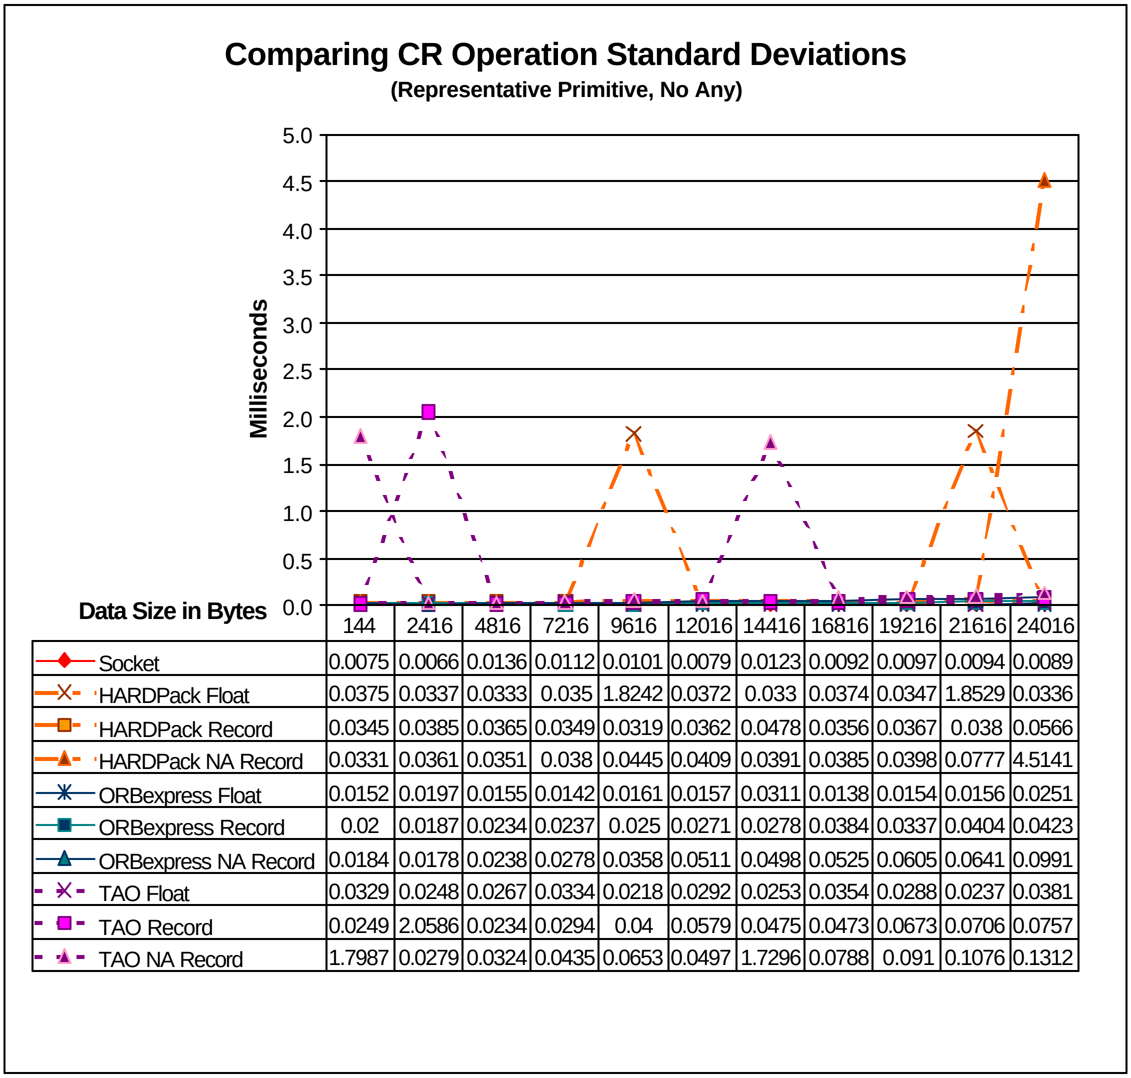

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
Socket	0.0075	0.0066	0.0136	0.0112	0.0101	0.0079	0.0123	0.0092	0.0097	0.0094	0.0089
HARDPack Float	0.0375	0.0337	0.0333	0.035	1.8242	0.0372	0.033	0.0374	0.0347	1.8529	0.0336
HARDPack Record	0.0345	0.0385	0.0365	0.0349	0.0319	0.0362	0.0478	0.0356	0.0367	0.038	0.0566
HARDPack NA Record	0.0331	0.0361	0.0351	0.038	0.0445	0.0409	0.0391	0.0385	0.0398	0.0777	4.5141
ORBexpress Float	0.0152	0.0197	0.0155	0.0142	0.0161	0.0157	0.0311	0.0138	0.0154	0.0156	0.0251
ORBexpress Record	0.02	0.0187	0.0234	0.0237	0.025	0.0271	0.0278	0.0384	0.0337	0.0404	0.0423
ORBexpress NA Record	0.0184	0.0178	0.0238	0.0278	0.0358	0.0511	0.0498	0.0525	0.0605	0.0641	0.0991
TAO Float	0.0329	0.0248	0.0267	0.0334	0.0218	0.0292	0.0253	0.0354	0.0288	0.0237	0.0381
TAO Record	0.0249	2.0586	0.0234	0.0294	0.04	0.0579	0.0475	0.0473	0.0673	0.0706	0.0757
TAO NA Record	1.7987	0.0279	0.0324	0.0435	0.0653	0.0497	1.7296	0.0788	0.091	0.1076	0.1312

Figure 5. CR Operations in Single Solaris Host: Standard Deviations (HARDPack Any Removed)

- Compared to the two other ORBs, ORB*express* shows very few discernable timing anomalies in this scenario. Its worst case behavior peaks in the Any series with a standard deviation at large message sizes that tops out at less than .5 milliseconds.
- HARDPack and TAO both show evidence of anomalous behaviors in selected data sets.
- For TAO, comparatively large standard deviations occur in three data sets: Any, NA Record, and Record. The performance for primitive data types appears to be solid and consistent. Detailed records available outside this report reveal more information about the anomalies. In the Any series, there is a disproportionate cost measured for the first sample in the series of transfers containing 4 (1st (smallest) data set), 225 (4th in increasing size), 600 (9th in

increasing size), and 750 (11th (largest) data set) structures. The anomalies in the NA Record series also occur on the first sample of each worrisome data set: the shortest data set (messages containing 4 NA records) and the 7th data set (messages containing 450 NA records). The Record series (aligned) showed only a single anomaly in the first sample of the data set with 75 records (2nd).

- For HARDPack, the unusual peaks are scattered liberally among primitive data types. Standard deviation peaks above 1 millisecond occur, at different message sizes, for shorts, longs, floats, and doubles. For the shorts, first sample maxima occur in data sets 3, 5, 7, and 11. For the longs, the same phenomenon occurs in data sets 5 and 8. For floats: data sets 5 and 10. For doubles: data sets 2 and 8.

These patterns occur with each repetition of the tests, but we have conducted no further analysis to determine their source. Neither have we run longer tests of more samples to see if such local maxima might recur later.

Since the peaks in CR operation times often occurred in the first sample, we recomputed the standard deviations for ORB*express* and TAO after removing the first five samples of each run. Figure 6 contains these “post startup” statistics.

Scenario 1a: Client, Server on Single Solaris Host

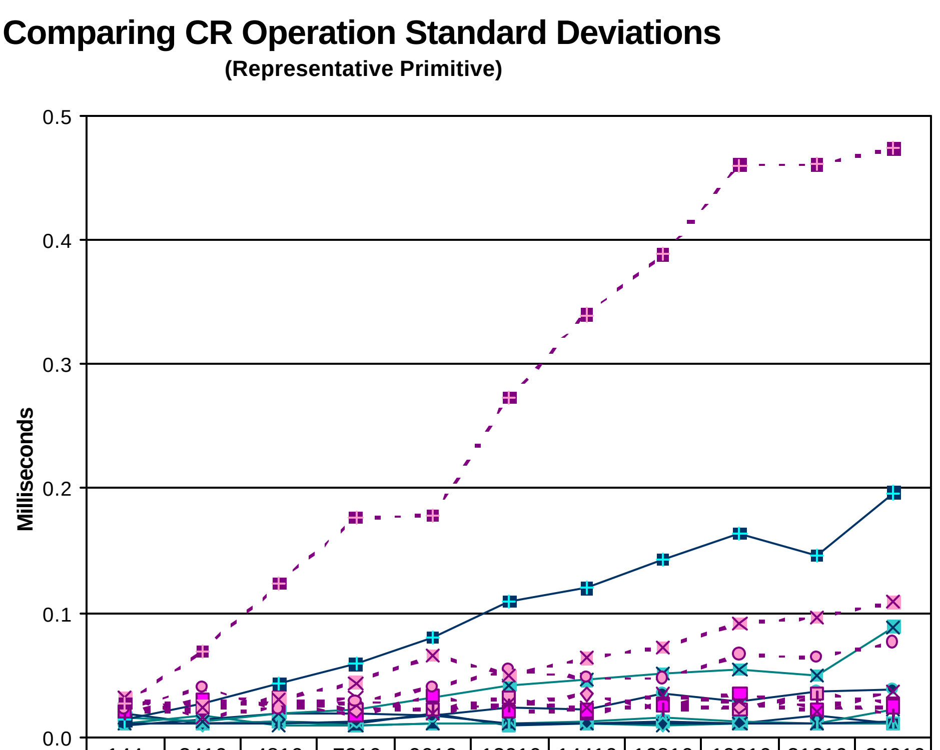

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
ORBexpress Short	0.0182	0.0114	0.013	0.0117	0.0193	0.0099	0.0117	0.0122	0.0112	0.0176	0.0113
ORBexpress Long	0.0108	0.0113	0.0124	0.0099	0.0115	0.0107	0.0119	0.0156	0.0129	0.0118	0.0104
ORBexpress Float	0.0115	0.0166	0.0099	0.0099	0.0107	0.0117	0.0107	0.01	0.0105	0.0115	0.0223
ORBexpress Double	0.0104	0.0102	0.0111	0.013	0.0175	0.0109	0.0117	0.0102	0.0118	0.0108	0.012
ORBexpress Record	0.0093	0.0148	0.0191	0.0191	0.0167	0.0243	0.0221	0.0357	0.0283	0.0371	0.0377
ORBexpress NA Record	0.0159	0.0127	0.0196	0.0214	0.0322	0.0416	0.0462	0.0513	0.0542	0.0498	0.0888
ORBexpress Any	0.0134	0.0264	0.0424	0.0586	0.0799	0.1092	0.1197	0.143	0.1639	0.1462	0.1967
TAO Short	0.0202	0.0234	0.0316	0.0214	0.0216	0.0325	0.0193	0.0262	0.0227	0.035	0.0274
TAO Long	0.02	0.0301	0.0259	0.0181	0.0337	0.0204	0.0215	0.0254	0.0358	0.0227	0.0246
TAO Float	0.0317	0.0155	0.0247	0.0322	0.0188	0.0275	0.0233	0.0342	0.0271	0.0213	0.0373
TAO Double	0.0192	0.0208	0.0316	0.0204	0.0232	0.0258	0.0343	0.0223	0.0235	0.0314	0.0173
TAO Record	0.0227	0.0401	0.0219	0.0283	0.0396	0.0541	0.0473	0.0472	0.0663	0.0637	0.076
TAO NA Record	0.0311	0.0244	0.0307	0.0437	0.0658	0.0499	0.0635	0.0723	0.0914	0.096	0.1083
TAO Any	0.027	0.0689	0.1237	0.1766	0.1782	0.2735	0.3396	0.3884	0.4601	0.4604	0.4736

Figure 6. CR Standard Deviations for ORB*express* and TAO with Startup Samples Removed

3. One-way Operations

3.1 Records and Primitives

Figure 7 summarizes the comparative performance of the three ORBs when the BasicIDL One-way methods execute with client, server, and background processes running in a single SPARC computer. As shown, the One-way operations for primitives for TAO and ORB*express* are relatively close in terms of average performance, with ORB*express* maintaining a small advantage. ORB*express* outperforms the other ORBs for Record and Non-aligned Record transfers. Direct comparison of One-way times, however, has relatively small value. The measurement on the client side indicates only how long it takes to queue a request for asynchronous handling and return to the caller. The measure does not document the total cost of the operation. Differences in ORB strategy regarding the amount of work to perform before returning to the caller may produce differences in timing that do not indicate how efficiently the ORB performs overall. Latency from client to server will often be a more valuable measurement and is reported below in server side data, although only for ORB*express* and TAO.

The measure of client latency for One-ways is *not* a useless piece of information, however. When it's important for a time-critical task to queue a less critical communication and continue to meet a deadline, this measure of One-way performance is of interest.

Scenario 1a: Client, Server on Single Solaris Host

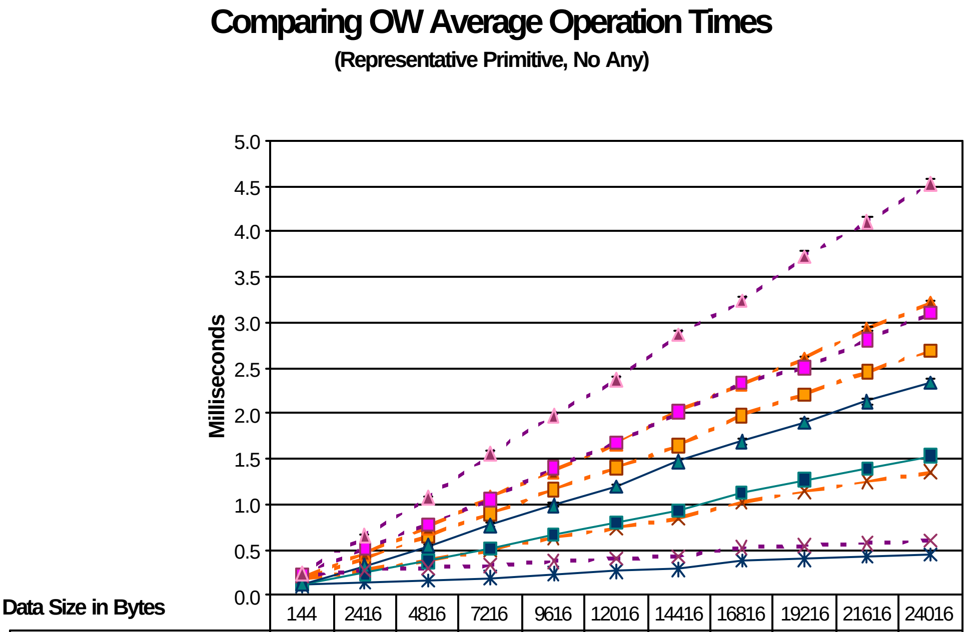

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
× HARDPack Float	0.1813	0.2749	0.3825	0.5033	0.6384	0.744	0.8532	1.0292	1.1409	1.2516	1.3562
■ HARDPack Record	0.1872	0.4084	0.6486	0.8987	1.1648	1.4058	1.6461	1.9815	2.2125	2.455	2.6915
▲ HARDPack NA Record	0.1897	0.4611	0.7526	1.078	1.3663	1.6683	2.0316	2.3226	2.6066	2.9311	3.2126
✱ ORBexpress Float	0.1154	0.143	0.1686	0.1958	0.2392	0.265	0.2901	0.3818	0.399	0.4253	0.4523
■ ORBexpress Record	0.118	0.2442	0.3784	0.5148	0.6717	0.7991	0.9366	1.1331	1.2721	1.3931	1.5348
▲ ORBexpress NA Record	0.1245	0.326	0.5437	0.7757	0.993	1.198	1.4753	1.6932	1.9022	2.1335	2.3419
× TAO Float	0.2218	0.2768	0.3059	0.33	0.3764	0.4046	0.4301	0.5247	0.5466	0.5738	0.5985
■ TAO Record	0.2305	0.5172	0.7819	1.0581	1.4075	1.6783	2.0139	2.3374	2.5	2.8077	3.1035
▲ TAO NA Record	0.2341	0.6556	1.0716	1.5536	1.9774	2.3706	2.8693	3.235	3.7253	4.1021	4.5252

Figure 7. One-way Operations on a Single Solaris Host: Average

One-way trend equations appear in Table 4. The similar x-coefficients for ORB*express* and TAO indicate that there are only small differences in the incremental data handling times for these ORBs. In One-way with primitive data types, as for Call & Return operations, the cost of each incremental byte under HARDPack is roughly three times the cost for the other ORBs. Again, we reiterate that these trends apply only to transfers of primitive data inside a single SPARC host.

Table 4. Comparative Trends in OW Operations with Primitives

Middleware used	Trend line equations for "float" operations
Socket	y = 0.00001x + 0.05229
ORB*express*	y = 0.000015x + 0.100169
TAO	y = 0.000016x + 0.223720
HARDPack	y = 0.000051x + 0.151081

3.2 Aligned Records

As shown in Figure 7 and Table 5, both the raw performance data and incremental trends give ORB*express* an advantage over TAO and HARDPack. HARDPack outperforms TAO in this series of tests, although the usual caveats regarding HARDPack performance still apply.

Table 5. Comparative Trends in OW Operations with (Aligned) Records

Middleware used	Trend line equations for "record" operations
Socket	y = 0.00001x + 0.05229
ORB*express*	y = 0.000060x + 0.093331
TAO	y = 0.000121x + 0.221608
HARDPack	y = 0.000107x + 0.146322

3.3 Non-aligned Records

The Record trends persist for Non-aligned Records with performance advantage falling to ORBexpress over HARDPack and HARDPack over TAO. As usual, however, the validity of the HARDPack performance is questionable until integrity issues are resolved.

Table 6. Comparative Trends in OW Operations with Non-Aligned Records

Middleware used	Trend line equations for "NA record" operations
Socket	y = 0.00001x + 0.05229
ORB*express*	y = 0.000094x + 0.099406
TAO	y = 0.000180x + 0.225404
HARDPack	y = 0.000128x + 0.153126

3.4 Standard Deviations

Figure 8 plots standard deviations calculated for the OW data sets of the scenario.

We find little cause for complaint in these numbers. We omitted the Any data from this graph for consistency and to show a little spread among these performance numbers. The standard deviations in the Any series were also very modest with all series exhibiting standard deviations under .6 milliseconds and most under .3 milliseconds.

Scenario 1a: Client, Server on Single Solaris Host

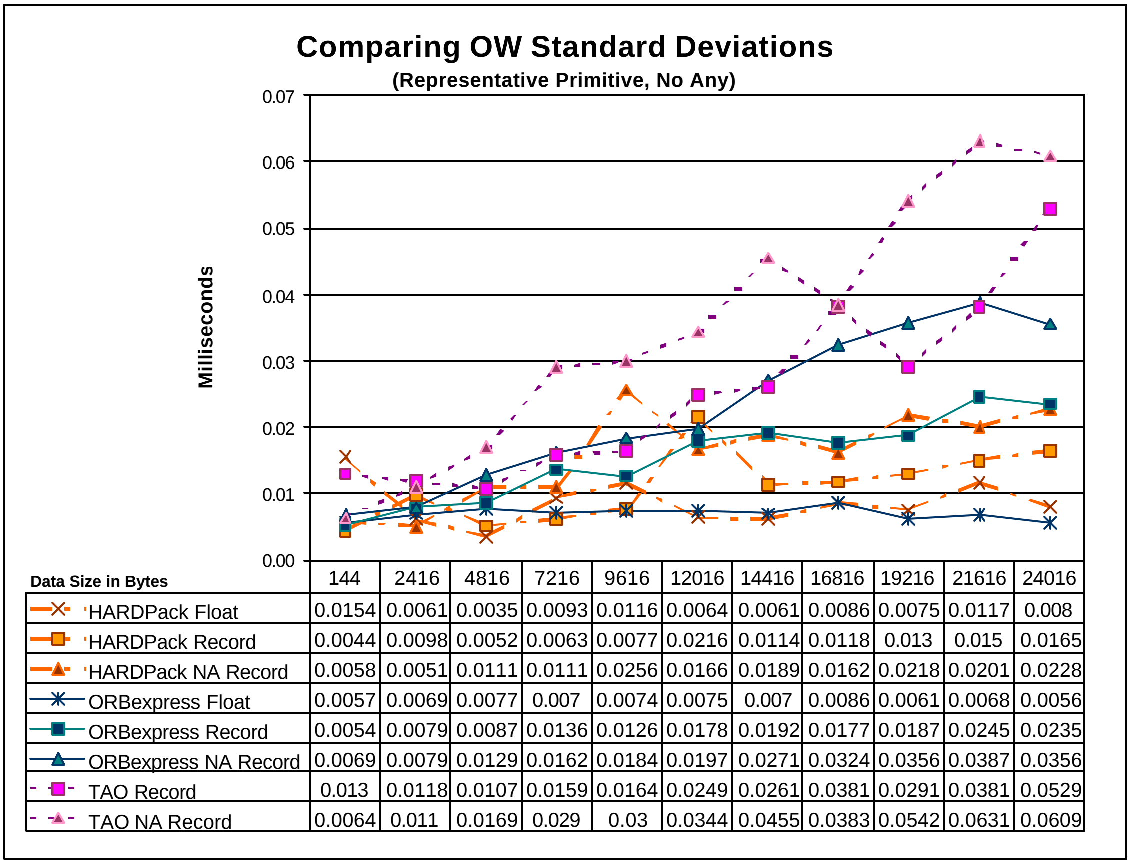

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
HARDPack Float	0.0154	0.0061	0.0035	0.0093	0.0116	0.0064	0.0061	0.0086	0.0075	0.0117	0.008
HARDPack Record	0.0044	0.0098	0.0052	0.0063	0.0077	0.0216	0.0114	0.0118	0.013	0.015	0.0165
HARDPack NA Record	0.0058	0.0051	0.0111	0.0111	0.0256	0.0166	0.0189	0.0162	0.0218	0.0201	0.0228
ORBexpress Float	0.0057	0.0069	0.0077	0.007	0.0074	0.0075	0.007	0.0086	0.0061	0.0068	0.0056
ORBexpress Record	0.0054	0.0079	0.0087	0.0136	0.0126	0.0178	0.0192	0.0177	0.0187	0.0245	0.0235
ORBexpress NA Record	0.0069	0.0079	0.0129	0.0162	0.0184	0.0197	0.0271	0.0324	0.0356	0.0387	0.0356
TAO Record	0.013	0.0118	0.0107	0.0159	0.0164	0.0249	0.0261	0.0381	0.0291	0.0381	0.0529
TAO NA Record	0.0064	0.011	0.0169	0.029	0.03	0.0344	0.0455	0.0383	0.0542	0.0631	0.0609

Figure 8. OW Operations on a Single Solaris Host: Standard Deviations

4. Server Side Data[1]

Our tests included measurements of latency from initiation of each operation by the client to receipt of the request by a servant. In Call & Return operations, the client suspends until the server returns a response, so the client-to-server latency is always shorter than the total operation time. Since the scenario 1a measurements are taken in a single machine, there are no issues regarding clock synchronization.

Figure 9 contains the average latency data for Call & Return operations in this single-SPARC scenario. Figure 10 shows the measurements for One-way operations.

Standard deviations for the same latency series appear in Figure 11 and Figure 12 for CR and OW operations, respectively. The Call & Return data tracks the client operation times. The One-way data has generally low standard deviations with the exception of a single data set for TAO. This disproportionate standard deviation derives from a 20 millisecond latency detected in

---

[1] Server side latency data was not available for HARDPack runs, so measurements for ORB*express* and TAO only are presented here.

the first transfer of the OW Float series, a behavior for which we have no reasonable explanation at this time.

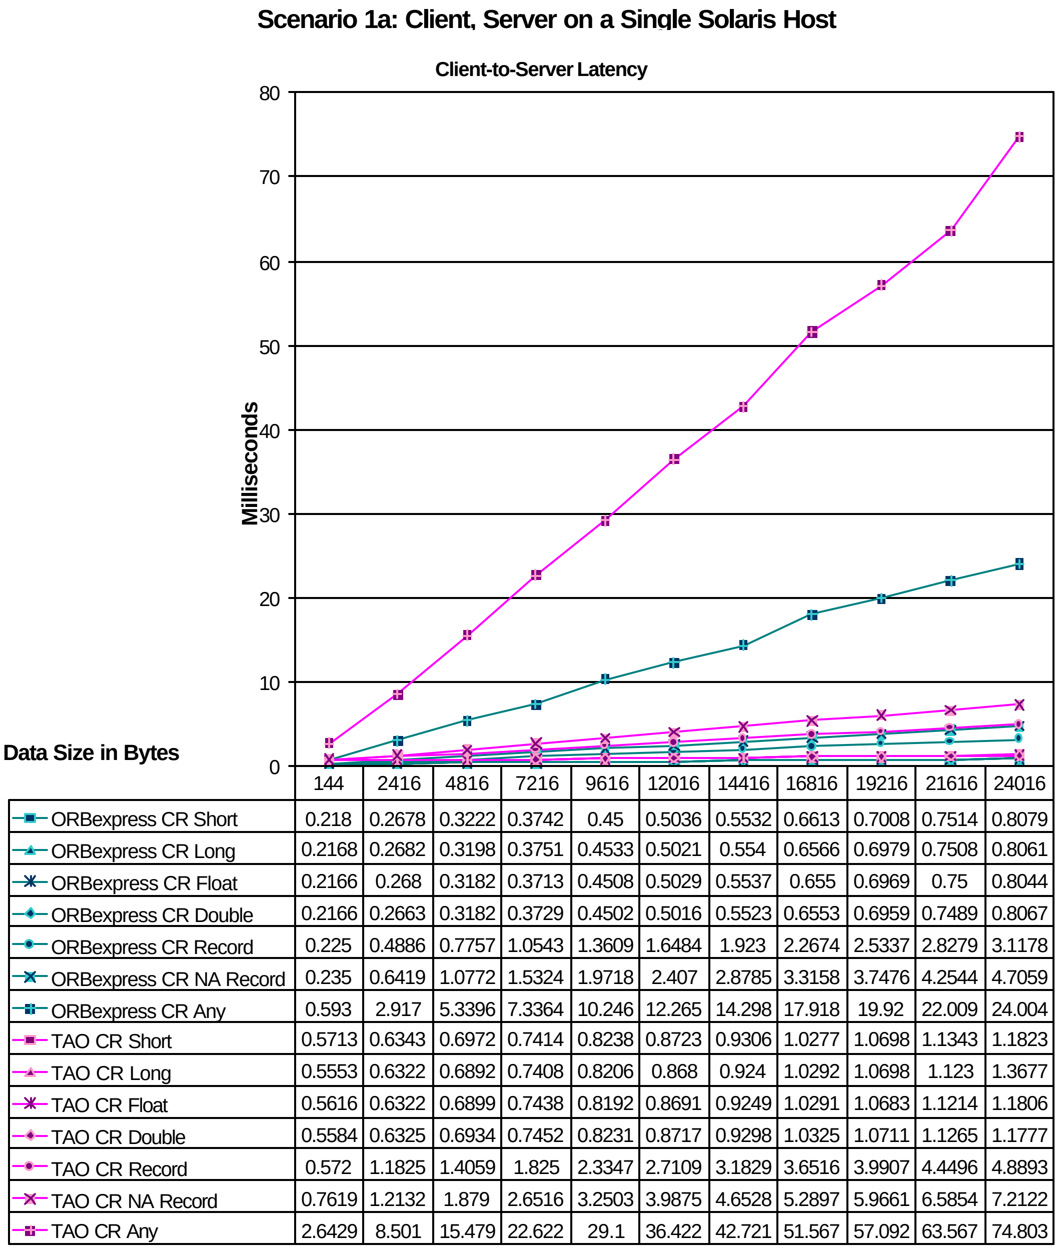

Scenario 1a: Client, Server on a Single Solaris Host

Client-to-Server Latency

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
■ ORBexpress CR Short	0.218	0.2678	0.3222	0.3742	0.45	0.5036	0.5532	0.6613	0.7008	0.7514	0.8079
▲ ORBexpress CR Long	0.2168	0.2682	0.3198	0.3751	0.4533	0.5021	0.554	0.6566	0.6979	0.7508	0.8061
✱ ORBexpress CR Float	0.2166	0.268	0.3182	0.3713	0.4508	0.5029	0.5537	0.655	0.6969	0.75	0.8044
◆ ORBexpress CR Double	0.2166	0.2663	0.3182	0.3729	0.4502	0.5016	0.5523	0.6553	0.6959	0.7489	0.8067
● ORBexpress CR Record	0.225	0.4886	0.7757	1.0543	1.3609	1.6484	1.923	2.2674	2.5337	2.8279	3.1178
✖ ORBexpress CR NA Record	0.235	0.6419	1.0772	1.5324	1.9718	2.407	2.8785	3.3158	3.7476	4.2544	4.7059
⊞ ORBexpress CR Any	0.593	2.917	5.3396	7.3364	10.246	12.265	14.298	17.918	19.92	22.009	24.004
■ TAO CR Short	0.5713	0.6343	0.6972	0.7414	0.8238	0.8723	0.9306	1.0277	1.0698	1.1343	1.1823
▲ TAO CR Long	0.5553	0.6322	0.6892	0.7408	0.8206	0.868	0.924	1.0292	1.0698	1.123	1.3677
✱ TAO CR Float	0.5616	0.6322	0.6899	0.7438	0.8192	0.8691	0.9249	1.0291	1.0683	1.1214	1.1806
◆ TAO CR Double	0.5584	0.6325	0.6934	0.7452	0.8231	0.8717	0.9298	1.0325	1.0711	1.1265	1.1777
● TAO CR Record	0.572	1.1825	1.4059	1.825	2.3347	2.7109	3.1829	3.6516	3.9907	4.4496	4.8893
✖ TAO CR NA Record	0.7619	1.2132	1.879	2.6516	3.2503	3.9875	4.6528	5.2897	5.9661	6.5854	7.2122
⊞ TAO CR Any	2.6429	8.501	15.479	22.622	29.1	36.422	42.721	51.567	57.092	63.567	74.803

Figure 9. Client to Server Latency for CR Operations: Average

Scenario 1a: Client, Server on Single Solaris Host

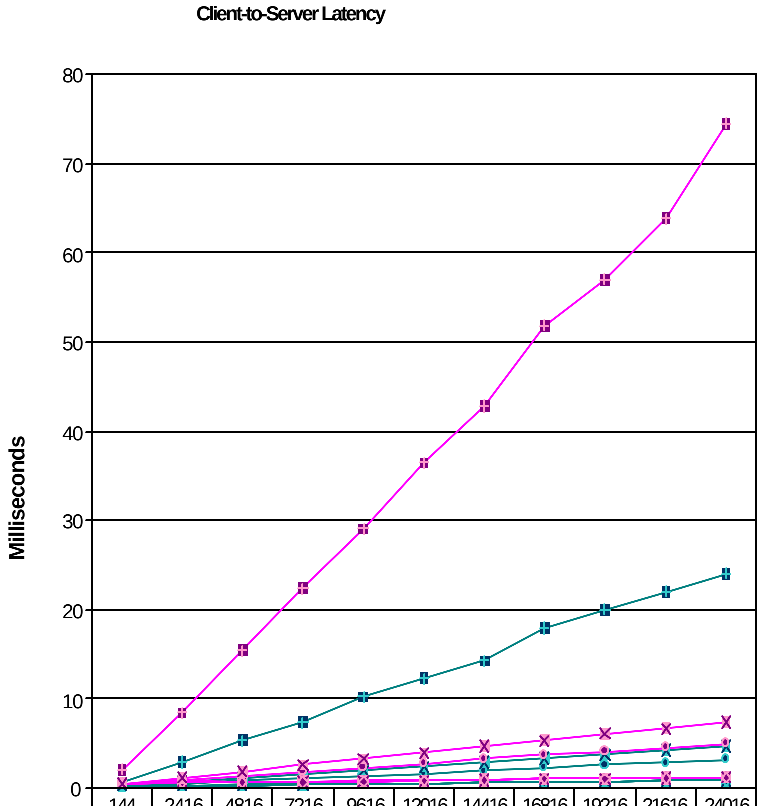

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
ORBexpress OW Short	0.2117	0.265	0.3154	0.3694	0.4453	0.5005	0.5494	0.6793	0.7174	0.7717	0.8277
ORBexpress OW Long	0.2137	0.2627	0.3194	0.3665	0.45	0.4951	0.5534	0.678	0.721	0.7721	0.8284
ORBexpress OW Float	0.2126	0.2642	0.3166	0.3702	0.4451	0.4994	0.5484	0.6818	0.7168	0.7717	0.8267
ORBexpress OW Double	0.214	0.2615	0.3163	0.3666	0.4461	0.4979	0.5496	0.677	0.7166	0.7712	0.8269
ORBexpress OW Record	0.2199	0.4852	0.7699	1.0509	1.3632	1.6435	1.9263	2.281	2.5778	2.848	3.1576
ORBexpress OW NA Record	0.2303	0.6409	1.0777	1.535	1.9714	2.3923	2.9055	3.3388	3.7715	4.2466	4.6816
ORBexpress OW Any	0.587	2.8962	5.3246	7.3353	10.239	12.253	14.258	17.92	19.987	22.032	24.025
TAO OW Short	0.4926	0.5735	0.6284	0.6781	0.7619	0.8135	0.8646	0.9927	1.0385	1.0944	1.1511
TAO OW Long	0.4914	0.5746	0.6261	0.6796	0.7569	0.8097	0.863	0.9934	1.0394	1.0881	1.149
TAO OW Float	0.4909	0.7671	0.6293	0.6769	0.7574	0.8104	0.8644	0.9949	1.0384	1.0956	1.148
TAO OW Double	0.492	0.5783	0.6291	0.6778	0.7592	0.8136	0.865	0.9969	1.0427	1.0931	1.1495
TAO OW Record	0.5071	0.941	1.3609	1.8012	2.3244	2.7166	3.2246	3.7264	4.0528	4.4783	4.9618
TAO OW NA Record	0.5136	1.1774	1.8476	2.5964	3.2698	3.894	4.6766	5.3066	6.0661	6.6794	7.3763
TAO OW Any	2.0435	8.4464	15.42	22.365	29.057	36.425	42.818	51.817	56.962	63.847	74.445

Figure 10. Client to Server Latency for OW Operations: Average

Scenario 1a: Client, Server on Single Solaris Host

Client-to-Server Latency: Standard Deviation

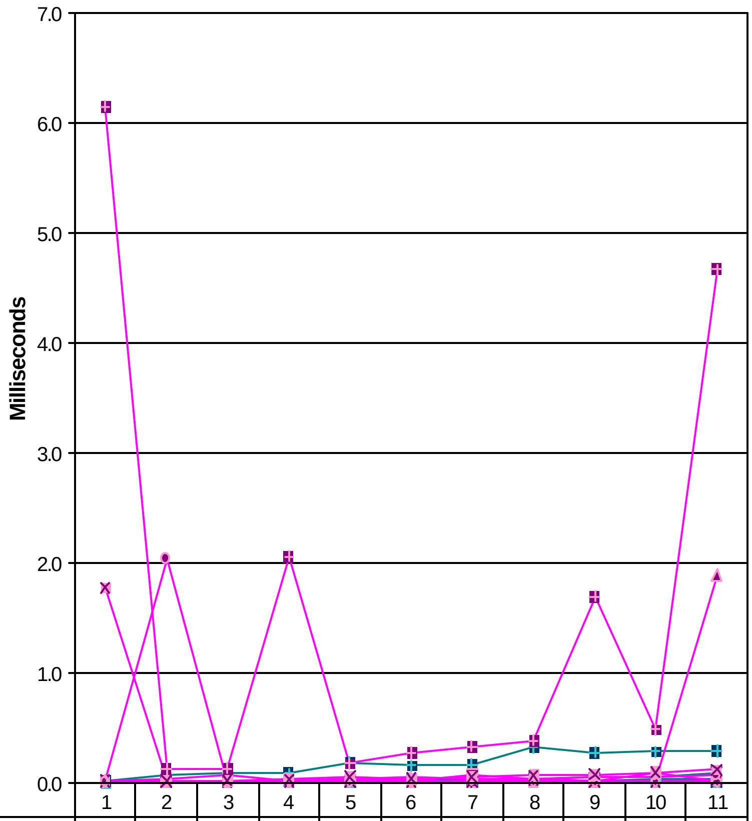

Data Size in Bytes	1	2	3	4	5	6	7	8	9	10	11
-■- ORBexpress CR Short	0.0173	0.0144	0.0321	0.0284	0.0212	0.0315	0.0154	0.0519	0.0315	0.018	0.0286
-▲- ORBexpress CR Long	0.0132	0.0152	0.0147	0.0224	0.0143	0.0154	0.0142	0.0182	0.0144	0.0161	0.0131
-✱- ORBexpress CR Float	0.0139	0.0177	0.0147	0.0132	0.0149	0.0141	0.0295	0.0128	0.0143	0.0142	0.0191
-◆- ORBexpress CR Double	0.013	0.014	0.0143	0.0158	0.0178	0.0144	0.0147	0.0141	0.0139	0.0153	0.0142
-●- ORBexpress CR Record	0.0126	0.018	0.0217	0.0226	0.0238	0.0253	0.0261	0.0375	0.0319	0.0382	0.0406
-✖- ORBexpress CR NA Record	0.0166	0.0173	0.0229	0.0265	0.0347	0.0489	0.0486	0.0509	0.0597	0.0619	0.0968
-⊞- ORBexpress CR Any	0.0297	0.0753	0.1055	0.095	0.1923	0.1638	0.1729	0.332	0.2882	0.2916	0.3008
-■- TAO CR Short	0.0203	0.046	0.0799	0.021	0.0419	0.0292	0.0814	0.0392	0.0215	0.1012	0.0334
-▲- TAO CR Long	0.0193	0.0275	0.024	0.0212	0.0302	0.0201	0.0206	0.0239	0.0312	0.0212	1.8999
-✱- TAO CR Float	0.0281	0.0218	0.0234	0.0283	0.0193	0.0253	0.0215	0.0301	0.025	0.0201	0.0328
-◆- TAO CR Double	0.0243	0.0216	0.028	0.0201	0.022	0.0239	0.0301	0.0213	0.0221	0.0276	0.0192
-●- TAO CR Record	0.0214	2.0441	0.0212	0.0267	0.0367	0.0544	0.0458	0.0451	0.0638	0.0652	0.0738
-✖- TAO CR NA Record	1.7836	0.0254	0.0297	0.0417	0.0614	0.0482	0.0618	0.0767	0.088	0.1038	0.1273
-⊞- TAO CR Any	6.1521	0.1335	0.1442	2.0687	0.1841	0.2809	0.3382	0.3918	1.702	0.4918	4.6754

Figure 11. Client to Server Latency for CR Operations: Standard Deviations

Scenario 1a: Client, Server on Single Solaris Host

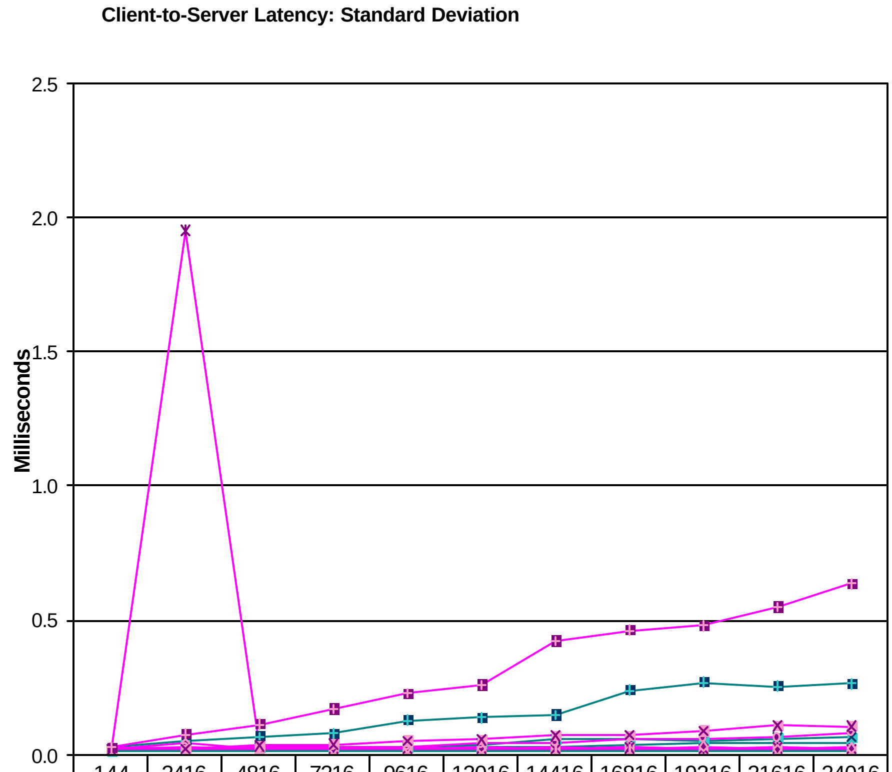

Data Size in Bytes	144	2416	4816	7216	9616	12016	14416	16816	19216	21616	24016
ORBexpress OW Short	0.0136	0.0175	0.0138	0.0134	0.0149	0.0167	0.014	0.0143	0.0122	0.0166	0.0177
ORBexpress OW Long	0.0133	0.0131	0.014	0.0134	0.0172	0.0137	0.0155	0.0129	0.0182	0.0151	0.0158
ORBexpress OW Float	0.0116	0.0143	0.016	0.0144	0.0143	0.0145	0.0139	0.0174	0.0128	0.0137	0.0132
ORBexpress OW Double	0.0163	0.014	0.0145	0.0131	0.0148	0.0214	0.0136	0.0138	0.014	0.0134	0.0201
ORBexpress OW Record	0.0115	0.0163	0.0183	0.0299	0.0221	0.028	0.0282	0.0347	0.0417	0.0421	0.041
ORBexpress OW NA Record	0.0137	0.0173	0.0245	0.0271	0.0288	0.0385	0.0573	0.054	0.0516	0.0609	0.0645
ORBexpress OW Any	0.0253	0.0463	0.0659	0.0763	0.1265	0.1376	0.1472	0.2375	0.2688	0.2542	0.2632
TAO OW Short	0.0194	0.0276	0.0191	0.0238	0.0304	0.02	0.0213	0.0256	0.0212	0.0235	0.0193
TAO OW Long	0.0257	0.0209	0.0209	0.0275	0.0205	0.0197	0.0258	0.0216	0.0215	0.0216	0.0281
TAO OW Float	0.0178	1.9515	0.0284	0.0193	0.0218	0.0264	0.021	0.0244	0.0192	0.0254	0.0209
TAO OW Double	0.0209	0.0409	0.0201	0.0186	0.026	0.022	0.0206	0.0209	0.0288	0.0176	0.0238
TAO OW Record	0.0264	0.0232	0.0243	0.031	0.0304	0.0397	0.043	0.0608	0.0609	0.0619	0.0831
TAO OW NA Record	0.0189	0.0203	0.0325	0.0383	0.0503	0.0543	0.0693	0.071	0.0872	0.1066	0.1048
TAO OW Any	0.0243	0.0739	0.1115	0.1697	0.228	0.2588	0.4229	0.462	0.4803	0.5492	0.6364

Figure 12. Client to Server Latency for OW Operations: Standard Deviations

Glossary

ACE	ADAPTIVE Communication Environment
ADAPTIVE	A Dynamically Assembled Protocol, Transformation and Validation Environment
AWACS	Airborne Warning and Control System
BDI	Basic data integrity
CORBA	Common Object Request Broker Architecture
CR	Call and return
DII COE	Defense Information Infrastructure Common Operating Environment
IDL	Interface definition language
IIOP	Internet inter-ORB protocol
IPT	Integrated Product Team
JTT	Joint Tactical Terminal
LMFS	Lockheed Martin Federal Systems (Produces and supports HARDPack)
NA	Non-aligned
OCI	Object Computing, Inc. (Supports TAO)
OIS	Objective Interface Systems (Produces and supports ORB*express*)
OMG	Object Management Group
ORB	Object request broker
OS	Operating system
OW	One way
POA	Portable Object Adapter
PPC	Power PC
RT	Real-time
RTOS	Real-time operating system
TAO	The ACE ORB
TWG	Technical Working Group

Active Page Record

Page Numbers	Revision Level	Revision Type (Added, Deleted)
I		
Ii		
Iii		
Iv		
v		
1		
2		
3		
4		
5		
6		
7		
8		
9		
10		
11		
12		
13		
14		
15		
16		
17		
Glossary-1		
APR-1		

Page Numbers	Revision Level	Revision Type (Added, Deleted)