APPENDIX C

MOBILE AGENT MIGRATIONS TIMING EXPERIMENTS

C.1. INTRODUCTION

In this appendix, we present the result of timing experiments that aim at providing a better understanding of the response time measured for Mobile Agent (MA) migrations. The objective is to investigate how the overall time is distributed among the individual phases of an MA migration and what is the effect of a number of factors such as the transport protocol used and the MA state size. In particular, the migration process is first divided into several independent, successive phases in order to assess the weight of their respective delays on the overall latency.

C.2. MOBILE AGENT MIGRATION TIMING MEASUREMENTS

The MA used for this experiment, termed ‘ping-pong’ MA, is launched by the manager application, visits a single remote Mobile Agent Server (MAS) server and then returns back to the manager, *without* performing any kind of task. Therefore, the phases identified are:

1. MA object creation (instantiation);
2. MA object serialisation by the manager application;
3. MA transfer to the MAS server;
4. MA object de-serialisation at the MAS side;
5. MA object serialisation by the MAS server;
6. MA transfer back to the manager station;
7. MA object de-serialisation at the manager side.

Both TCP and UDP transport protocols are considered for the MA transfers, while the serialised MA state can either be compressed before its transmission or non-compressed. Another investigation issue is the effect of the serialised MA's state size on the overall migration delay. In particular, one of the experiment's objectives is to examine to what extend does the increment of the state size affects the distribution of the individual migration phases delay. The average times corresponding to each migration phase are presented in Table C.1 and Table C.2, with the former corresponding to an MA object with relatively moderate state size (476 bytes when compressed, 678 bytes when uncompressed) and the latter to an MA with larger state size (1152 bytes when compressed, 3970 bytes when uncompressed). The state size growth is realised by increasing the number of the MA's non-transient variables [ARN96].

Migration Phase	*MA Creation*	*Manager Ser/tion*	*MAS Deser/tion*	*MAS Ser/tion*	*Manager Deser/tion*	*MA Transfer*	*Overall Time*
TCP, compressed	6.12	3.24	8.06	5.19	5.37	56.04	84.02
TCP, non-compressed	4.08	0.74	6.22	0.46	6.32	59.15	76.98
UDP, compressed	4.92	3.67	4.56	6.02	6.23	53.71	79.10
UDP, non-compressed	8.28	2.66	2.89	1.74	2.31	55.79	73.68

Table C.1: Time measurements (in msec) depicting the distribution of delays for a "ping-pong" MA with a 'small' state size (476 bytes compressed / 678 bytes uncompressed) during its lifetime (each table value represents 100 individual time measurements)

Migration Phase	*MA Creation*	*Manager Ser/tion*	*MAS Deser/tion*	*MAS Ser/tion*	*Manager Deser/tion*	*MA Transfer*	*Overall Time*
TCP, compressed	4.82	8.03	151.51	7.41	23.97	66.68	262.43
TCP, non-compressed	9.81	3.09	148.13	2.84	29.10	89.20	282.17
UDP, compressed	8.34	7.12	149.42	6.60	18.43	61.31	251.22
UDP, non-compressed	11.22	8.85	149.25	4.28	7.11	80.18	260.88

Table C.2: Time measurements (in msec) depicting the distribution of delays for a "ping-pong" MA with a 'large' state size (1152 bytes compressed / 3970 bytes uncompressed) during its lifetime (each table value represents 100 individual time measurements)

The main reason for conducting time measurement experiments with 'ping-pong' rather than with single-hop MAs is that it would not be feasible to measure MA transfers over UDP, as explained in Section 4.6.1.1. It is noted that the column representing the transfer delay represents the summation of the time needed to transfer the MA object from the manager host to the MAS site and vice-versa.

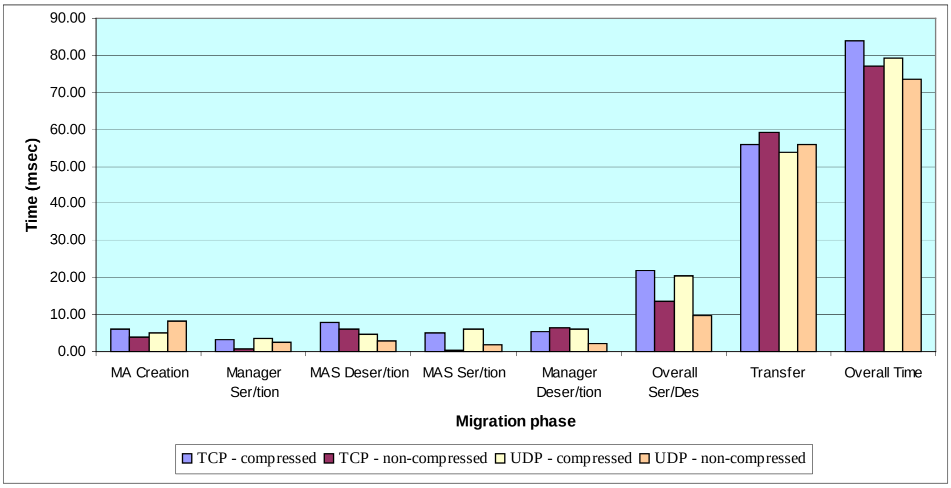

(a)

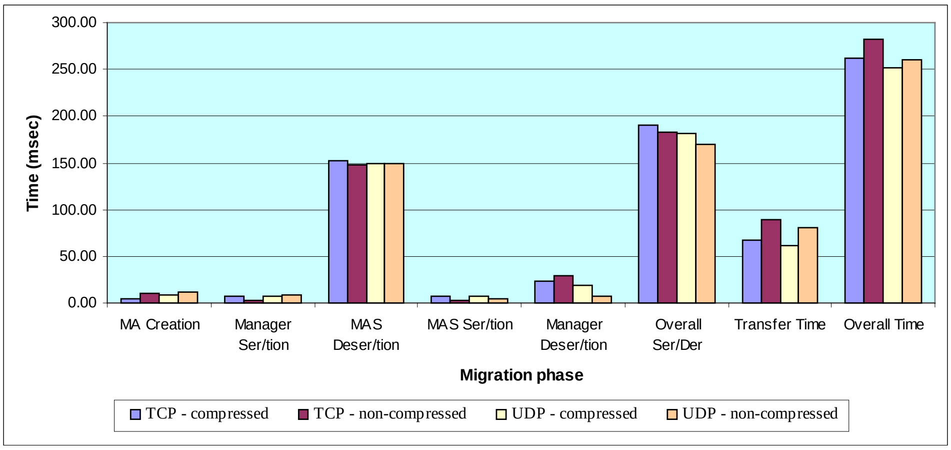

(b)

Figure C.1. Comparison of the delays experienced within the individual migration phases of an MA with (a) moderate or (b) large state size, depending on the transport protocol and on whether the serialised state is compressed prior to its transfer

The distribution of the overall response time among the phases that compose the migration process are graphically illustrated for both the types of MAs used in these experiments, i.e. with moderate (Figure C.1a) or large state size (Figure C.1b). It is interesting to observe the way that compression affects the response time as a function of the volume of compressed data. In particular, regardless of the utilised transport protocol (TCP or UDP), MA's state compression always causes reduction of transfer time (the volume of transmitted data is decreased) and increment of the overall serialisation/se-serialisation time (compression/de-compression is regarded as part of the serialisation/se-serialisation process). However, the effect of compression on the overall response time largely depends on the MA's state size. For serialised state of moderate size, compression increases the overall time. Conversely, when the migration of MAs with larger state sizes is considered, compression results in a reduction of the response time. Therefore, compression reduces the latency of the migration process in case that the original size of the serialised state and the compression ratio is such so as to justify the time penalty involved in the compression process. To illustrate, in our first type of MA, compression saves only 678-476 = 202 bytes from being transferred through the network (29.8% of the original data volume), while in the second type the saving becomes 3970 - 1152 = 2818 (71% of the original data volume).

Another interesting aspect of the MA migration time measurements is the alterations that state size causes on the distribution of the overall response time among the individual migration phases. The 'pie' diagrams shown in Figure C.2 confirm that increasing the MA's state size results in shifting response time "centre of mass" from transfer to the serialisation/de-serialisation process. For instance, the transfer latency of an MA with moderate state size covers the 66.7% of the overall response time, with the sum of the individual serialisation/de-serialisation times comprising the 26.1%. In contrast, when dealing with the MA with large state size, these percentages become 25.4% and 72.1% respectively.

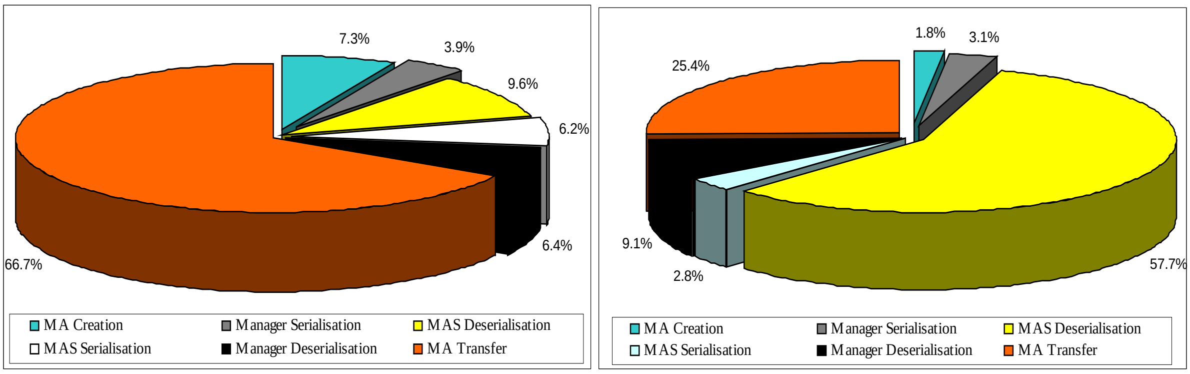

(a) (b)

Figure C.2. Distribution (percentage of the overall response time) of the delays incurred within the individual migration phases for MAs with either (a) moderate or (b) fairly large state size (the serialised state is compressed)

Interestingly, de-serialisation process is proved more time consuming than its inverse process of serialisation. In addition, a significant portion of the response time is covered by the MA's state de-serialisation, which takes place at the MAS server side (see Figure C.2b). This is due to the fact that MAS entities use the customised ClassLoader (MACL) described in Section 4.4.3, which is slower than the default JVM ClassLoader used by the manager application. The serialisation process taking place at the MAS server is significantly faster than the de-serialisation, since it uses the MA class definition that has been already loaded by the MACL during the serialisation phase.

In general, the response time difference observed between the two types of MAs is mainly due to the difference in the overall serialisation/de-serialisation time, as shown in Figure C.3. Certainly, the transfer time also increases in the case of the MA with larger state size; however, this is not the decisive factor.

A possible optimisation would be to save the time needed to create (instantiate) the MA object, as suggested in Appendix B. That could be realised through creating the MA beforehand, that is, before the actual management operation commences. Such an optimisation would result in performance gain, especially in the case of MAs with small state size, when MA creation represents a considerable proportion of the overall time (7.3%).

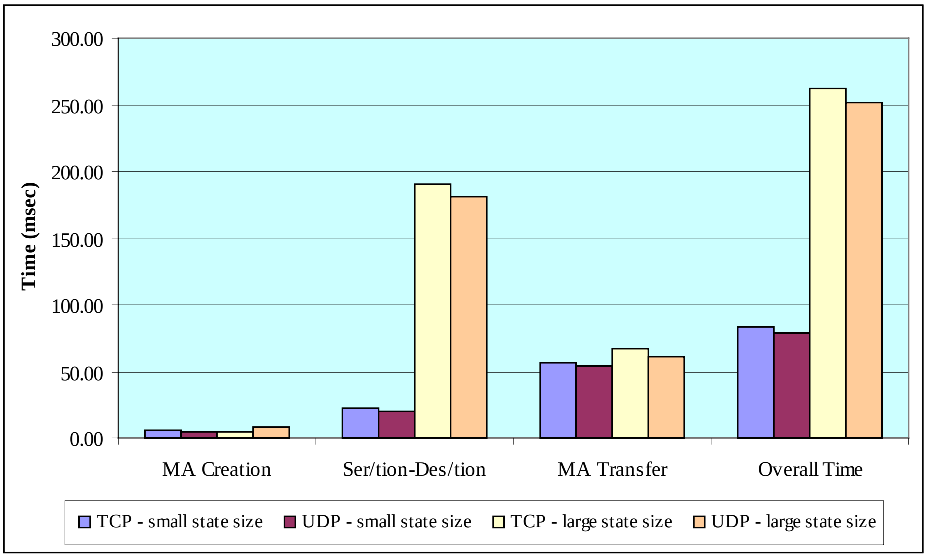

Figure C.3. Comparison of the individual migration phases delays for MAs with moderate vs. large state size (the serialised state is compressed prior to its transmission)