Performance Evaluation of Plausible Clocks

Francisco J. Torres-Rojas
Centro de Investigaciones en Computación (CIC), Costa Rica Institute of Technology
Costa Rica
torres@ic-itcr.ac.cr

Abstract. Plausible Clocks do not *characterize* causality [6] but, under appropriate circumstances, their accuracy is close to vector clocks. This paper explores the effects that several factors have on the performance of these clocks.

1 Introduction

Vector clocks capture the causality relation between events in a distributed history H [2, 4]. However, these clocks require one entry for each one of the N sites of the system [1], which causes scalability problems [6, 7]. *Plausible clocks* are an scalable alternative to vector clocks that do not characterize causality but whose accuracy can be close to vector clocks under appropriate circumstances [8]. A detailed description of these clocks is presented in [7, 8]. This paper ponders the effects that several factors have on the performance of the plausible clock *Comb* [8] and identifies the conditions under which it performs better. The parameter ρ (proportion of cases whose causal relationship is incorrectly reported by *Comb*) is used as the response variable. Among other results, it was found that in client/server systems with high communication frequency, this clock can order more than 90% of the event pairs in the same way as vector clocks. Section **2** describes and analyzes a multifactorial experiment that considers the effects of several factors on the performance of *Comb*. Section **3** presents an experiment where the scalability of *Comb*, under conditions for best performance, is tested. Section **4** gives the conclusions of this paper.

2 Experiment 1: Multifactorial Design

We explore the effects that the following factors have on ρ: system size (20, 50 and 80 sites), clock size (6, 12 and 18 entries), communication pattern (random and client/server), local history size (10, 35 and 60 events), and message interval mean (exponential distribution with mean $1/\lambda = 3$, 14 and 25 events). For each one of the 162 possible combinations of factor levels, 3 distributed histories are generated, i.e., there are 486 trials of the experiment. Each history is executed, sending and receiving messages and keeping the timestamps assigned to each event by *Comb*, together with vector clocks. Then, we determine the causal relationships between all the ordered pairs in the set $H \times H$ from the point of view of *Comb* [8], compare them with the actual causal relationships, and compute ρ. A total of 2,969,821,897 pairs of events were considered. The data set for this experiment can be found in [9]. The statistical procedure Analysis of Variance (ANOVA) checks k samples obtained from k populations, and determines if there is statistical evidence that some of them have different means [5]. The assumptions of ANOVA (random error normally and independently distributed with mean 0.0 and constant variance) are verified in [9] and,

R. Sakellariou et al. (Eds.): Euro-Par 2001, LNCS 2150, pp. 476-481, 2001.

hence, its conclusions are statistically valid. Each one of the 162 combinations of factor levels defines a population characterized by a particular mean of ρ. We test if all these populations have the same mean, or if there is evidence that some of them have different means (i.e., that the investigated factors have significant effects on the behavior of *Comb*). Figure 1 shows the ANOVA output of the statistical package *Minitab* for this experiment. Each row corresponds to a possible source of variation in ρ. Column `P` presents the probability p of observing the analyzed set of samples when the particular source of variation has **no** significant effect on ρ. Usually, a factor whose value of p is greater than 0.01, is considered as not significant [5].

MTB > ANOVA 'RATE' = C1!C2!C3!C4!C5.

Factor	Type	Levels	Values		
Sites	fixed	3	20	50	80
Clock	fixed	3	6	12	18
Pattern	fixed	2	0	1	
Events	fixed	3	10	35	60
Mean	fixed	3	3	14	25

Analysis of Variance for RATE

Source	DF	SS	MS	F	P
Sites	2	0.495746	0.247873	388.71	0.000
Clock	2	0.973416	0.486708	763.26	0.000
Pattern	1	0.167167	0.167167	262.15	0.000
Events	2	0.072725	0.036362	57.02	0.000
Mean	2	0.210761	0.105381	165.26	0.000
Sites*Clock	4	0.002875	0.000719	1.13	0.344
Sites*Pattern	2	0.015301	0.007650	12.00	0.000
Sites*Events	4	0.004351	0.001088	1.71	0.148
Sites*Mean	4	0.003962	0.000990	1.55	0.187
Clock*Pattern	2	0.147005	0.073503	115.27	0.000
Clock*Events	4	0.001375	0.000344	0.54	0.707
Clock*Mean	4	0.035201	0.008800	13.80	0.000
Pattern*Events	2	0.308820	0.154410	242.15	0.000
Pattern*Mean	2	0.713878	0.356939	559.75	0.000
Events*Mean	4	0.933027	0.233257	365.79	0.000
Sites*Clock*Pattern	4	0.005532	0.001383	2.17	0.072
Sites*Clock*Events	8	0.007138	0.000892	1.40	0.196
Sites*Clock*Mean	8	0.003945	0.000493	0.77	0.627
Sites*Pattern*Events	4	0.064593	0.016148	25.32	0.000
Sites*Pattern*Mean	4	0.079552	0.019888	31.19	0.000
Sites*Events*Mean	8	0.010190	0.001274	2.00	0.046
Clock*Pattern*Events	4	0.011423	0.002856	4.48	0.002
Clock*Pattern*Mean	4	0.019763	0.004941	7.75	0.000
Clock*Events*Mean	8	0.082474	0.010309	16.17	0.000
Pattern*Events*Mean	4	0.227229	0.056807	89.09	0.000
Sites*Clock*Pattern*Events	8	0.029874	0.003734	5.86	0.000
Sites*Clock*Pattern*Mean	8	0.013133	0.001642	2.57	0.010
Sites*Clock*Events*Mean	16	0.013237	0.000827	1.30	0.196
Sites*Pattern*Events*Mean	8	0.056771	0.007096	11.13	0.000
Clock*Pattern*Events*Mean	8	0.048146	0.006018	9.44	0.000
Sites*Clock*Pattern*Events*Mean	16	0.017959	0.001122	1.76	0.035
Error	324	0.206606	0.000638		
Total	485	4.983176			

Figure 1. Analysis of Variance for Experiment 1 (*Minitab* Output).

All main factors are significant. The average effects are: ρ increases when the system is larger, ρ decreases when the clocks have more entries, and ρ is lower in client/server systems. There are 6 significant second level interactions, 4 of them include the communication pattern, while the system size is present in just one interaction. Figure 2 presents the average effect of the second level interactions. The value of ρ reduces when there are less sites in the system or when the clocks get larger (upper part of Figure 2). These effects tend to flatten after a certain point and, in client/server systems, this effect is not as noticeable. The performance of *Comb* improves when the local history grows and when the level of communications is higher (i.e., smaller values of $1/\lambda$). There are 6 significant third level interactions, 5 of them include the communication pattern and the system size appears in 2 of the interaction. The interactions between system size, communication pattern and local history size can be analyzed with the plots presented in Figure 3. In general, when the system has fewer sites, the value of ρ is smaller. However, with a client/server system, an increment in the size of the local histories overrides or at least reduces the effects that the system size has on ρ. The average effects of the interactions between system size, communication pattern and 1/λ are plotted in Figure 4. The system size affects ρ, but, in a client/server system, a high level of communications reduces its influence. Figure 5 presents the interactions

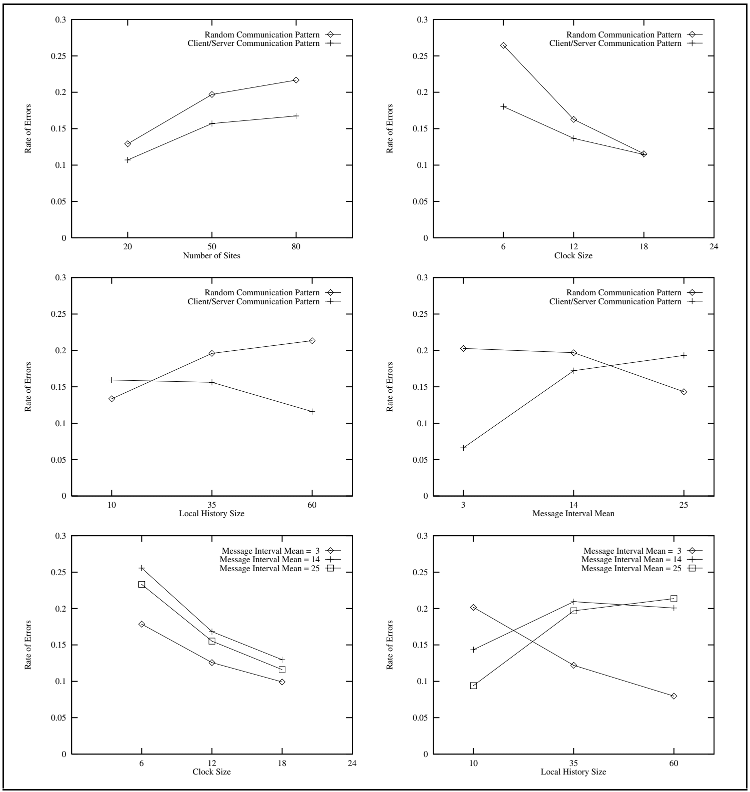

Figure 2. Average Effects of Second Level Interactions

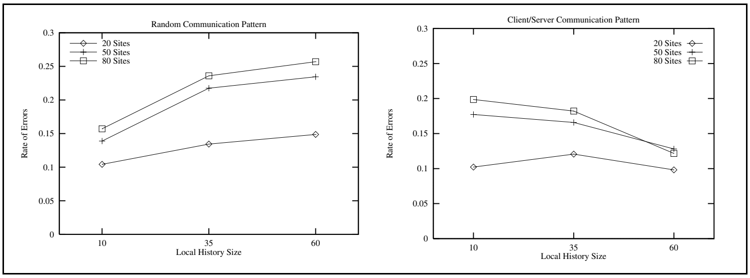

Figure 3. Interactions between Sites in the System, Communication Pattern and Local History Size

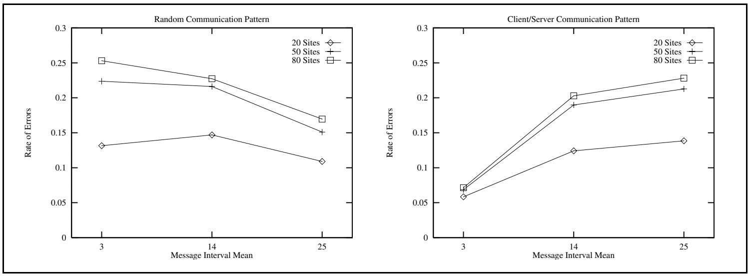

Figure 4. Interactions between Sites in the System, Communication Pattern and $1/\Lambda$

between clock size, communication pattern and local history size. Figure 6 presents the interactions between clock size, communication pattern and $1/\lambda$. In general, larger plausible clocks obtain better performance, but this effect is minimized in client/server systems when the local history size grows or when the level of communications increase. Figure 7 shows the interaction of clock size, local history size and $1/\lambda$. Larger logical clocks have better performance, and the values of ρ improve when the local histories are larger and the frequency of communications is higher. Figure 8 presents the interactions between communication pattern, local history size and $1/\lambda$. A client/server system has better results than a system with a random communication pattern. In particular, the minimum value of ρ is obtained by a client/server system, with a high rate of communications and large local histories.

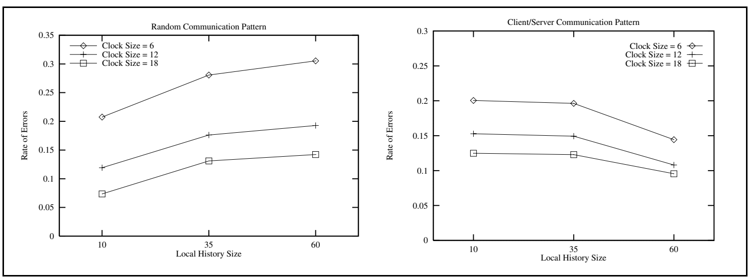

Figure 5. Interactions between Clock Size, Comm. Pattern and Local History Size

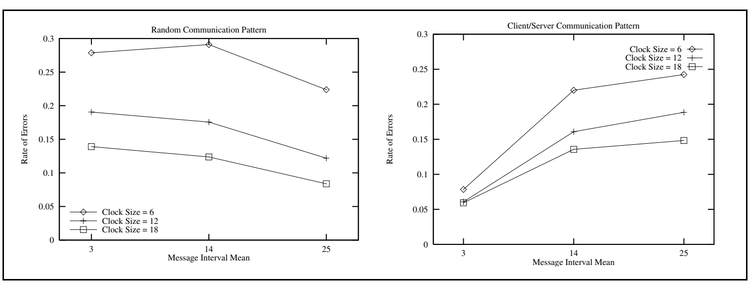

Figure 6. Interactions between Clock Size, Communication Pattern and $1/\Lambda$

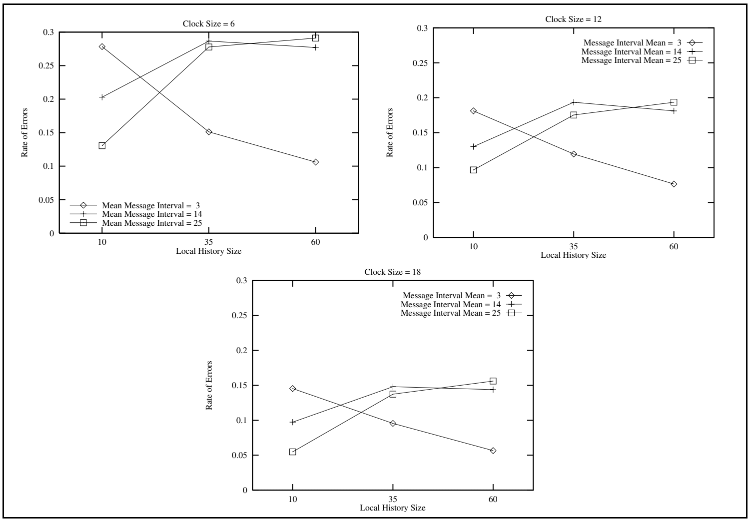

Figure 7. Interactions between Clock Size, Local History Size and 1/Λ

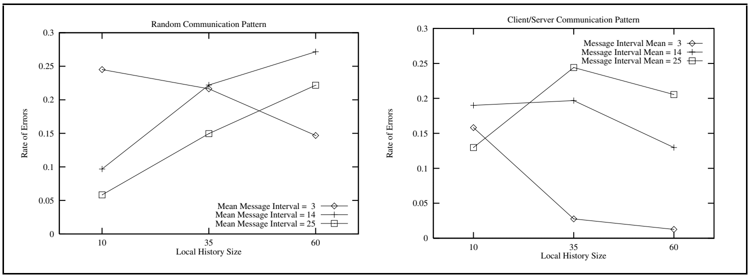

Figure 8. Interactions between Communication Pattern, Local History Size and 1/Λ

3 Experiment 2: Scalability of Plausible Clocks

The conditions for optimal performance of *Comb* identified in Section **2** are maintained as the system size is varied. We simulate a client/server system where each local history has 70 events with $1/\lambda = 3$. A 6 entry plausible clock *Comb* is used. Five system sizes are used (20, 40, 60, 80 and 100 sites) and 5 replications of the experiment are done. Figure 9 shows the average values of ρ at each system size. In a client/server system with a high level of communication and large local histories, the performance of *Comb* is not affected substantially by the number of sites in the system, i.e., this clock is scalable under the conditions mentioned.

4 Conclusions

The effects that system size, logical clock size, communication pattern, local history size and $1/\lambda$ have on the performance of plausible clocks were explored using computer simulations. Experiment 1 proved that all the main factors have a statistically significant effect on ρ. There are 15 significant interactions of factors and 12 of them include the communication patterns studied, while only 5 include the

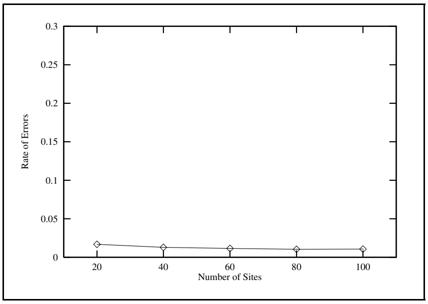

Figure 9. Scalability of Plausible Clocks under optimal conditions.

system size and in every case combined with the communication pattern. Based on the analysis of the second and third level interactions, we claim that even when the number of sites in the system and the size of the logical clock have a clear effect on the performance of plausible clocks, the interactions between the communication pattern and the other factors studied override their effects in many cases, especially with client/server systems. *Comb* performs better in client/server systems with high communications rates and large local histories. Experiment 2 shows that, under these conditions, the ability of *Comb* to accurately capture the causal relationships between a very high percentage (more than 98%) of the possible pairs of events is unaffected by the number of sites in the distributed system, even with a small number of entries.

References

1 B. Charron-Bost, "Concerning the size of logical clocks in Distributed Systems", Information Processing Letters 39, pp. 11-16. 1991.

2 C.J. Fidge, "Fundamentals of Distributed Systems Observation", IEEE Software, vol 13, No. 6, November 1996.

3 L. Lamport, "Time, clocks and the ordering of events in a Distributed System", Communications of the ACM, vol 21, pp. 558-564, July 1978.

4 F. Mattern, "Virtual Time and Global States in Distributed Systems", Conf. (Cosnard et al (eds)) Proc. Workshop on Parallel and Distributed Algorithms, Chateau de Bonas, Elsevier, North Holland, pp. 215-226. October 1988.

5 D. Montgomery, "Design and Analysis of Experiments", John Wiley and Sons, Fourth Edition. 1997.

6 R. Schwarz and F. Mattern, "Detecting causal relationships in distributed computations: in search of the holy grail", Distributed Computing, Vol. 7, 1994.

7 F. Torres-Rojas, "Scalable Approximations to Causality and Consistency of Distributed Objects", Ph.D. dissertation, College of Computing, Georgia Institute of Technology. July 1999.

8 F. Torres-Rojas and Mustaque Ahamad,"Plausible Clocks: Constant Size Logical Clocks for Distributed Systems", Distributed Computing, pp. 179-195, December 1999.

9 F. Torres-Rojas, "Performance Evaluation of Plausible Clocks", Technical Report, Computer Science Department, Costa Rica Institute of Technology, October 2000.