Clearing control policies for MAP inventory process with partially satisfied demand

Yonit Barron

Stochastic systems with clearing policies model many different applications, including certain bulk server queues and buffers in communication systems and manufacturing. Clearing models were first studied mainly for the stock level of cumulative inventory systems. Other applications include public transportation models, where arriving customers form a queue until a server, e.g., bus or train, arrives and removes all the customers present. In the queuing context, systems with work removals were attracted attention and more.

In this paper, we consider a production/clearing model in a random environment where a single machine produces a certain product into a buffer continuously. Customers (e.g., retailers) generate the demand for the product. The system is totally cleared at stationary renewal times, i.e., the model characterized by stochastic inputs (production) and outputs (retailers) and an additional total "clearings" at certain random times. After such a clearing, the system is assumed to ready to start working again immediately; thus, the associated content process is assumed to be semi-regenerative, starting almost-anew at level zero at every clearing time.

Three core models are considered: (i) Each time a pre-deterministic control limit level q is reached by the content level, the buffer is cleared. In the language of inventory theory, the inventory level is controlled under *continuous* review because it is observed continuously over time and the controller "sees" the specific level once is reached. We denote the first clearing time by T_q. (ii) The clearing times form a Poisson process which is independent of inventory process; thus the clearing time for the first cycle is an $exp(\zeta)$-distributed random time, say T_ζ. Alternatively, this *sporadic* review can reflect situations where management may have a little control or information on the process (in the case of disasters or obsolescence). (iii) We also deal with a combination of the above policies, $min(T_q , T_\zeta)$; clearing takes place at exponential time unless the level q is reached before, in which case the controller gets an emergency call, arrives immediately and clears the system.

The *continuous* models can be used to describe the control of epidemics, in which the quantity of interest is the number of susceptibles and the clearing corresponds to mass vaccination whenever the number of susceptibles exceed a specific number. Further use of these models includes inventory control. In many factories, due to physical capacity limits or safety and fire hazard regulations, inventory levels cannot exceed a specific level; thus the inventory is transferred to other locations.

Examples of *sporadic* models are disasters or obsolescence. In models with disasters all items stored are subject to external unexpected events that instantaneously bring the utility of all items on the shelf to 0, e.g., spoilage because of extreme weather conditions or a malfunction of a refrigerator that stores them. The family of models with obsolescence assumes that items become obsolete due to the introduction of a new product that replaces them in the market. Thus, the entire stock will simultaneously become obsolete at some (typically random) time. Examples are ample; in some industries like wireless chips the lifecycles of parts can be less than one year. Further examples range from avionics and military sectors, high tech products, communications, construction equipment, medical devices, transportations and supply chain networks.

In this paper we assume that the arrival times of demands form a Markov Additive Process (MAP) governed by a continuous-time Markov chain (environment), and the demand sizes are independent and have phase-type distributions depending on the type of arrival. The production process switches between predetermined rates which depend on the state of the environmental (note that the case of a fixed production rate is a special case of the latter). Negative inventory is not allowed- that is, there is no backlogging- so that level 0 is a reflecting barrier and we assume that the demand is completely or partially satisfied.

For the three policies we assume the following costs: (i) fixed cost for each clearing; (ii) holding cost for the inventory and (iii) penalty cost for unsatisfied demand. Our objective is to obtain tractable formulas for the appropriate cost functionals under all of these policies. In this paper we focus on the discounted cost criterion using a discount factor $\beta>0$; however, we also deal with the long run average cost criterion. To the best of our knowledge, such a general model has been never investigated in clearing/disaster literature and thus our research differs. Furthermore, while most of the papers involve analytic derivations of the quantities of interest, we emphasize our study using a more probabilistic approach; this enables a simple derivation of quantities of interest and obtains easy-to-implement explicit formulas for each policy. Our analysis is based on martingale techniques and on exit-time results for Markov-modulated fluid flow (MMFF) models. Starting from this analysis, one can minimize the cost of the system with respect to the parameters (e.g. q, ζ, the production rates, the demand sized and the costs) under each policy.

In this paper, we show (through numerical examples) that the discounted expected cost is a convex function of q (for T_q policy) and ζ (for T_ζ policy) and increases in production rates; we also showed the trade-off between the fixed and lost demand costs and the holding costs as a function of q and ζ; however, due to their different behavior on the total cost, their impact is opposite. Under the *sporadic* control policy, a comparative study between the T_ζ policy and the $min(T_q , T_\zeta)$ policy observes that for small values of ζ, the combined policy $min(T_q , T_\zeta)$ performs better; thus it is worthwhile the controller to determine the optimal q^* in order to reduce the total cost. However, for high values of ζ, the performance of both policies coincides and it would be easier for the controller to implement the *sporadic* policy only. Finally, the average case was investigated, and it was shown that both the discounted and the average cost functionals have a similar behavior.