Strategic interactions in a system of intensive care units

Anand Bharadwaj
The University of Melbourne

Introduction

In this project, building on the work of [2], we model a system of two interacting intensive care units (ICUs) as a continuous-time Markov chain (CTMC), where each unit may choose a threshold occupancy level above which all incoming arrivals will be diverted. [2] quantified the welfare loss (measured as the reduction in system throughput) due to selfish behaviour in this game using the price of anarchy (PoA) metric. We further develop their analysis by exploring the effect of heterogeneity in service rates on the equilibrium strategies of each ICU and the PoA of the system. This idea of service rate heterogeneity then motivates the development of a novel extension of the [2] model. In essence, we consider a Bayesian game where each ICU is privately informed about their service rate (type). his modified setup allows us to model the effect of inter-unit communication by comparing welfare in the private information case with the welfare when the players' types are common knowledge. We then proceed to compute the price of anarchy (PoA), price of stability (PoS) and price of communication (PoC), which are the respective measures of system inefficiency when the interacting ICUs do not communicate, communicate but do not cooperate, and communicate and cooperate.

Model

We build on the basic model of [2]. The game consists of two players, namely, the Nevill Hall (NH) and Royal Gwent (RG) hospitals, with respective bed capacities $c_{NH} = 8$ and $c_{RG} = 16$. We model the evolution of the system as a continuous-time Markov chain (CTMC), with the arrival rates for the two ICUs denoted by λ_{NH} and λ_{RG} respectively and the service rates denoted by μ_{NH} and μ_{RG}. With the assumption of no queueing, the state space of the model becomes the set $S = \{(u,v) \in \mathbb{Z} \mid 0 <= u <= c_{NH}, 0 <= v <= c_{RG}\}$, where u denotes the number of filled beds at NH and v denotes the number of occupied beds at RG. We now assume that each unit H chooses a diversion threshold K_H, such that $0 <= K_H <= c_H$, where if the occupancy at that unit exceeds this threshold, that unit will be in diversion. The arrival rates of patients are then dependent on the diversion status of the two players. More formally, we assume that the arrival rates of rates are of the form λ_H^r for $H \in \{NH, RG\}$ and $r \in \{(l,l),(l,h),(h,l),(h,h)\}$. Here, l denotes low demand (i.e. where a given ICU is not diverting patients), while h denotes high demand (i.e. where a given ICU is in diversion). The generator of the CTMC thus is:

$$q_{(u_i,v_i),(u_j,v_j)} = \begin{cases} u_i\mu_{NH} & \text{if } (u_i,v_i)-(u_j,v_j)=(1,0) \\ v_i\mu_{RG} & \text{if } (u_i,v_i)-(u_j,v_j)=(0,1) \\ \lambda_{NH}^{(l,l)} & \text{if } (u_i,v_i)-(u_j,v_j)=(-1,0) \text{ and } u_i < K_{NH} \text{ and } v_i < K_{RG} \\ \lambda_{RG}^{(l,l)} & \text{if } (u_i,v_i)-(u_j,v_j)=(0,-1) \text{ and } u_i < K_{NH} \text{ and } v_i < K_{RG} \\ \lambda_{NH}^{(l,h)} & \text{if } (u_i,v_i)-(u_j,v_j)=(-1,0) \text{ and } u_i < K_{NH} \text{ and } v_i \geq K_{RG} \\ \lambda_{RG}^{(l,h)} & \text{if } (u_i,v_i)-(u_j,v_j)=(0,-1) \text{ and } u_i < K_{NH} \text{ and } v_i \geq K_{RG} \\ \lambda_{NH}^{(h,l)} & \text{if } (u_i,v_i)-(u_j,v_j)=(-1,0) \text{ and } u_i \geq K_{NH} \text{ and } v_i < K_{RG} \\ \lambda_{RG}^{(h,l)} & \text{if } (u_i,v_i)-(u_j,v_j)=(0,-1) \text{ and } u_i \geq K_{NH} \text{ and } v_i < K_{RG} \\ \lambda_{NH}^{(h,h)} & \text{if } (u_i,v_i)-(u_j,v_j)=(-1,0) \text{ and } u_i \geq K_{NH} \text{ and } v_i \geq K_{RG} \\ \lambda_{RG}^{(h,h)} & \text{if } (u_i,v_i)-(u_j,v_j)=(0,-1) \text{ and } u_i \geq K_{NH} \text{ and } v_i \geq K_{RG} \\ 0 & \text{otherwise} \end{cases}$$

Finally, defining P^H to be the steady-state probability distribution, $U_H = \frac{\sum_{n=0}^{c_H} nP^H(n)}{c_H}$ to be the utilisation and $T_H = \mu_H \sum_{n=0}^{c_H} nP^H(n)$ to be the throughput for each hospital $H \in \{NH, RG\}$, each ICU aims to attain a utilisation rate that is as close as possible to a utilisation target t. More precisely, we assume that the ICUs simultaneously choose diversion thresholds K_H, $H \in \{NH, RG\}$, such that $0 <= K_H <= c_H$ to minimise a quadratic loss function $(U_H - t)^2$. Then, for a given choice of strategies, we define social welfare as the total system throughput (i.e. $T_{NH} + T_{RG}$).

Effect of Service Rate Variation on the PoA

In their paper, [2] considered two variants of their basic model. In Case 1 (strict diversion), if both ICUs declare that they are in diversion, the demand is lost from the system; however, in Case 2 (soft diversion), each ICU services its own demand if both units decide to divert incoming arrivals. Specifically, the state-dependent arrival rates in the two cases are the following:

Strict Diversion

$$\lambda_{RG}^{(r)} = \begin{cases} \lambda_{NH} & \text{if } r \in (l,l) \\ \lambda_{NH}+\lambda_{RG} & \text{if } r \in (l,h) \\ 0 & \text{if } r \in (h,l),(h,h) \end{cases}$$

$$\lambda_{RG}^{(r)} = \begin{cases} \lambda_{RG} & \text{if } r \in (l,l) \\ \lambda_{NH}+\lambda_{RG} & \text{if } r \in (h,l) \\ 0 & \text{if } r \in (l,h),(h,h) \end{cases}$$

Soft Diversion

$$\lambda_{NH}^{(r)} = \begin{cases} \lambda_{NH} & \text{if } r \in (l,l),(h,h) \\ \lambda_{NH}+\lambda_{RG} & \text{if } r \in (l,h) \\ 0 & \text{if } r \in (h,h) \end{cases}$$

$$\lambda_{RG}^{(r)} = \begin{cases} \lambda_{RG} & \text{if } r \in (l,l),(h,h) \\ \lambda_{NH}+\lambda_{RG} & \text{if } r \in (h,l) \\ 0 & \text{if } r \in (h,h) \end{cases}$$

We now perform some numerical experiments to test the effect of heterogeneity in service rates. The parameter values used in the original paper were $\mu_{NH} = 0.262$ and $\mu_{RG} = 0.198$. The effect of service rate variation was ascertained by choosing different values of μ_{NH} and μ_{RG} summing to $0.262 + 0.198 = 0.46$. To show this effect more clearly, it was decided to set the capacities and arrival rates at the two ICUs to be equal instead of using the original parameter values in [2]. We choose $c_{NH} = c_{RG} = 8$ and $\lambda_{NH} = \lambda_{RG} = 1.5$.

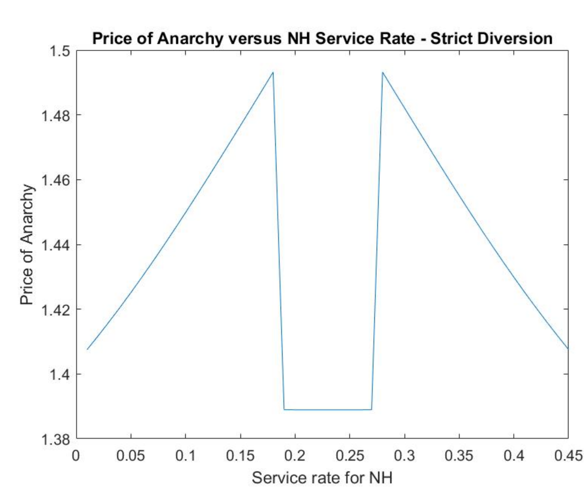

Fig. 1: Strict Diversion

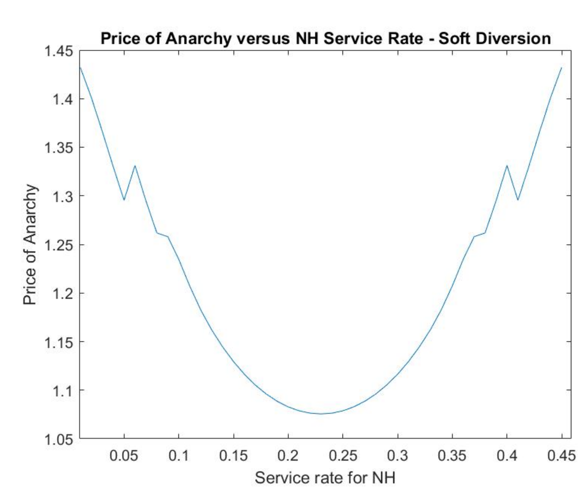

Fig. 2: Soft Diversion

In both cases, we see that differences in the service rates of the two ICU units have a significant effect on the price of anarchy of the system, even when the total service rate is constant. However, the effects are somewhat different in the two cases. In the soft diversion case, the minimal price of anarchy occurs when $\mu_{NH} = \mu_{RG} = 0.23$, with the PoA increasing sharply on both sides as the service rates become more asymmetric. In contrast, with strict diversion, although the minimum PoA once again occurs when the service rates are equal, asymmetry does not necessarily lead to a higher PoA. Thus, it certainly appears that skewed service rates can accentuate the effects of selfish incoordination in this game, although this effect is more pronounced when demand is conserved in the system.

Extension to Bayesian Model

The insight from the previous section that service rate heterogeneity can amplify the effects of competitive incoordination motivates the development of a new model incorporating private information and communication. We adapt the model of [1] and assume that the two ICUs are privately informed about their service rate μ_H (type). For analytical simplicity, suppose that each player's type is randomly and independently chosen at the start of the game by Nature from the type space $\{\mu_{NH}, \mu_{RG}\}$ with equal probability. Each player knows their own type, but not the type of the other ICU unit. Now, we compute social welfare (system throughput) under three scenarios:

Case 1: Communication and cooperation. Both ICUs reveal their service rates and then cooperatively choose strategies to maximise system throughput.

Case 2: Communication, but no cooperation. The ICUs reveal their types at the start of the game, but do not play cooperatively. We compute Nash equilibria and social welfare for the perfect-information game defined by each pair of types, and take the expectation to calculate the expected system throughput.

Case 3: No communication or cooperation. Here, the agents' types are private information, and their strategies involve specifying a diversion threshold for each possible value of their type. We then look for a Bayes Nash equilibrium and compute the welfare corresponding to such a strategy pair.

We define the price of anarchy (PoA) as the ratio of welfare in Case 1 and the lowest welfare at a Nash equilibrium in Case 3; the price of stability (PoS) as the ratio of welfare in Case 1 and optimal welfare at a Nash equilibrium in Case 2; and the price of communication (PoC) as the ratio of optimal social welfare in Cases 2 and 3.

Results

We compute the PoA, PoS and PoC for various service demands and utilisation targets (see [2] for the relevant notation). For brevity, only the strict diversion case is shown.

x	$t=0.2$	$t=0.5$	$t=0.8$
0	2.9322	1.3621	1.0000
1	4.3683	1.8064	1.1636
2	4.4936	2.1485	1.2629

x	$t=0.2$	$t=0.5$	$t=0.8$
0	3.2572	1.3127	1.0000
1	4.5814	1.9281	1.1510
2	4.7291	1.9783	1.1946

x	$t=0.2$	$t=0.5$	$t=0.8$
0	0.9002	1.0377	1.0000
1	0.9535	0.9369	1.0110
2	0.9502	1.0860	1.0571

Fig. 3: PoA, PoS and PoC for different target and demand rates

Examining the trends in the tables, we see that the PoA and PoS increase with demand for a fixed target, and decrease with the target for fixed demand. These findings largely mirror those in [2]. In essence, greater demand exacerbates the effect of strategic incoordination, while a higher value of the target tends to increase system throughput, thereby lowering the PoA and PoS. The most interesting finding is the PoC is sometimes less than 1, suggesting that the absence of communication can be a countervailing influence against the competitive incentives of the players, thereby increasing equilibrium welfare.

Acknowledgements

This project has taught me a great deal about how to apply mathematical tools and techniques to relevant practical problems and about the process of doing research. Given my interdisciplinary interests in mathematics, statistics and economics, I feel well placed to apply the skills that I have learnt in my future studies and research. I am immensely grateful to my supervisor Dr Mark Fackrell for stimulating my interest in this topic and for his guidance, support and insight over the course of this project.

References

[1] Mark Fackrell et al. "The value of communication and cooperation when servers are strategic". In: *ANZIAM Journal* 61 (2019), pp. 349–367.

[2] Vincent Knight, Izabela Komenda, and Jeff Griffiths. "Measuring the price of anarchy in critical care unit interactions". In: *Journal of the Operational Research Society* 68.6 (2017), pp. 630–642.