The Price of Anarchy, Stability, and Communication in Interacting Intensive Care Units

By Ashley Hanson   Supervised by Dr Mark Fackrell
The University of Melbourne

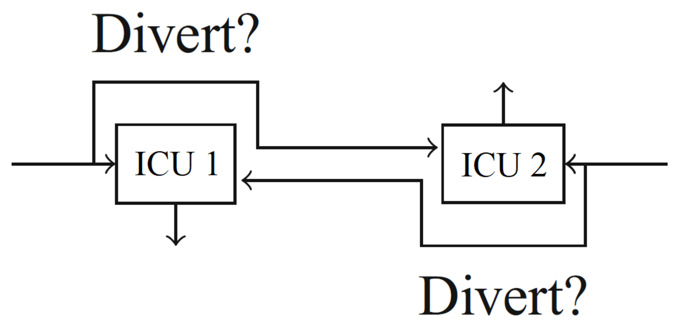

Figure 1: Diagrammatic representation of ICU interaction through patient diversion for a system of 2 ICUs. Patients arrive at one of the ICUs, spend some time in that ICU, and then leave. But in particular circumstances, the ICU can divert patient arrivals to the other ICU. [1]

Model

We consider a system of N interacting intensive care units (ICUs). Let n refer to ICU n, where $n \in \{1, 2, ..., N\}$. Each ICU has a bed capacity c_n, patient arrivals which follow a Poisson process with rate λ_n, and patient length of stay times exponentially distributed with mean $1/\mu_n$.

Each ICU chooses a capacity threshold K_n, where $0 \leq K_n \leq c_n$. v_n is the number of patients in ICU n at a given point in time. If $v_n < K_n$, ICU n is in low demand (l). If $v_n \geq K_n$, ICU n is in high demand (h). r_n is the status of ICU n, where $r_n \in \{l, h\}$. $\mathbf{r} = (r_1, r_2, ..., r_N)$ represents the status of the N ICUs. For example, $N = 5$ and $\mathbf{r} = (h, h, l, l, l)$. At a given time, L (H) is the set of ICUs with low (high) demand. Given a status $\mathbf{r}$, $\lambda_n^{\mathbf{r}}$ is the arrival rate of patients that ICU n admit to their ward. See Figure 2.

Strict Diversion (non-cooperative framework)

Arrival rates with this framework:

$$\lambda_n^{\mathbf{r}} = \begin{cases} \lambda_n + \dfrac{1}{|L|}\displaystyle\sum_{d \in H} \lambda_d, & \text{if } r_n = l \\ 0, & \text{if } r_n = h \end{cases}.$$

Patients are lost from the system when all ICUs are in high demand.

Soft Diversion (cooperative framework)

Similar set-up to strict diversion, except that ICUs must accept their own patients when all ICUs are in high demand. This can be viewed as a form of cooperation.

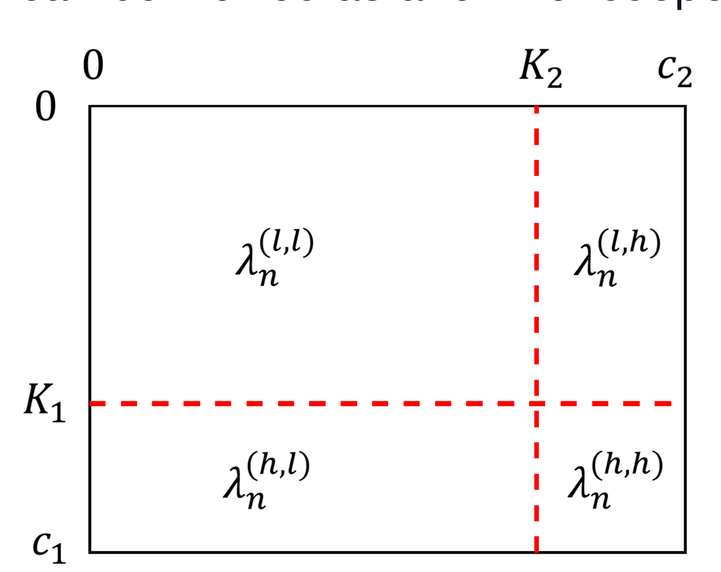

Figure 2: General arrival rates at each region for a system of 2 ICUs and given thresholds K_1 and K_2. [1]

Calculations

For every possible configuration of diversion thresholds $(K_1, K_2, ..., K_N)$, a continuous-time Markov chain (CTMC) is set-up. For each CTMC the stationary distribution is determined. This leads us to find the expected number of occupied beds in each ICU n. For a given configuration of thresholds, define the utilisation rate and throughput (respectively) in ICU n:

$$U_n = (1/c_n) \times \mathbb{E}(\text{occupancy in ICU } n), \quad T_n = \mu_n \times \mathbb{E}(\text{occupancy in ICU } n).$$

Each ICU calculates its optimal threshold (best response) to the thresholds of the other ICUs. The ICUs best response is found by solving the optimisation problem $\min\ (U_n - t)^2$ such that $0 \leq K_n \leq c_n$ and $K_n \in \mathbb{Z}$, where t is the bed utilisation target, set by central control. A Nash equilibrium is found at an intersection of best responses. See Figure 4. At a Nash equilibrium, unilateral deviation does not benefit any ICU.

For every possible configuration of thresholds, the sum of throughputs $T_1 + T_2 + ... + T_N$ is determined. T^* is the maximum sum of throughputs (independent of t). T^+ (T^-) is the largest (smallest) sum of throughputs when considering only threshold configurations which correspond to Nash equilibria. The Price of Anarchy [1], Price of Stability [2], and Price of Communication [3] (respectively):

$$\text{PoA} = T^*/T^-, \qquad \text{PoS} = T^*/T^+, \qquad \text{PoC} = T^+/T^-.$$

The PoS (PoA) is relevant when communication is (is not) possible between the N ICUs.

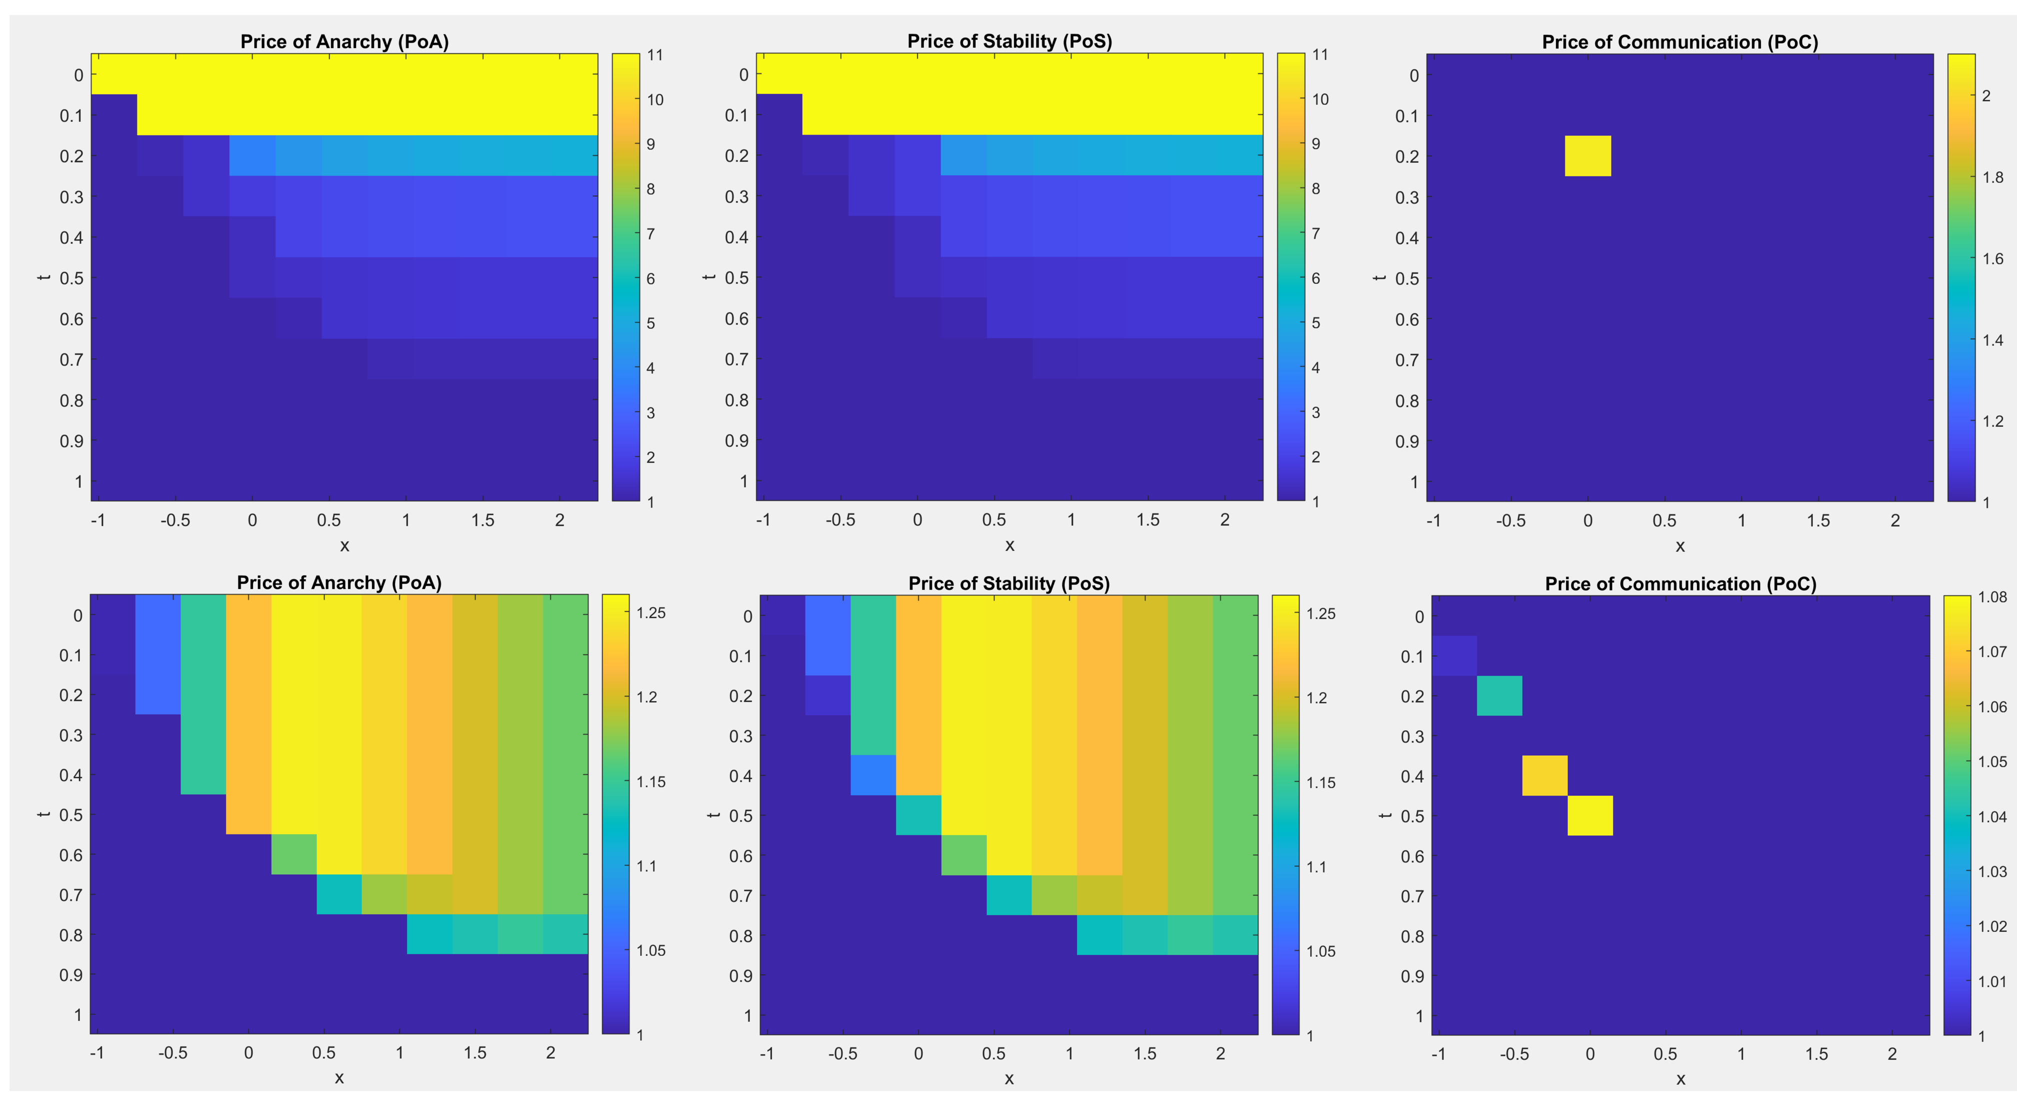

Figure 3: PoA, PoS and PoC plots for strict (top row) and soft (bottom row) diversion with 4 ICUs and bed capacities 2, 2, 3 and 3. x refers to a modified demand rate found by the transformation $\lambda_n \to \lambda_n(1 + x)$ for all $n \in \{1, 2, ..., N\}$. t is the bed utilisation target.

Experience

It was very rewarding to apply all the knowledge I had developed in my undergraduate studies to a practical application of mathematics. I learnt to appreciate that an initial idea for a research project can quickly lead into the exploration of many interesting extensions.

I would like to thank my supervisor, Dr Mark Fackrell, for the commitment he has shown towards my project. Mark's guidance provided a valuable insight into how research is conducted and presented. I always learnt something new during our regular discussions.

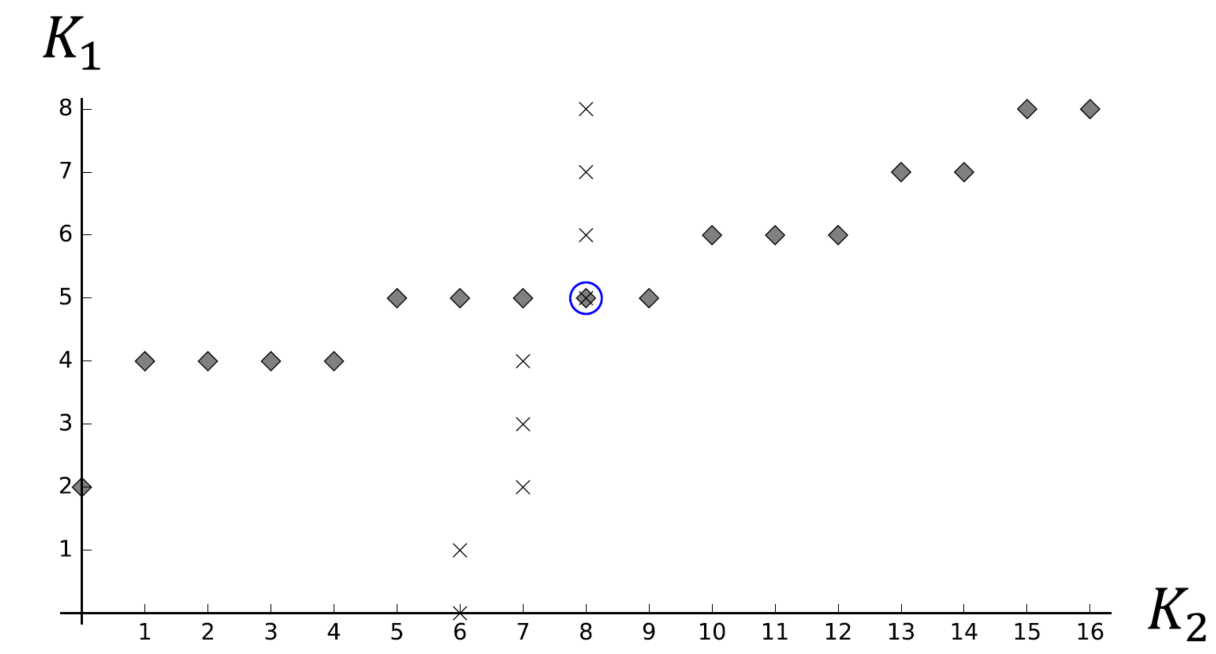

Figure 4: Example of best responses for a system of 2 ICUs. Diamonds (crosses) are the best responses for ICU 1 (2). [1]

Analysis

Since there are often multiple Nash equilibria, communication becomes more useful when the ICUs are working cooperatively (soft diversion) and there are more ICUs. The PoA and PoS are generally higher in the strict diversion framework (no cooperation).

With soft diversion, uncoordinated behaviour typically occurs for medium arrival rates. Alignment of interests always occurs for $t = 1$. But the optimal bed utilisation target occurs for the minimum value of t such that PoA $= 1$ (or PoS $= 1$). Increasing the number of ICUs tends to lower the optimal t value.

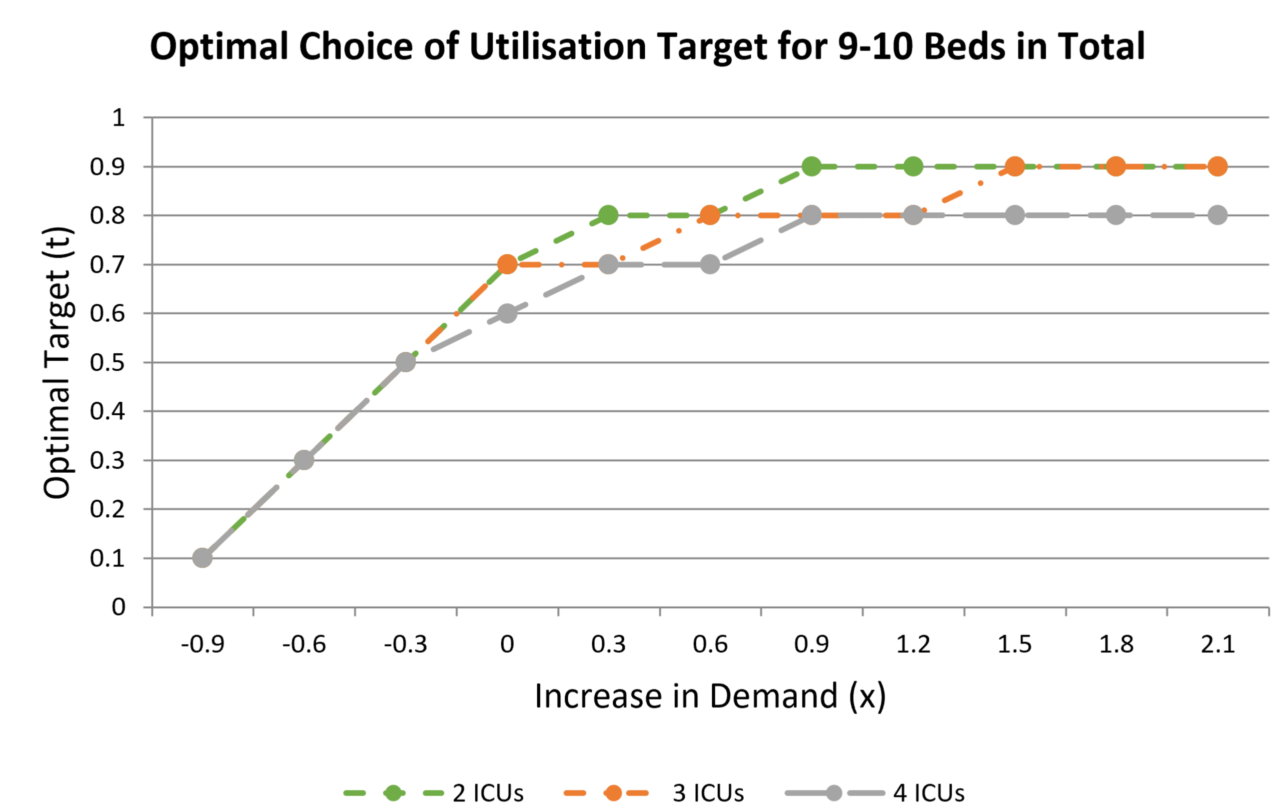

Figure 5: Optimal choice of utilisation target for 2 ICUs (with 3, 6 beds), 3 ICUs (with 2, 3, 4 beds) and 4 ICUs (with 2, 2, 3, 3 beds). The framework is strict diversion (and no communication).

References

1. Knight, V, Komenda, I & Griffiths, J, 2017, 'Measuring the price of anarchy in critical care unit interactions', *Journal of Operations Research Society*, vol. 68, no. 6, pp. 630-642.

2. Anshelevich, E, Dasgupta, A, Kleinberg, J, Tardos, E, Wexler, T & Roughgarden, T, 2004, 'The Price of Stability for Network Design with Fair Cost Allocation', *SIAM Journal on Computing*, vol. 38, no. 4, pp. 1602-1623.

3. Li, C, 2017, *Bayesian Game Theory with Application to Service Industry*, Master of Science Thesis, The University of Melbourne, Parkville, Vic.