# BCAM WORKSHOP ON
# BANDIT MODELS AND APPLICATIONS

## May 10, 2011

**John GITTINS**, Oxford University
*Forwards Induction, with particular reference to Pharmaceutical Research*

**Urtzi AYESTA**, BCAM and IKERBASQUE
*Properties of the Gittins index with application to optimal scheduling in an M/G/1 queue*

**Martin ERAUSQUIN**, EHU/UPV and BCAM
*Near-optimality of Whittle's index rule in a wireless downlink network scheduling problem*

**Peter JACKO**, BCAM
*Knapsack problem for perishable items, index-knapsack heuristic, and nearly-optimal revenue management*

bcam
basque center for applied mathematics

**BCAM** – Basque Center for Applied Mathematics
Bizkaia Technology Park, Building 500
E-48160 Derio - Basque Country - Spain
Tel. +34 946 567 842 | www.bcamath.org
( *matematika mugaz bestalde* )

**10:00-11:00 John Gittins**, Oxford University

*Forwards Induction, with particular reference to Pharmaceutical Research*

A forwards induction policy is a type of greedy algorithm for Markov decision processes. It is straightforward to implement and is optimal for a large class of models, especially in stochastic resource allocation. Instances of forwards induction will be discussed, with particular reference to its application in the selection of candidate drugs in pharmaceutical research.

Gittins, J. C., Charalambous, C., Chen, B., Oreskovich, A. M., Qu, S., and Yu, J. Y. (2011), OPRRA user-guide. Gittins, J. C., Glazebrook, K. D., and Weber, R. (2011), 'Multi-armed bandit allocation Indices', 2nd edition, Wiley. Qu, S., and Gittins, J.C. (2011), 'A forwards induction approach to candidate drug selection', to appear in Advances in Applied Probability, September 2011.

11:00-11:30 Coffee break

**11:30-12:00 Urtzi Ayesta**, BCAM and IKERBASQUE

*Properties of the Gittins index with application to optimal scheduling in an M/G/1 queue*

For an M/G/1 queue with the objective of minimizing the mean number of jobs in the system, the Gittins index rule is known to be optimal among the set of non-anticipating policies. We develop properties of the Gittins index. For a single-class queue it is known that when the service time distribution is of type Decreasing Hazard Rate (New Better than Used in Expectation), the Foreground-Background (First-Come-First-Served) discipline is optimal. By utilizing the Gittins index approach, we show that in fact, Foreground-Background and First-Come-First-Served are optimal if and only if the service time distribution is of type Decreasing Hazard Rate and New Better than Used in Expectation, respectively. For the multi-class case, where jobs of different classes have different service distributions, we obtain new results that characterize the optimal policy under various assumptions on the service time distributions. We also investigate distributions whose hazard rate and mean residual lifetime are not monotonic.

This is a joint work with S. Aalto (Aalto University, FI) and R. Righter (University of California at Berkeley, US).

**12:00-12:30 Martin Erausquin**, EHU/UPV and BCAM

*Near-optimality of Whittle's index rule in a wireless downlink network scheduling problem*

We present a wireless downlink network model, and its associated flow-level scheduling problem. Based on the Whittle's index we design a novel policy for this problem, which is an opportunistic scheduling rule. We show numerically that it outperforms many other well known scheduling policies. We prove that the proposed rule is stable in all stabilizable systems. We also analyze the model in the fluid-limit asymptotic regime and show that the proposed rule is optimal.

This is a joint work with U. Ayesta (BCAM, ES), P. Jacko (BCAM, ES), M. Jonckheere (CONICET, AR) and I.M. Verloop (BCAM, ES).

**12:30-13:00 Peter Jacko**, BCAM

*Knapsack problem for perishable items, index-knapsack heuristic, and nearly-optimal revenue management*

We introduce the knapsack problem for perishable items, which concerns the optimal dynamic allocation of a limited knapsack to a collection of perishable items. Such a problem is motivated by retail revenue management, where products often have an associated lifetime after which they cannot be sold. Due to its intractability, we derive a relaxation of the problem, which leads to its decomposition into single-product parametric subproblems. We employ the restless bandit indexation methodology in order to derive an index for the dynamic (single-)product promotion problem in closed form and to establish optimality of time-monotone promotion policy. We design a new, index-knapsack heuristic that uses the index values to define a proxy for the price of promotion and relies on solving an associated knapsack problem. The performance of the proposed policy is evaluated in a systematic computational study, showing an exceptional near-optimality and a significant superiority over the well-known index rule and over the benchmark earlier-deadline-first policy.