Supervisor: Dr. Nikitas J. Dimopoulos

# Abstract

With the availability of fast microprocessors and small-scale multiprocessors, inter-node communication has become an increasingly important factor that limits the performance of parallel computers. Essentially, message-passing parallel computers require extremely short communication latency such that message transmissions have minimal impact on the overall computation time. This thesis concentrates on issues regarding hardware communication latency in reconfigurable networks (optical or electronic), and software communication latency regardless of the type of network.

The first contribution of this thesis is the design and evaluation of two different categories of prediction techniques for message-passing systems. This thesis utilizes the communications locality property of message-passing parallel applications to devise a number of heuristics that can be used to predict the target of subsequent communication requests, and to predict the next consumable message at the receiving ends of communications.

Specifically, I propose two sets of predictors: C*ycle-based* predictors, which are purely dynamic predictors, and *Tag-based* predictors, which are static/dynamic predictors. The performance of the proposed predictors, specially Better-cycle2 and Tag-bettercycle2, are very good on the application benchmarks studied in this thesis. The proposed predictors could be easily implemented on the network interface due to their simple algorithms and low memory requirements.

As the second contribution of this thesis, I show that the majority of reconfiguration delays in reconfigurable networks can be hidden by using one of the proposed high hit ratio predictors. The proposed predictors can be used in establishing a communication pathway between a source and a destination before this pathway is to be used.

The third contribution of this thesis is the analysis of a broadcasting algorithm that utilizes latency hiding and reconfiguration in a single-hop reconfigurable network to speed the broadcasting operation. The analysis brings up closed formulations that yields the termination time.

The fourth contribution of this thesis is a new total exchange algorithm in single-hop reconfigurable networks. I conjecture that this algorithm ensures a better termination time than what can be achieved by either the direct or standard exchange algorithms.

The fifth contribution of this thesis is the use and evaluation of the proposed predictors to predict the next consumable message at the receiving ends of communications. This thesis contributes by claiming that these message predictors can be efficiently used to drain the network and cache the incoming messages even if the corresponding receive calls have not been posted yet. This way, there is no need to copy the early arriving messages into a temporary buffer. The performance of the proposed predictors, Single-cycle, Tag-cycle2 and Tag-bettercycle2, on the parallel applications are quite promising and suggest that prediction has the potential to eliminate most of the remaining message copies.