Doctoral Thesis Defense

Speaker:	Zunce Wei
Supervisors:	Drs. D. Goswami, H. F. Li
Examining Committee:	Drs. A. Boukerche, F. Khendek, S. P. Mudur, J. Rilling
Title:	Locality-Driven Checkpoint and Recovery
Date:	Wednesday, April 14, 2010
Time:	10:00 a.m.
Place:	EV 1.162

ABSTRACT

Checkpoint and recovery are important fault-tolerance techniques for distributed systems. Existing strategies incur either considerable checkpoint/logging cost or global-wise recovery effect, when applied to large-scale distributed systems. This thesis proposes a locality-driven strategy for such systems, with both affordable runtime overhead and controllable failure recoverability.

Messages establish dependencies between distributed processes, which can be handled either by coordinated checkpoints or via logging. Existing strategies enforce a uniform handling policy for all message dependencies, and hence gains advantage at one end but bears disadvantage at the other. In this thesis, a generic theory of Quasi-Atomic Recovery has been formulated to accommodate the two message handling methods, thus giving rise to a hybrid strategy with combined advantages from both ends.

A group checkpoint strategy has been proposed, based on the observation of message localization around 'locality regions' in distributed systems. Group-wise coordinated checkpoints can be created around such regions and only the few inter-region messages are logged subsequently. Runtime overhead is hence largely reduced and recovery spread is as localized as region-wise. Various checkpoint protocols and recovery techniques have been developed to provide trade-offs between flexibility and performance.

Distributed executions exhibit locality of message interactions. Such locality originates from resolving distributed dependency localization via message passing, and appears as a hierarchical 'region-transition' pattern. A bottom-up approach has been proposed to identify such 'locality regions' by composing 'locality intervals' from individual processes, based on their tight message coupling relations between each other. Experiments conducted on real-life applications have shown the existence of hierarchical locality regions and have justified the feasibility of this approach.

Performance of group checkpoint strategies is subject to their uses of locality regions. An abstract performance measure has been proposed to properly integrate both runtime overhead and failure recoverability in a region-wise manner. Good locality leads to good optimized performance measure, and such a locality pattern can serve as a good candidate for the set of optimal regions. Example checkpoint protocols have been proposed to efficiently identify optimized regions of a given pattern, with assistance of design-time or runtime knowledge.