



Contribution ID: 195

Type: **Plenary**

Concurrent Collections: Parallel Programming Without Worrying About Parallelism

Concurrent Collections is a different way of writing parallel applications. Its major contribution is to isolate the task of specifying the application semantics from any consideration of its parallel execution. This isolation makes it much easier for the domain-expert, the physicist for example, to specify the application. It also makes the task of the tuning-expert, mapping the application to a parallel platform, more effective. These experts may be the same person at different times, distinct people, or the tuning-expert may be automated.

The application is written in terms of high-level operators, the chunks of computation in the application. The specification indicates neither an explicitly serial ordering nor an explicitly parallel one. It simply specifies constraints on the order among these high-level operators. These constraints are based on the flow of data and control in the application.

Concurrent Collections supports (and we have built) a wide range of different mapping approaches. The grain, the distribution among processors and the scheduling within a processor are done statically or dynamically and by a person or automatically.

Primary author: KNOBE, Kathleen (Intel)

Presenter: KNOBE, Kathleen (Intel)

Track Classification: 1. Computing Technology