



Figure 1: KTA is a system for structuring and storing legal discourse. Each box above presents the portion of a distributed database that is stored on a lawsuit participant's computer. The court clerk's computer holds the primary record of the case, i.e., all the information that has been filed, most of which is publicly accessible (in this case, there is only one letter, an exhibit or a motion to dismiss, that is sealed and protected via encryption). Note that a p in front of a "document" indicates that it was filed by the plaintiff, a D by the defendant, a J by the judge.

The plaintiff's local database contains only a draft motion for summary judgment, currently hidden from all other parties in this action.

The dynamic trial/settlement model on the defendant's machine reflects his concern about the estimate cost of this case. This model is updated when events in the litigation change probabilities. For example, if an motion to exclude evidence succeeds, that may make it difficult for the plaintiff to prove a critical element of a claim. The probability of loss on that claim will be reduced, thus reducing the expected cost of the case. Because over 90% of cases settle before trial, this kind of information can be very valuable in a complex case.

The judge in this case has retained her draft denial of the defendant's motion to dismiss because she may want to use the same presumptions of fact in hearing a summary judgment motion. KTA assists judges by outlining cases and could quickly show which claims cannot succeed even under certain assumptions about the facts.