ference procedures to more general graphs. It operates in three stages. In the first stage, edges are added to $\mathcal{G}$ until it is *triangulated*.[8] Then, a tree is formed from the maximal cliques of the triangulated graph. Finally, an exact inference algorithm, which is equivalent to belief propagation, is performed on the tree of cliques. The triangulation step ensures that the clique tree satisfies the *running intersection property*, which requires that any variables shared by two clique nodes also be members of each clique on their unique connecting path. This property *must* be satisfied for local computations on the clique tree to produce globally consistent estimates. Unfortunately, for most interesting graphs with cycles, triangulation greatly increases the dimension d of the resulting clique variables. As tree–based inference procedures require $\mathcal{O}(Nd^3)$ operations, the junction tree algorithm is often no more tractable than direct matrix inversion.

Because direct methods are generally intractable, a wide range of iterative inference algorithms have been developed for graphs with cycles. One of the most popular is known as *loopy belief propagation* [79]. The loopy BP algorithm operates by iterating the parallel BP message passing equations, as derived in §2.3.1, on a graph with cycles. Because of the presence of loops, the messages lose their strict probabilistic interpretation, and the exact answer will not be achieved after any finite number of iterations. In some cases, the iterations may not converge at all. However, for many graphs, especially those arising in error–correcting codes [34, 51], loopy BP converges to beliefs which very closely approximate the true conditional marginal distributions [33, 52]. The standard BP derivation, as given in §2.3.1, provides no justification for loopy BP, other than the vague intuition that belief propagation should perform well for graphs whose cycles are "long enough."

For Gaussian Markov random fields, loopy BP has been analyzed in some detail. It was examined by Rusmevichientong and Van Roy [62] for the specific graphs used by turbo codes [51], and for general MRFs by Weiss and Freeman [80]. They show that when loopy BP does converge, it always calculates the correct conditional

[8] In a triangulated graph, every cycle of length four or greater has a chord, that is, an edge joining two nonconsecutive vertices of the cycle.