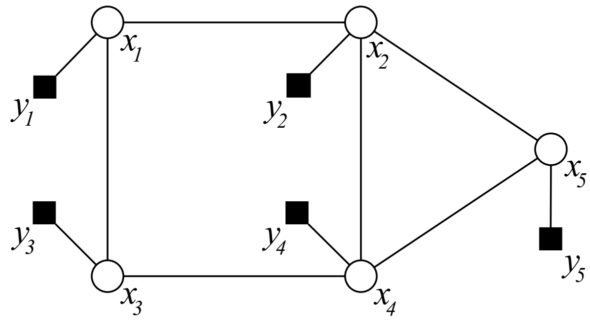

Figure 2-3: Graphical model with observation nodes explicitly shown. Circles represent hidden nodes x_s, while squares represent local observations y_s.

set $\{y_s\}_{s=1}^N$ of *local*, conditionally independent observations of the hidden variables $\{x_s\}_{s=1}^N$, so that $p(y \mid x) = \prod_{s=1}^N p(y_s \mid x_s)$. In this case, $C = \text{diag}(C_1, C_2, \ldots, C_N)$ and $R = \text{diag}(R_1, R_2, \ldots, R_N)$ must be block diagonal matrices, and for all $s \in \mathcal{V}$,

$$y_s = C_s x_s + v_s \tag{2.18}$$

where $v_s \sim \mathcal{N}(0, R_s)$. As shown in Figure 2-3, it is sometimes convenient to augment $\mathcal{G}$ by a set of *observed* nodes representing the local observations $\{y_s\}_{s=1}^N$.

Given this setup, we would like to compute the conditional *marginal* distributions $p(x_s \mid y) \sim \mathcal{N}\left(\widehat{x}_s, \widehat{P}_s\right)$ for all $s \in \mathcal{V}$. Note that each $\widehat{x}_s$ is simply a subvector of $\widehat{x}$, while each $\widehat{P}_s$ is a block diagonal element of $\widehat{P}$. We begin in §2.3.1 by deriving the general integral form of an algorithm, known as *belief propagation* (BP), which efficiently computes $p(x_s \mid y)$ for any tree–structured graph. In §2.3.2, we specialize the BP algorithm to the Gaussian case. Finally, in §2.3.3 we discuss existing exact and approximate inference algorithms for graphs with cycles.

2.3.1 Exact Inference for Tree–Structured Graphs

For graphs whose prior distribution is defined by a Markov chain, there exist efficient recursive algorithms for exactly computing the single–node conditional marginal distributions $p(x_s \mid y)$. If the variables are jointly Gaussian, one obtains a family of state–space smoothing algorithms which combine a standard Kalman filtering recur-

of P^{-1}, and can be easily accounted for by any of the inference algorithms discussed in this thesis.